

# FIA-Edit: Frequency-Interactive Attention for Efficient and High-Fidelity Inversion-Free Text-Guided Image Editing

Kaixiang Yang<sup>1, †</sup>, Boyang Shen<sup>1, †</sup>, Xin Li<sup>1</sup>, Yuchen Dai<sup>1</sup>, Yuxuan Luo<sup>1</sup>, Yueran Ma<sup>1</sup>  
Wei Fang<sup>2</sup>, Qiang Li<sup>1, \*</sup>, Zhiwei Wang<sup>1, \*</sup>

<sup>1</sup>Huazhong University of Science and Technology

<sup>2</sup>Wuhan United Imaging Healthcare Surgical Technology Co., Ltd  
{kxyang, boyangshen, liqiang8, zwwang}@hust.edu.cn

## Abstract

Text-guided image editing has advanced rapidly with the rise of diffusion models. While flow-based inversion-free methods offer high efficiency by avoiding latent inversion, they often fail to effectively integrate source information, leading to poor background preservation, spatial inconsistencies, and over-editing due to the lack of effective integration of source information. In this paper, we present **FIA-Edit**, a novel inversion-free framework that achieves high-fidelity and semantically precise edits through a **F**requency-**I**nteractive **A**ttention. Specifically, we design two key components: (1) a Frequency Representation Interaction (FRI) module that enhances cross-domain alignment by exchanging frequency components between source and target features within self-attention, and (2) a Feature Injection (FIJ) module that explicitly incorporates source-side queries, keys, values, and text embeddings into the target branch’s cross-attention to preserve structure and semantics. Comprehensive and extensive experiments demonstrate that FIA-Edit supports high-fidelity editing at low computational cost ( $\sim 6s$  per  $512 \times 512$  image on an RTX 4090) and consistently outperforms existing methods across diverse tasks in visual quality, background fidelity, and controllability. Furthermore, we are the first to extend text-guided image editing to clinical applications. By synthesizing anatomically coherent hemorrhage variations in surgical images, FIA-Edit opens new opportunities for medical data augmentation and delivers significant gains in downstream bleeding classification.

**Code** — <https://github.com/kk42yy/FIA-Edit>

## 1 Introduction

Text-guided image editing aims to modify an image according to a given textual description while preserving content unrelated to the edit. This task has witnessed significant progress in recent years, driven by the development of powerful generative models such as Denoising Diffusion Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020), Latent Diffusion Models (LDMs) (Rombach et al. 2022), and Diffusion Transformers (DiTs) (Peebles and Xie 2023; Labs 2024). These models

<sup>†</sup>Co-first authors.

<sup>\*</sup>Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: FIA-Edit is capable of handling a wide range of image editing tasks, including object modification, addition and removal, color transformation, and text replacement.

have been widely applied in real-world scenarios containing video editing (Geyer et al. 2023; Yoon et al. 2024; Gao et al. 2025; Yang et al. 2025; Cai et al. 2025; Li et al. 2025), visual effects production, and social media content creation.

Among these advances, tuning-free diffusion-based methods (e.g., DDIM-based sampling (Hertz et al. 2022; Tumanyan et al. 2023; Cao et al. 2023; Koo et al. 2024; Wu et al. 2024; Brooks, Holynski, and Efros 2023) and Rectified Flow (Avrahami et al. 2024; Deng et al. 2024; Rout et al. 2024; Wang et al. 2024; Hu et al. 2025)) have gained increasing attention. These approaches eliminate the need for per-instance fine-tuning, allowing for flexible and efficient zero-shot editing. Most existing methods fall into one of two categories, each presenting a fundamental trade-off




Figure 2: Overview of inversion-based and inversion-free image editing methods. (a) Inversion-based methods first invert the source image to noise, then edit from noise using the target prompt, often injecting source features during denoising. (b) Inversion-free methods bypass inversion by estimating velocity fields from noisy latent to source and noisy latent to target. Their difference defines the editing direction from source to target. However, this does not guarantee that the result preserves both background and target semantics (dark red ellipse). (c) Our method incorporates source-aware constraints (*i.e.*, FIA Constraint) during the computation of the noisy latent to target velocity field  $v_t^{tar}$ , effectively guiding the editing trajectory toward regions that preserve background fidelity while achieving semantic accuracy. Dashed arrows in (c) indicate FlowEdit, which lacks this guidance and fails to reach such optimal region.

between editing fidelity and computational efficiency.

The first and more established category adopts an *inversion-first* paradigm (Hertz et al. 2022; Tumanyan et al. 2023; Cao et al. 2023; Wu et al. 2024), where the source image is first projected into a latent prior distribution, typically Gaussian noise, using techniques such as DDIM inversion (Song, Meng, and Ermon 2020), Rectified Flow inversion (Lipman et al. 2022; Liu, Gong, and Liu 2022; Esser et al. 2024), or more advanced schemes (Miyake et al. 2025; Mokady et al. 2023). The editing process is then carried out in two stages: reconstructing the source image from the latent using the source prompt, and navigating from the same latent point to the target using the edited prompt. In this design, the prior distribution serves as a central “waypoint” that connects the source and target domains. To better align content and structure across the two branches, various feature interaction strategies including attention replacement and prompt injection have been proposed (see Fig. 2a). However, the inversion step is computationally intensive and significantly slows down the overall editing pipeline.

To improve efficiency, recent works have explored *inversion-free* approaches (Xu et al. 2023; Kulikov et al. 2024; Kim, Hong, and Ye 2025). These methods avoid explicit mapping to the latent prior and instead aim to directly construct the source-to-target trajectory. Since the source-to-target path is not directly accessible, these methods introduce virtual intermediate states by injecting noise into the source image. From this noisy reference point, two velocity fields are estimated: one pointing back to the source using source-prompt and the other toward the target using target-prompt. The vector difference between these two flows serves as an approximation of the editing direction, which implicitly encodes the source-to-target semantic transforma-

tion without ever performing an actual inversion. While this design enables much faster inference, it lacks explicit integration of source features during the editing process. As shown in Fig. 2b, this often results in poor content preservation in non-edited regions, leading to semantic drift, spatial inconsistency, and over-editing artifacts.

To address these limitations, we propose **FIA-Edit** (Fig. 2c), a novel inversion-free image editing framework that achieves high editing quality, strong background preservation, and fast generation. Instead of relying on an implicit feature transformation, FIA-Edit introduces an *explicit feature-level interaction* mechanism between the source and target representations throughout the editing trajectory. This design improves structural consistency and mitigates semantic drift in background regions.

The core of our method lies in a lightweight *Frequency-Interactive Attention* architecture. It contains two key modules: (1) The Frequency Representation Interaction (FRI) module, which fuses source and target features in the frequency domain within self-attention blocks, promoting cross-domain alignment without additional memory cost; (2) The Feature Injection (FIJ) module, which injects source-side queries, keys, values, and text embeddings into the cross-attention layers of the target branch, enhancing spatial and semantic consistency.

Our main contributions are summarized as follows:

- We propose **FIA-Edit**, an efficient and inversion-free image editing framework that achieves high-fidelity edits while explicitly preserving background structures.
- We introduce a unified Frequency-Interactive Attention mechanism, consisting of the FRI and FIJ modules, which enable explicit feature-level interaction between source and target to improve content alignment and structural consistency.
- We conduct extensive experiments on the PIE-Bench benchmark, and demonstrate that FIA-Edit achieves state-of-the-art performance across diverse editing tasks, with superior background preservation and semantic controllability.
- To the best of our knowledge, we are the first to apply general-purpose text-guided image editing methods to clinical images. Moving beyond artistic manipulation, FIA-Edit enables anatomically meaningful modifications, such as adjusting bleeding severity in surgical scenes. This opens up new opportunities for using image editing tools in medical data augmentation and downstream clinical tasks.

## 2 Related Works

### 2.1 Inversion-based Methods

Inversion-based image editing methods typically rely on first inverting the source image back into noise through an inversion process, and then performing editing conditioned on target prompts. Broadly, existing work in this category can be divided into three main directions:

**Improvements to the inversion process.** A number of approaches aim to enhance the quality, stability, and ac-

curacy of the inversion process. In the DDIM-based setting, Null-Text Inversion (Mokady et al. 2023) demonstrates that effective inversion can be achieved without any textual prompt to suppress irrelevant content during reconstruction. For rectified flow-based approaches, methods such as RF-Inv (Rout et al. 2024) and FireFlow (Deng et al. 2024) focus on refining the inversion process to reduce reconstruction artifacts. Other methods, including Direct Inversion (Ju et al. 2023) and DNAEdit (Xie et al. 2025), aim to minimize the discrepancy between the actual and ideal inversion outputs, thereby boosting the quality of reconstruction.

**Feature injection during editing.** To improve controllability and fidelity, many tuning-free approaches incorporate feature injection mechanisms during the generation process. Prompt-to-Prompt (P2P) (Hertz et al. 2022) explores direct feature replacement within cross-attention layers, whereas Plug-and-Play (PnP) (Tumanyan et al. 2023) injects source features between residual and attention blocks to enhance background preservation. FTEdit (Xu et al. 2024) introduces semantic feature replacement within adaptive layer normalization modules, enabling more precise and disentangled control over the generated content.

**Frequency-aware latent processing.** Recent studies have increasingly explored the integration of frequency operations. FlexiEdit (Koo et al. 2024) suppresses high-frequency components in DDIM latents associated with editable regions, enabling non-rigid edits. FDS (Ren et al. 2025) adopts wavelet decomposition to adaptively select frequency bands according to the editing task, enabling fine-grained control. Despite their effectiveness, these methods require an inversion process, increasing editing time and involving multiple task-specific hyperparameters.

## 2.2 Inversion-free Methods

To reduce the computational overhead of image editing, a natural direction is to eliminate the time-consuming inversion process. Several recent approaches have explored this idea by bypassing the explicit mapping of source images into the noise space. InfEdit (Xu et al. 2023) introduces the Denoising Diffusion Consistent Model (DDCM), which adopts a multi-step consistency sampling strategy that enables image editing without requiring explicit inversion. FlowEdit (Kulikov et al. 2024) further proposes an inversion-free framework by leveraging the velocity field to construct a direct trajectory from the source image to the edited target, avoiding inversion to Gaussian noise. Building on this, FlowAlign (Kim, Hong, and Ye 2025) introduces trajectory regularization to achieve more consistent and controllable text-driven editing within this inversion-free paradigm.

While these methods significantly reduce editing time, they often underutilize source image features, leading to insufficient background preservation and noticeable inconsistencies in non-edited regions. In contrast, we propose FIA-Edit, which enhances inversion-free editing by introducing frequency-aware feature interaction directly within the velocity field. This design effectively retains high-fidelity background information while ensuring both editing quality and runtime efficiency.

## 3 Method

Fig. 3 illustrates the overall architecture of our method. As shown in Fig. 3a, our approach builds upon the inversion-free FlowEdit paradigm (Kulikov et al. 2024) as backbone. On this basis, we enhance the integration of relevant information from the source velocity field into the target velocity field, thereby improving both background preservation and semantic consistency during editing. Then, we detail the backbone and the proposed FIA Constraint, which consists of two sub-modules: Frequency Representation Interaction (FRI) and Feature Injection (FIJ).

### 3.1 Backbone

Our backbone is built upon Rectified Flow, which enables direct progression from the source domain to the target domain by estimating the difference between their respective velocity fields.

At a given discrete editing time step  $\sigma_t$  with index  $t$ , a linear interpolation is first employed between the source image  $\mathbf{X}^{src}$  and Gaussian noise  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ , following the Rectified Flow formulation:

$$\mathbf{x}_t^{src} = (1 - \sigma_t) \cdot \mathbf{X}^{src} + \sigma_t \cdot \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (1)$$

Using the source prompt  $\mathcal{P}^{src}$ , corresponding source velocity  $v_\theta(\mathbf{x}_t^{src}, \mathcal{P}^{src}, t)$  is computed.

Next, to obtain the corresponding target velocity field, the current editing feature  $\mathbf{x}_t^{FE}$  and the additive relationship of vectors  $\mathbf{x}^{FE} = \mathbf{X}^{src} + \mathbf{x}^{tar} - \mathbf{x}^{src}$  are leveraged. Accordingly, the target representation at step  $t$  can be expressed as:

$$\mathbf{x}_t^{tar} = \mathbf{x}_t^{FE} + \mathbf{x}_t^{src} - \mathbf{X}^{src}. \quad (2)$$

Target velocity  $v_\theta(\mathbf{x}_t^{tar}, \mathcal{P}^{tar}, t)$  is obtained with target prompt  $\mathcal{P}^{tar}$ . It is worth noting that during the editing process,  $\mathbf{x}_t^{src}$  and  $\mathbf{x}_t^{tar}$  progressively move toward the source and target domains, respectively. At the initial step, the two are identical.

The direction is then determined by velocity difference:

$$v_t^\Delta = v_\theta(\mathbf{x}_t^{tar}, \mathcal{P}^{tar}, t) - v_\theta(\mathbf{x}_t^{src}, \mathcal{P}^{src}, t). \quad (3)$$

As evident, the backbone relies solely on the implicit interaction between  $\mathbf{x}_t^{src}$  and  $\mathbf{x}_t^{tar}$ , without any explicit guidance from the source image features. This often causes the editing process to deviate toward the target domain too freely, leading to weak constraints from  $\mathbf{X}^{src}$  and suboptimal results, especially in preserving source-relevant content.

Finally, the editing feature is updated iteratively according to the rectified flow stepping rule:

$$\mathbf{x}_{t-1}^{FE} = \mathbf{x}_t^{FE} + (\sigma_{t-1} - \sigma_t) \cdot v_t^\Delta. \quad (4)$$

After completing all time steps, the edited image is synthesized from the final state  $\mathbf{x}_0^{FE}$ .

### 3.2 FIA Constraint

As shown in Fig. 3b, to better preserve background content and ensure semantic alignment, we explicitly incorporate source features into the computation of target velocity fields

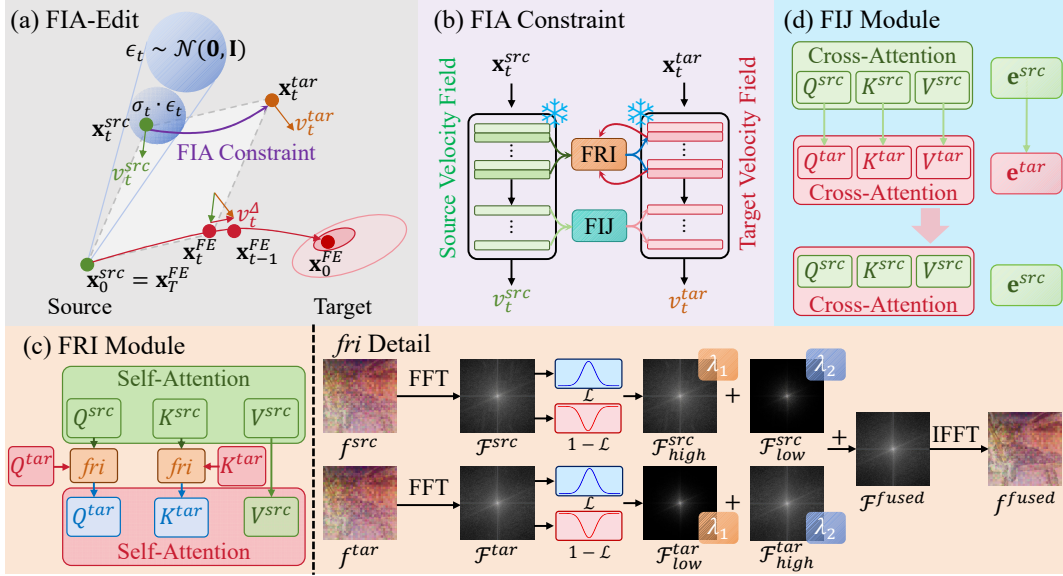


Figure 3: Details of our framework. (a) Overview of FIA-Edit. During the computation of source and target velocity fields, we introduce the FIA constraint to enable interaction between source and target features. (b) FIA constraint. (c) Frequency Representation Interaction (FRI). FRI is integrated into the self-attention layers. Both source and target Q/K features are fused in the frequency domain, and the fused output replaces the target Q/K. The right side shows the detailed structure of the frequency-domain fusion module *fri*. (d) Feature Injection (FIJ). FIJ is used in the cross-attention layers in the latter of DiT.

through a proposed **FIA Constraint**. This constraint consists of two key components: (1) the FRI module, which operates within *self-attention* to enable frequency-domain interaction between source and target features, and (2) the FIJ module, applied in *cross-attention* to inject source features directly.

We denote the attention features extracted from  $v_\theta(x_t^{src}, \mathcal{P}^{src}, t)$  and  $v_\theta(x_t^{tar}, \mathcal{P}^{tar}, t)$  as  $\{f_t^{src}\}$  and  $\{f_t^{tar}\}$ , respectively. Then, Eq. 3 can be reformulated as:

$$v_t^\Delta = v_\theta(x_t^{tar}, \mathcal{P}^{tar}, t, \text{FIA}(\{f_t^{src}\}, \{f_t^{tar}\})) - v_\theta(x_t^{src}, \mathcal{P}^{src}, t) \quad (5)$$

### 3.3 Frequency Representation Interaction

To preserve structural fidelity while enabling meaningful semantic transformation, we introduce the Frequency Representation Interaction (FRI) module, which performs cross-domain feature fusion in the frequency domain, as illustrated in Fig. 3c.

FRI is motivated by the observation that structure and semantics are more naturally disentangled in the frequency space: low-frequency components primarily encode coarse spatial layouts and background structures, while high-frequency components capture fine-grained textures and semantic details. Based on this insight, we propose a cross-domain frequency fusion strategy that enhances the high-frequency components of the source and the low-frequency components of the target, while suppressing low-frequency content in the source and high-frequency signals in the target. This selective fusion effectively leverages source information through frequency-domain interaction.

We first compute the velocity fields of  $x_t^{src}$  and  $x_t^{tar}$ , and extract the intermediate features  $f_t^{src} \in \mathbb{R}^{C \times H \times W}$  and  $f_t^{tar} \in \mathbb{R}^{C \times H \times W}$  from DiT. A 2D Fast Fourier Transform (FFT)  $\text{FFT}(\cdot)$  is then applied:

$$\mathcal{F}^{src} = \text{FFT}(f_t^{src}), \quad \mathcal{F}^{tar} = \text{FFT}(f_t^{tar}). \quad (6)$$

Using a Gaussian low-pass filter  $\mathcal{L}$ , we decompose each into high- and low-frequency components:

$$\mathcal{F}_{high}^{src} = \mathcal{F}^{src} \cdot (1 - \mathcal{L}), \quad \mathcal{F}_{low}^{src} = \mathcal{F}^{src} \cdot \mathcal{L}, \quad (7)$$

$$\mathcal{F}_{high}^{tar} = \mathcal{F}^{tar} \cdot (1 - \mathcal{L}), \quad \mathcal{F}_{low}^{tar} = \mathcal{F}^{tar} \cdot \mathcal{L}. \quad (8)$$

The fused spectrum is computed by applying cross-weighted fusion:

$$\mathcal{F}^{fused} = \lambda_1 \cdot (\mathcal{F}_{high}^{src} + \mathcal{F}_{low}^{tar}) + \lambda_2 \cdot (\mathcal{F}_{low}^{src} + \mathcal{F}_{high}^{tar}), \quad (9)$$

where  $\lambda_1 = 0.8$ ,  $\lambda_2 = 0.2$  are weighting coefficients, emphasizing structure and semantics from source image while suppressing conflicting signals.

The fused feature is then reconstructed via inverse FFT:

$$f^{fused} = \text{IFFT}(\mathcal{F}^{fused}), \quad (10)$$

which is then injected into the self-attention layers, guiding the target velocity updates. By aligning complementary information across domains, FRI allows our model to perform semantically accurate edits while preserving the source’s visual structure, improving both realism and control.

### 3.4 Feature Injection

To further improve background preservation, we draw inspiration from inversion-based methods (Tumanyan et al. 2023;

Cao et al. 2023) that inject source features to maintain spatial consistency and fine-grained control. We propose a Feature Injection (FIJ) module (Fig. 3d) to explicitly introduce source features into the editing process.

Unlike prior works that inject only Q or K across the entire network, FIJ operates within the cross-attention layers of the later DiT blocks (*i.e.*, layers 13  $\sim$  23). Specifically, we inject source-side query ( $Q^{src}$ ), key ( $K^{src}$ ), value ( $V^{src}$ ), and text embedding ( $e^{src}$ ) into the target attention computation:

$$Q^{tar} \leftarrow Q^{src}, K^{tar} \leftarrow K^{src}, V^{tar} \leftarrow V^{src}, e^{tar} \leftarrow e^{src}. \quad (11)$$

The injection is applied only during the early generation steps, when  $x_t^{tar}$  and  $x_t^{src}$  are still similar. This early fusion allows  $x_t^{tar}$  to absorb source information smoothly under  $\mathcal{P}^{tar}$ 's guidance, enabling coherent edits and avoiding the abrupt changes seen in other methods. Additionally, FIJ stabilizes the semantic alignment between source and target prompts.

## 4 Experiment

### 4.1 Experiment Design

**Dataset and Baselines.** To thoroughly evaluate the effectiveness of our proposed method, we conduct experiments on the PIE-Bench (Ju et al. 2023) benchmark, which comprises 700 image–prompt pairs spanning 10 diverse editing categories. We compare our approach with a comprehensive set of baselines, including: LDM-based methods (P2P (Hertz et al. 2022), PnP (Tumanyan et al. 2023), MasaCtrl (Cao et al. 2023), FlexiEdit (Koo et al. 2024), and FreeDiff (Wu et al. 2024)), FLUX-based approaches (RF-Inv (Rout et al. 2024), StableFlow (Avrahami et al. 2024), RF-Edit (Wang et al. 2024), and DCEdit (Hu et al. 2025)), and DiT-based methods (FTEdit (Xu et al. 2024), FlowEdit (Kulikov et al. 2024), and DNAEdit (Xie et al. 2025)). All models are evaluated using their publicly released implementations and default settings to ensure fair and consistent comparison.

**Metrics.** To comprehensively evaluate both editing performance and background preservation, we adopt six complementary metrics. Structure Distance (Tumanyan et al. 2022) quantifies structural consistency between the edited and original images. PSNR, LPIPS (Zhang et al. 2018), MSE, and SSIM (Wang et al. 2004) jointly assess content fidelity in unedited regions. For text-image alignment, we calculate CLIP similarity (Radford et al. 2021) over both the entire image and the specifically edited regions. Note that region masks provided by the dataset are used solely for evaluation purposes to isolate the edited areas.

**Implementation Details.** FIA-Edit is based on the SD3.5-Medium (Stability AI 2025) model with 50 sampling steps. Velocity fields are computed using CFG scales of 3.5 (source) and 13.5 (target). The FRI module is applied across all 50 steps, while the FIJ module is activated only during the first 27 steps to constrain early velocity field. During the  $x_t^{FE}$  stepping process, we add the reused interpolation noise  $\epsilon_t$  with a scaling factor of  $\sigma_t$ . Details of the Gaussian low-pass filter are provided in the Appendix. All experiments are

run on a single NVIDIA RTX 4090 GPU. For comparison, we reproduce results of open-source baselines using their official code and default settings (see Appendix for details), while results of non-released methods are directly reported from their original papers.

### 4.2 Comparisons with Other Editing Methods

**Quantitative Comparison.** As shown in Table 1, we conduct comprehensive evaluations on PIE-Bench across representative LDM-, FLUX-, and DiT-based methods. FIA-Edit achieves the best performance in background preservation while also demonstrating strong semantic alignment. Compared to the inversion-free baseline FlowEdit, our method preserves background details more accurately, highlighting the effectiveness of integrating source-target feature interactions during velocity field computation. Among inversion-based methods, P2P maintains relatively good background consistency but suffers from weak prompt alignment, suggesting that it overly retains source content. Overall, FIA-Edit delivers both superior background fidelity and precise semantic edits, leading to the best average ranking across all metrics, which validates the effectiveness of our approach.

**Qualitative Comparison.** Visual results are shown in Fig. 4, covering content alterations, object addition, and pose change. Our method produces high-quality, semantically accurate edits while maintaining background integrity. In contrast, other methods either fail to achieve the intended semantic change or suffer from noticeable background distortion (*e.g.*, FlowEdit), clearly demonstrating the superiority of our approach.

**GPU Memory and Runtime.** We report GPU memory and runtime of open-source methods in Table 2, measured on a single RTX 4090 for fair comparison. Due to high memory demands, RF-Inv, StableFlow, and RF-Edit are tested on an A100 (80GB). Runtime is averaged over 10 samples (image size:  $512 \times 512$ ), covering the full pipeline from loading to saving. As shown, inversion-free methods are notably faster. Compared to FlowEdit, our method incurs slightly more runtime due to feature interaction, while keeping memory usage comparable. Overall, FIA-Edit strikes a strong balance between quality and efficiency.

### 4.3 Ablation Study

We conduct ablations to evaluate the effectiveness of Frequency Representation Interaction (FRI) and Feature Injection (FIJ), as shown in Table 3. In FRI, *freq* denotes our frequency-domain fusion, while *add* refers to simple feature addition. Row 1 vs. Row 2 shows that FIJ greatly improves background preservation by injecting source features during velocity estimation. Row 2 vs. Row 4 indicates that combining FRI on top of FIJ further enhances semantic alignment while retaining strong background consistency, demonstrating the value of frequency-guided interaction. Row 3 vs. Row 4 shows that our structured *freq* design outperforms naive *add*. Note that CLIP score alone does not reflect editing quality, ideal edits require both high semantic fidelity and accurate background preservation. See Appendix for visual examples illustrating this balance.

Method	Model	Structure	Background Preservation				CLIP Similarity		Rank
		Distance $\times 10^3$ ↓	PSNR ↑	LPIPS $\times 10^3$ ↓	MSE $\times 10^4$ ↓	SSIM $\times 10^2$ ↑	Whole ↑	Edited ↑	Avg. ↓
P2P	SD1.4	11.65 <sup>2</sup>	27.22 <sup>2</sup>	54.55 <sup>1</sup>	32.86 <sup>3</sup>	84.76	25.02	22.10	5.3
PnP	SD1.5	24.29	22.46	106.06	80.45	79.68	25.41	22.62	9.6
MasaCtrl	SD1.4	24.70	22.64	87.94	81.09	81.33	24.38	21.35	10.7
FlexiEdit	SD1.4	22.13	25.74	80.45	58.45	82.62	25.15	22.87 <sup>2</sup>	6.0
FreeDiff	SD1.5	18.70	24.73	89.76	55.32	81.68	25.03	22.12	7.9
RF-Inv	FLUX	48.76	19.51	195.85	155.74	68.95	25.11	22.50	11.6
StableFlow	FLUX	19.24	23.04	76.94	84.85	87.22	24.30	21.28	8.9
RF-Edit	FLUX	27.70	23.22	131.18	75.00	81.44	25.22	22.40	9.4
DCEdit	FLUX	22.36	25.41	94.17	48.09	85.60	25.47	22.71	6.1
FTEdit	SD3.5	18.17	26.62	80.55	40.24	91.50 <sup>1</sup>	25.74 <sup>3</sup>	22.27	4.4 <sup>3</sup>
FlowEdit	SD3.5	23.62	23.21	93.81	69.95	85.09	26.78 <sup>1</sup>	23.73 <sup>1</sup>	6.1
DNAEdit	SD3.5	14.19 <sup>3</sup>	26.66 <sup>3</sup>	74.57 <sup>3</sup>	32.76 <sup>2</sup>	88.63 <sup>3</sup>	25.63	22.71	3.1 <sup>2</sup>
Ours	SD3.5	10.34 <sup>1</sup>	27.32 <sup>1</sup>	55.02 <sup>2</sup>	28.66 <sup>1</sup>	89.21 <sup>2</sup>	25.89 <sup>2</sup>	22.82 <sup>3</sup>	1.7 <sup>1</sup>

Table 1: Quantitative comparison on PIE-Bench. Rank denotes the average ranking across all evaluation metrics. Our method achieves strong performance in both background preservation and semantic alignment, yielding the best average rank. Super-scripts <sup>1</sup>, <sup>2</sup>, and <sup>3</sup> denote the best, second-best, and third-best performance, respectively.

Method	P2P	PnP	MasaCtrl	FlexiEdit	FreeDiff	<i>RF-Inv</i>	<i>StableFlow</i>	<i>RF-Edit</i>	FlowEdit	Ours
GPU(GB)	10.95	8.99	11.42	18.73	6.08	<i>69.22</i>	<i>35.39</i>	<i>32.91</i>	17.93	17.93
Time(s)	34.84	18.09	21.71	38.97	17.41	<i>76.74</i>	<i>26.07</i>	<i>34.51</i>	3.49	6.30

Table 2: Memory and runtime comparison. RF-Inv, StableFlow, and RF-Edit were run on an A100 80GB GPU, while all other methods were tested on a single RTX 4090. Our approach achieves a favorable balance between speed and editing quality.

Module		Structure	Background Preservation				CLIP Similarity	
FIJ	FRI	Distance $\times 10^3$ ↓	PSNR ↑	LPIPS $\times 10^3$ ↓	MSE $\times 10^4$ ↓	SSIM $\times 10^2$ ↑	Whole ↑	Edited ↑
×	×	23.62	23.21	93.81	69.95	85.09	26.78	23.73
✓	×	14.89	25.59	70.18	41.74	87.51	26.30	23.12
✓	<i>add</i>	16.50	25.93	85.44	38.72	86.51	26.05	22.68
✓	<i>freq</i>	10.34	27.32	55.02	28.66	89.21	25.89	22.82

Table 3: Ablation study on key components of FIA-Edit. FRI and FIJ denote the proposed Frequency Representation Interaction and Feature Injection modules, respectively. Within FRI, *freq* refers to our frequency-domain fusion design, while *add* denotes direct addition of source and target features.

Method	ConvNeXt-T	Aug	PnP	MasaCtrl	FlexiEdit	FreeDiff	FlowEdit	Ours
AUC (%)	81.54	82.10	81.98	<u>84.22</u>	82.05	82.25	83.83	<b>85.05</b>
PR-AUC (%)	38.66	38.81	38.53	38.82	37.97	38.60	<u>40.34</u>	<b>43.81</b>
Precision (%)	50.90	49.18	53.03	51.92	52.62	<u>53.68</u>	50.88	<b>54.01</b>
Recall (%)	29.49	30.84	26.57	27.17	26.03	25.65	<u>31.44</u>	<b>32.90</b>
F1-score (%)	37.35	37.91	35.40	35.67	34.83	34.71	<u>38.86</u>	<b>40.89</b>
Accuracy (%)	91.93	91.76	92.09	92.01	92.05	<u>92.13</u>	91.93	<b>92.24</b>

Table 4: Comparison of bleeding classification performance. All methods except ConvNeXt-T augment bleeding data with an additional  $\sim 5,000$  images. Aug denotes traditional augmentation. Ours significantly improves Recall, showing the value of editing-based augmentation. **Bold**: best; underline: second-best.

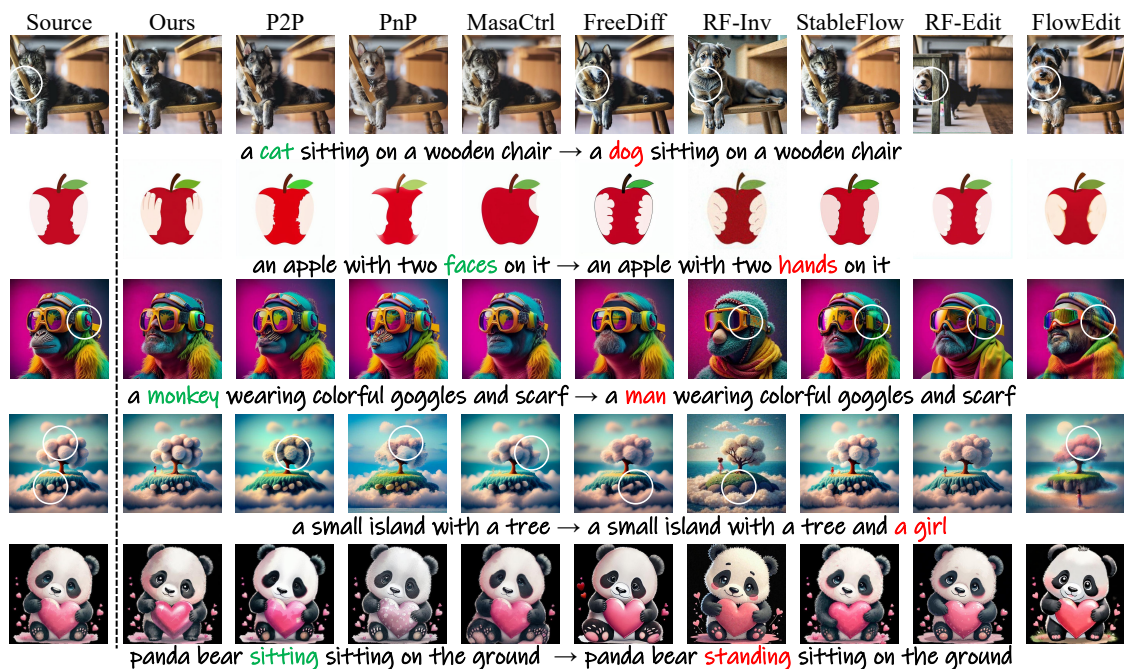


Figure 4: Qualitative comparison. Our method preserves the background while accurately reflecting the target semantics. White circles highlight cases where other methods poorly preserve non-editing regions.

## 5 Editing for Bleeding Classification Task

**Clinical Motivation.** Early detection of abnormal intraoperative bleeding is crucial yet challenging. Identifying early bleeding from surgical videos can assist surgeons in rapidly locating bleeding sites. However, such cases are rare, leading to severe data imbalance. Existing efforts on surgical image or video synthesis rarely address bleeding scenarios, let alone via image editing. To the best of our knowledge, we are the first to explore *text-guided image editing* for surgical bleeding augmentation. Our approach aims to enrich bleeding variations and mitigate data imbalance, ultimately improving downstream classification performance.

**Experimental Setup.** We use the Laparoscopic Roux-en-Y Gastric Bypass dataset (Bose et al. 2025) with 140 videos (100 for training, 40 for testing), sampled at 1fps for 770K frames. The training set includes 512K normal and 44K bleeding frames; the test set has 197K normal and 17K bleeding frames, indicating severe imbalance. We adopt ConvNeXt-T (Liu et al. 2022) as the classification backbone. From the training set, we extract 4,803 early-stage bleeding frames (~50 per video) and edit them into varying bleeding levels (see Appendix). To ensure editing quality and efficiency, we compare four LDM-based and two inversion-free methods, plus standard augmentations (e.g., flipping, rotation). All editing models use SD 1.5 or 3.5 checkpoints without task-specific finetuning. The augmented images are incorporated into training set for downstream classification.

**Results and Analysis.** Quantitative results are shown in Table 4, with visualizations provided in the Appendix. Traditional augmentation methods yield only marginal gains despite generating an additional 5K frames. The four

inversion-based methods improve precision but degrade recall, indicating reduced sensitivity to bleeding. This may stem from limited generation quality of SD1.5, potentially requiring domain-specific finetuning. In contrast, FlowEdit and our method both improve recall, suggesting better utility in enhancing bleeding classification. However, FlowEdit’s poor background preservation may lead to false positives. Benefiting from both semantic fidelity and structural consistency, our method achieves the most balanced performance, demonstrating the potential of image editing techniques for surgical data augmentation and downstream applications.

## 6 Conclusion

We present FIA-Edit, a novel and efficient inversion-free image editing framework that introduces Frequency-Interactive Attention for improved semantic alignment and background preservation. By explicitly modeling source-target interactions in both the frequency and spatial domains, our approach addresses key limitations of existing inversion-free methods, achieving high-quality, controllable image edits without costly inversion. To the best of our knowledge, we are the first to explore the use of generative image editing for *clinical data augmentation*. Specifically, we apply FIA-Edit to synthesize plausible variations of surgical bleeding images, resulting in improved performance on downstream classification task. This demonstrates the potential of controlled image editing in medical data scenarios, opening new avenues for future research. We hope our work contributes a new perspective to tuning-free editing and inspires broader exploration of generative techniques in real-world applications.

## Acknowledgments

This work was supported in part by National Key R&D Program of China (Grant No. 2023YFC2414900), National Natural Science Foundation of China (Grant No.62202189), and research grants from Wuhan United Imaging Healthcare Surgical Technology Co., Ltd.

## References

- Avrahami, O.; Patashnik, O.; Fried, O.; Nemchinov, E.; Aberman, K.; Lischinski, D.; and Cohen-Or, D. 2024. Stable Flow: Vital Layers for Training-Free Image Editing. *arXiv preprint arXiv:2411.14430*.
- Bose, R.; Nwoye, C. I.; Lazo, J.; Lavanchy, J. L.; and Padoy, N. 2025. Feature Mixing Approach for Detecting Intraoperative Adverse Events in Laparoscopic Roux-en-Y Gastric Bypass Surgery. *arXiv preprint arXiv:2504.16749*.
- 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.
- Cai, L.; Zhao, K.; Yuan, H.; Wang, X.; Zhang, Y.; and Huang, K. 2025. DFVEdit: Conditional Delta Flow Vector for Zero-shot Video Editing. *arXiv preprint arXiv:2506.20967*.
- 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*, 22560–22570.
- Deng, Y.; He, X.; Mei, C.; Wang, P.; and Tang, F. 2024. Fire-Flow: Fast Inversion of Rectified Flow for Image Semantic Editing. *arXiv preprint arXiv:2412.07517*.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Gao, J.; Yang, K.; Yao, X.; and Hu, Y. 2025. Unity in Diversity: Video Editing via Gradient-Latent Purification. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 23401–23411.
- Geyer, M.; Bar-Tal, O.; Bagon, S.; and Dekel, T. 2023. Tokenflow: Consistent diffusion features for consistent video editing. *arXiv preprint arXiv:2307.10373*.
- 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*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hu, Y.; Peng, J.; Lin, Y.; Liu, T.; Qu, X.; Liu, L.; Zhao, Y.; and Wei, Y. 2025. DCEdit: Dual-Level Controlled Image Editing via Precisely Localized Semantics. *arXiv preprint arXiv:2503.16795*.
- Ju, X.; Zeng, A.; Bian, Y.; Liu, S.; and Xu, Q. 2023. Direct inversion: Boosting diffusion-based editing with 3 lines of code. *arXiv preprint arXiv:2310.01506*.
- Kim, J.; Hong, Y.; and Ye, J. C. 2025. FlowAlign: Trajectory-Regularized, Inversion-Free Flow-based Image Editing. *arXiv preprint arXiv:2505.23145*.
- Koo, G.; Yoon, S.; Hong, J. W.; and Yoo, C. D. 2024. Flex-iedit: Frequency-aware latent refinement for enhanced non-rigid editing. In *European Conference on Computer Vision*, 363–379. Springer.
- Kulikov, V.; Kleiner, M.; Huberman-Spiegelglas, I.; and Michaeli, T. 2024. Flowedit: Inversion-free text-based editing using pre-trained flow models. *arXiv preprint arXiv:2412.08629*.
- Labs, B. F. 2024. FLUX. <https://github.com/black-forest-labs/flux>.
- Li, G.; Yang, Y.; Song, C.; and Zhang, C. 2025. FlowDirector: Training-Free Flow Steering for Precise Text-to-Video Editing. *arXiv preprint arXiv:2506.05046*.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2022. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*.
- Liu, X.; Gong, C.; and Liu, Q. 2022. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11976–11986.
- Miyake, D.; Iohara, A.; Saito, Y.; and Tanaka, T. 2025. Negative-prompt inversion: Fast image inversion for editing with text-guided diffusion models. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2063–2072. IEEE.
- 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.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- 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.
- Ren, Y.; Jiang, Z.; Zhang, T.; Forchhammer, S.; and Süsstrunk, S. 2025. FDS: Frequency-Aware Denoising Score for Text-Guided Latent Diffusion Image Editing. *arXiv preprint arXiv:2503.19191*.
- 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.

Rout, L.; Chen, Y.; Ruiz, N.; Caramanis, C.; Shakkottai, S.; and Chu, W.-S. 2024. Semantic image inversion and editing using rectified stochastic differential equations. *arXiv preprint arXiv:2410.10792*.

Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

Stability AI. 2025. Stable Diffusion 3.5 Medium.

Tumanyan, N.; Bar-Tal, O.; Bagon, S.; and Dekel, T. 2022. Splicing vit features for semantic appearance transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10748–10757.

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, J.; Pu, J.; Qi, Z.; Guo, J.; Ma, Y.; Huang, N.; Chen, Y.; Li, X.; and Shan, Y. 2024. Taming rectified flow for inversion and editing. *arXiv preprint arXiv:2411.04746*.

Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.

Wu, W.; Fan, Q.; Qin, S.; Gu, H.; Zhao, R.; and Chan, A. B. 2024. FreeDiff: Progressive Frequency Truncation for Image Editing with Diffusion Models. In *European Conference on Computer Vision*, 194–209. Springer.

Xie, C.; Li, M.; Li, S.; Wu, Y.; Yi, Q.; and Zhang, L. 2025. DNAEdit: Direct Noise Alignment for Text-Guided Rectified Flow Editing. *arXiv preprint arXiv:2506.01430*.

Xu, P.; Jiang, B.; Hu, X.; Luo, D.; He, Q.; Zhang, J.; Wang, C.; Wu, Y.; Ling, C.; and Wang, B. 2024. Unveil Inversion and Invariance in Flow Transformer for Versatile Image Editing. *arXiv preprint arXiv:2411.15843*.

Xu, S.; Huang, Y.; Pan, J.; Ma, Z.; and Chai, J. 2023. Inversion-free image editing with natural language. *arXiv preprint arXiv:2312.04965*.

Yang, X.; Zhu, L.; Fan, H.; and Yang, Y. 2025. Videograin: Modulating space-time attention for multi-grained video editing. In *The Thirteenth International Conference on Learning Representations*.

Yoon, S.; Koo, G.; Hong, J. W.; and Yoo, C. D. 2024. Dni: Dilutional noise initialization for diffusion video editing. In *European Conference on Computer Vision*, 180–195. Springer.

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.