

Editing Everything Everywhere All at Once

Fabio Quattrini^{*1}, Carmine Zaccagnino^{*1}, Enis Simsar², Marta Tintoré Gazulla³, Rita Cucchiara¹, Alessio Tonioni³, and Silvia Cascianelli¹

¹ University of Modena and Reggio Emilia

² ETH Zurich

³ Google

<https://mice.silviacascianelli.com>

Abstract. Editing multiple elements of an image in a single forward pass is a practical alternative to multi-turn image manipulation, offering improved efficiency and potentially better harmonization. However, when several instructions target different regions, semantic interference often leads to attribute leakage and poor edit disentanglement, especially as the number of edits increases. In this work, we propose **MICE** (Multi-Instance Concurrent Editing), a training-free strategy for scalable multi-instance image editing with Multimodal Diffusion Transformers. MICE modifies the additive bias of joint attention to regulate interactions between instance-specific edit instructions, latent, and context tokens identified via user-provided segmentation masks. Specifically, MICE allows intra-instance attention, penalizes interactions between neighboring region tokens, and suppresses unrelated cross-instance attention. As a result, our method enforces attribute binding while preserving global visual consistency. We evaluate MICE on LoMOE-Bench and introduce **MICE-Bench**, a more challenging benchmark with an average of 8.5 concurrent edits per image. The experiments demonstrate that our approach outperforms strong baselines and recent competitors in terms of visual quality preservation and faithfulness to the editing instructions.

Keywords: Multi-Instance Editing · Image Editing · Flow Matching

1 Introduction

Concurrent modification of multiple elements within a single image has recently emerged as a compelling extension of text-guided image editing [1, 3, 10, 43]. A possible solution would entail multi-turn pipelines [42, 47], where edits are applied sequentially. However, editing multiple instances one at a time is time-consuming, and the final output can depend on the order in which the edits are performed in case of overlapping instances. Concurrent editing enables all transformations to be performed in one pass, offering clear advantages in efficiency and usability. Nonetheless, existing approaches implementing this setting [19, 38] face a fundamental challenge: when several instructions target different regions or objects, their semantic signals may interfere, leading to *attribute leakage* across

* Equal contribution, order determined by a coin flip.




Fig. 1: Our MICE approach modifies the joint attention maps of MMDiT-based flow matching generative models to concurrently edit multiple elements at inference time.

instances. As the number of edits increases, ensuring edit disentanglement, spatial locality, and global visual coherence becomes progressively more difficult (a quantitative analysis of this phenomenon is reported in the supplementary).

To address this problem, existing approaches primarily entail inference-time adaptations of diffusion and flow-based generative models. A first family of methods performs gradient descent on the image latents to progressively correct undesired interactions between edits [3, 10]. A second line of work intervenes directly in the attention mechanisms of Multimodal Diffusion Transformers [9] (MMDiTs), introducing biasing strategies at selected layers of the velocity field estimator to regulate token interactions [29, 43, 46]. Compared to gradient descent, attention-based solutions are generally more efficient and more readily applicable to transformer-based backbones. However, prior biasing strategies [43, 46] typically require defining multiple bias matrices to be applied to architecture-dependent layers determined by hyperparameters (e.g., a strict bias to isolate edits in certain layers and a softer bias to promote harmonization in other layers). This makes performance sensitive to model-specific design choices and reduces portability across architectures.

In this work, we propose a training-free and architecture-agnostic strategy for multi-instance concurrent editing, dubbed **MICE**. Our approach operates directly on the joint attention mechanism of MMDiTs by modifying the additive bias applied to the attention logits (which is commonly set to 0). Rather than enforcing hard structural constraints, we reshape token interactions in a principled manner to balance instance isolation with global coherence. Specifically, we modify the additive bias in the joint attention operation so that the interactions among prompt, latent, and context tokens relative to each edit instance are isolated from those of other edit instances, ensuring attribute binding, while still maintaining visibility to relevant out-of-instance tokens to enhance harmonization. We identify the instances via segmentation masks that can be provided by the user or the output of a promptable segmentation model [17, 39], and use

such masks to identify instance-specific tokens among the text prompt, latent, and context source image tokens given as input to the MMDiT. Then we compute the additive bias of the joint attention map using the masks, smoothed with a parametric Gaussian kernel. The effect of our strategy is that the tokens relative to a specific instance attend to each other with no bias, can attend to neighboring tokens with a decreasing negative bias (accommodating for small shape variations and helping harmonize the overall appearance of the image), and are prevented from attending faraway, non-relevant tokens (possibly relative to different editing instances). As a result, our approach prevents attribute leakage between editing instances while obtaining a harmonized, visually appealing overall image.

We instantiate our method on Flow Matching-based backbones and evaluate it on LoMOE-Bench [3], a recent benchmark designed for parallel multi-instance editing. Note that the samples in LoMOE-Bench typically contain 3 concurrent edits, rendering it insufficient for evaluating performance under challenging conditions. Therefore, to evaluate performance under more complex conditions, we introduce **MICE-Bench**, a new benchmark featuring samples with approximately 8.5 concurrent edits on average.⁴ Across both benchmarks, MICE demonstrates superior instruction-following capability compared to strong baselines and recent competitors, both in terms of the number of edits correctly executed, background preservation, and the faithfulness of each modification to its corresponding textual instruction. For reproducibility, we release the code.⁵

2 Related Work

Conditional generative image editing methods have seen great progress and consolidated around Autoregression [12, 20, 34, 35] or Diffusion [19, 24, 33, 37, 38] approaches, with Flow Matching (FM) [23] emerging as a dominant framework for efficient, straight-path sampling [36]. The conditioning signal is generally multimodal, comprising guidance [45], text instructions [9], and context images [4, 19, 38]. However, these techniques rely on globally conditioned velocity fields and struggle with local, fine-grained localization. Recent works address this by using attention regularization [32, 44] or constrained cross-token interactions [7, 16] but primarily target single-region or global edits and often fail to maintain coherence when applied to multiple disjoint regions simultaneously. While some methods investigate region-aware generation [5, 8] and causal masking [14] to inject spatial control, most FM editors [19, 38] still rely on globally conditioned velocity fields and suffer attribute leakage between regions during multi-instance editing [26, 46], a limitation we address via our instance-disentangled attention mechanism. To address the aforementioned challenges of multi-instance editing, current approaches primarily cluster into three categories: scaling inference time (*multi-turn*), scaling the architecture (*multi-branch*), or regularizing the internal generation process.

⁴ https://hf.co/datasets/blowing-up-groundhogs/mice_bench

⁵ <https://github.com/Blowing-Up-Groundhogs/mice>

Multi-Turn. State-of-the-art editing models are trained to reduce visual drift and show promise in iterative editing workflows [19, 38], where edits are performed sequentially. While serving as a functional baseline, multi-turn editing differs fundamentally from concurrent editing. Technically, visual degradation accumulates as the number of iterations exceeds a few cycles (around six [19]), and inference time scales linearly with the number of instructions, making this approach infeasible as instance edit counts increase. Methodologically, the final result is path-dependent, especially with overlapping regions. We therefore focus on *concurrent* editing, where all changes occur in a single pass.

Multi-Branch. To enforce disentanglement, some approaches split the generation process into multiple parallel branches, typically one per edit or region. MultiDiffusion [1], the seminal method in this paradigm, binds multiple generation paths by merging independent denoising steps into a least-squares optimal solution. LoMOE [3] adopts this framework to apply localized influence on the N target regions using foreground masks, while other approaches further refine this by using regional branches [5], performing independent inversion searches [41] or grouping edit "aspects" into specialized branches [16]. Finally, LayerEdit [10] takes this paradigm to its extreme by decomposing the scene into $N + 2$ independent layers. While multiple generation paths allow instance isolation, these methods suffer from severe scaling bottlenecks in time and memory usage. Moreover, computational overhead is worsened by the latent optimization often required. For instance, LayerEdit and LoMOE rely on iterative gradient descent and perform a backward pass over N parallel branches. Furthermore, these methods rely on heuristic reassembly to smooth visual boundaries, which often fails to maintain global coherence in dense, overlapping scenes.

Regularization. Another class of methods focuses on end-to-end, single-pass editing by regularizing the internal generation process. One prominent direction involves inference-time optimization. MDE-Edit [48] performs gradient descent on the latents at each denoising step to minimize specialized losses. While these methods maintain a single-pass memory footprint, they rely on aggregating $O(N)$ loss terms, introducing significant computational overhead and numerical instability during backpropagation, particularly when handling multiple instances of varying spatial scales, where large-scale objects can dominate the gradient signal and suppress smaller targets. Other approaches perform fine-tuning or train additional modules to target quantity perception [21] or personalization [25]. Closer to our work, some approaches tweak the attention operator. IDAttn [43] introduces binary disentanglement masks $(0, -\infty)$ to enforce instance isolation. However, this approach relies on heuristics to balance different masking regimes in different model layers (early, middle, and late). This rigidity makes performance sensitive to model architecture and hinders porting to other baselines. In contrast, MICE offers a principled, architecture-agnostic solution by replacing hard binary masks with a parametric additive bias, enabling a soft, continuous transition between instance isolation and global harmonization without per-layer tuning or optimization overhead.

3 MICE

In this section, we present our approach for multi-instance image editing. The task consists in modifying the content of multiple areas of an image (instances) according to user-provided editing instructions, while preserving its original visual style and layout consistency. Our approach builds upon a recent MMDiT-based flow matching model capable of text-conditioned image editing, which we extend with a novel biased joint attention mechanism. Unlike prior work that relies on strict hard masking strategies, which often result in harsh boundaries and poor coherence, our method introduces a continuous, spatially-smoothed attention bias formulation, as detailed below.

3.1 Preliminaries

Formally, let $I^{\text{ref}} \in \mathbb{R}^{H \times W \times 3}$ denote the input image and consider a set of N segmentation masks $M = \{m_n\}_{n=1}^N$, each enclosing an object, and a set of corresponding editing instructions $T = \{t_n\}_{n=1}^N$, *e.g.*, “Change the mouse to a black chick”. The masks can either be provided by a user or by auxiliary modules for Instance Segmentation. The task entails generating an edited image I^{edit} such that each instance within m_n is replaced according to t_n , while maintaining the global style and visual coherence of I^{ref} .

Our framework operates within a MMDiT [9, 19] architecture trained with conditional rectified flow matching [23, 36]. Specifically, given a target data sample $I_1^{\text{edit}} \sim p_{\text{data}}$, a source sample $I_0^{\text{edit}} \sim p_0$, where $p_0 \sim \mathcal{N}(0, \mathbf{I})$, and conditioning signals $\mathcal{C} = \{I^{\text{ref}}, T, M\}$, the model learns a time-dependent velocity field $v_\theta(I_t^{\text{edit}}, t, \mathcal{C})$ to denoise I_0^{edit} into I_1^{edit} . The input image I^{ref} is encoded by a Variational Autoencoder (VAE), which compresses it to $\mathbb{R}^{\frac{H}{d} \times \frac{W}{d} \times c}$ where d is the compression ratio and c is the number of channels of the VAE. The editing instructions t_i are tokenized, padded to a fixed length l , and then passed through a text encoder, usually a language model like T5 [30] or Qwen [40], which returns a sequence of l tokens encoding the information contained in the editing instruction. Then, these are fed to a velocity field estimator, alongside the noisy latents for I_t^{edit} . After a predefined number of estimation steps, the resulting latents are passed to a VAE decoder to obtain the final I_0^{edit} .

In state-of-the-art FM models, the velocity field estimator is an MMDiT, where tokens from various modalities (textual prompt, noisy visual latents, and conditioning context) are concatenated and processed via joint attention. Let $\mathbf{Z} = \mathbf{Z}^{\text{prompt}} \parallel \mathbf{Z}^{\text{latent}} \parallel \mathbf{Z}^{\text{context}}$ denote the joint token sequence. From $\mathbf{Z}$ we compute the queries $\mathbf{Q}$, keys $\mathbf{K}$, and values $\mathbf{V}$ for the joint attention operation defined as:

$$\text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}} + \mathbf{B} \right) \mathbf{V},$$

where $\mathbf{B} \in \mathbb{R}_{\leq 0}^{|\mathbf{Z}| \times |\mathbf{Z}|}$ is an attention bias matrix, commonly set to 0. Our approach defines the values of the bias matrix based on the types of tokens that interact to achieve editing isolation and preserve background and overall visual quality.



Fig. 2: Instance-aware smoothing strategy. Each segmentation mask is smoothed independently, and the smoothed values are zeroed when touching the edge of a neighboring, non-smoothed mask (*e.g.*, the red and yellow masks). If an instance is sufficiently distant from the others, its mask is smoothed with no constraints (*e.g.*, the blue mask).

3.2 Smoothly-Disentangled Attention

Prior instance-disentangled attention methods [29, 43, 46] enforce isolation by resorting to binary attention bias matrices $\mathbf{B}$, with 0 allowing interactions and $-\infty$ blocking them. This approach is effective for semantic isolation but can lead to noticeable seams, incorrect local lighting, and a lack of smooth blending with the nearby areas of the image. To contrast this, some approaches [43, 46] relax disentanglement by allowing full latent and context self- and cross-interactions in some layers of the MMDiT. However, this strategy entails identifying the optimal layers where to apply the disentangled and relaxed regimes for each architecture, limiting the method portability. To tackle this, we propose to apply the same bias matrix in all layers by designing it to ensure a trade-off between disentanglement and harmonization. Specifically, we not only set the values of the matrix to 0 to allow interactions and $-\infty$ to prevent them, but also use intermediate negative values to dampen the attention without fully blocking it.

Editing Areas Definition. To determine the values of the bias, we start from the segmentation masks of each instance, which can be either provided directly by the user or obtained from an instance segmentation model. We smooth such masks via a Gaussian kernel so that the attention will be increasingly penalized as the spatial distance from the instance mask increases. Specifically, given a binary 2D segmentation mask $m_n \in \{0, 1\}^{H \times W}$ for instance n , we downsample it to account for VAE compression, then we apply a 2D Gaussian convolution with kernel size k and obtain a smoothed mask s_n . To prevent interference between instances due to the smoothing, we devise a specific instance-aware smoothing strategy. We consider each instance separately and apply the Gaussian smoothing to its segmentation mask. Given instance n , if an element in the latent-space feature map is part of the mask of any other instance, its value is explicitly overridden to 0 (see Fig. 2).

The resulting continuous values lie in $[0, 1]$, therefore, we map them into the log-probability space $(-\infty, 0]$ to serve as an attention bias. Once the smoothed masks are defined and corrected to ensure instance-awareness, we unroll them in raster order. Each of them becomes a row in the bias matrix $\mathbf{B}$, whose values are defined as follows. Given a visual token Z_j (either in $\mathbf{Z}^{\text{latent}}$ or $\mathbf{Z}^{\text{context}}$) outside the strict bounds of instance n , the bias in the attention with a latent or context token Z_i belonging to instance n is defined as:

$$\mathcal{B}(j) = \frac{1}{\tau} \log(s_n[j] + \epsilon)$$

where τ is a temperature hyperparameter controlling the falloff intensity, and $\epsilon = 10^{-30}$ prevents numerical instability. Intuitively, k and τ control the spatial extent and the strength of the blending between the edited instance and the surrounding regions. Large values for k allow tokens to interact with a wider neighborhood, which is beneficial when the edit modifies the object shape or when the mask is irregular, and large values for τ typically yield smoother blending of the edited instance with the surrounding background. On the other hand, smaller values for these hyperparameters keep the editing more localized. A deeper analysis of the effect of k and τ is reported in the supplementary material.

The definition of the overall matrix is detailed below and schematized in Fig. 3.

Bias Matrix. To define the bias matrix, we logically partition the sequence of tokens $\mathbf{Z}$ according to the N instances and their corresponding target regions in the source context image and the generated image. Specifically:

- The **prompt tokens** $\mathbf{Z}^{\text{prompt}}$ are grouped in sequences corresponding to each editing prompt for the instances $\mathbf{Z}_n^{\text{prompt}}$ and an auxiliary sequence that represents the global prompt $\mathbf{Z}_g^{\text{prompt}}$. This can either be empty to keep the original overall appearance of the image, or instantiated to change its global attributes (*e.g.*, style, lighting, color palette).
- The **latent tokens** $\mathbf{Z}^{\text{latent}}$ are organized into background tokens $\mathbf{Z}_u^{\text{latent}}$ that should stay unedited, and local instance tokens $\mathbf{Z}_n^{\text{latent}}$ corresponding to the area that will be occupied by each edited instance in the final image.
- The **context tokens** $\mathbf{Z}^{\text{context}}$ are distinguished between tokens belonging to background $\mathbf{Z}_u^{\text{context}}$ and instances $\mathbf{Z}_n^{\text{context}}$ similarly to the latent tokens.

To formally align the global token sequence with the spatial dimensions of the masks, we define a mapping function that yields the spatial raster index of any visual token $Z_j \in \mathbf{Z}^{\text{latent}} \cup \mathbf{Z}^{\text{context}}$, *i.e.*, $\pi(\cdot) : \{0, \dots, |\mathbf{Z}|\} \rightarrow \{0, \dots, \frac{H}{d} \times \frac{W}{d}\}$. The attention bias matrix is thus defined as:

$$\mathbf{B}_{ij} = \begin{cases} 0 & (Z_i, Z_j) \in \mathbf{Z}_g^{\text{prompt}} \vee (Z_i, Z_j) \in \mathbf{Z}_n^{\text{prompt}} \quad \forall n \in [1, N] & (1) \\ 0 & Z_i \in \mathbf{Z}_g^{\text{prompt}} \cup \mathbf{Z}_u^{\text{latent}} \wedge Z_j \notin \mathbf{Z}^{\text{prompt}} & (2) \\ 0 & Z_i \in \mathbf{Z}_g^{\text{latent}} \cup \mathbf{Z}^{\text{context}} \wedge Z_j \in \mathbf{Z}_g^{\text{prompt}} & (3) \\ 0 & Z_i \in \mathbf{Z}_n^{\text{context}} \wedge Z_j \in \mathbf{Z}_n^{\text{prompt}} \quad \forall n \in [1, N] & (4) \\ 0 & Z_i \in \mathbf{Z}^{\text{context}} \wedge Z_j \notin \mathbf{Z}^{\text{prompt}} & (5) \\ \mathcal{B}(\pi(j)) & Z_i \in \mathbf{Z}_n^{\text{latent}} \cup \mathbf{Z}_n^{\text{prompt}} \quad \forall n \in [1, N] \wedge Z_j \notin \mathbf{Z}^{\text{prompt}} & (6) \\ \mathcal{B}(\pi(i)) & Z_i \in \mathbf{Z}^{\text{latent}} \wedge Z_j \in \mathbf{Z}_n^{\text{prompt}} \quad \forall n \in [1, N] & (7) \\ -\infty & \text{otherwise} & (8) \end{cases}$$

In other words, we enforce strict binding between the tokens relative to the prompts (*Case 1*) to prevent attribute leakage within the textual domain. Then, we keep the interaction of the global prompt and the background latents with all the latent and context tokens unbiased (*Case 2*) to favor visual coherence in the unedited background regions, and do the same where the background latents and the context tokens attend to those relative to the global prompt (*Case 3*).




Fig. 3: Our MICE approach entails defining a bias map (whose logic representation is on the right) that regulates the interaction between background and instance-specific prompt, latent, and context tokens, based on instances’ localization information in the form of smoothed segmentation masks.

Furthermore, we apply a bias based on instance-aware smoothed masks to the interactions between prompt tokens and latent tokens, relative to each instance and to all latent tokens and context tokens (*Case 6*). This ensures that the local, instance-specific prompt and the noisy image area enclosing each instance can interact with areas in the context image and the noisy image that are larger than the exact ones occupied by the instance, with a strength that is decreasing proportionally to the distance from the instance. In this way, the edited instance will blend better with its surroundings in the final image. The same biasing strategy is applied to the interaction between the latent tokens and the instance-specific prompts (*Case 7*) to ensure coherent localization of the editing instruction. As for the context tokens, we prevent those relative to the background from interacting with the instance prompts, and let the context tokens for each instance interact only with the corresponding instance-specific prompt (*Case 4*). This choice is intended to better localize the instance-specific prompts within the reference image given as context. Finally, we keep the attention from the context tokens to the latent and context tokens unbiased (*Case 5*) to favor layout, colors, and style preservation in the final image.

4 MICE-Bench

Existing benchmarks for multi-instance concurrent editing on natural images lack settings with a high number of concurrent edits. For example, the benchmark proposed by [41] contains only one or two edits per image, and [3] is capped at seven edits, averaging only three. These benchmarks, therefore, do not provide significant avenues for verifying a model’s ability to perform an arbitrarily high number of concurrent edits.

To bridge this gap, we introduce **MICE-Bench** (Multi-Instance Concurrent Editing Benchmark), which we specifically annotate for high-density editing. Notably, the benchmark features 260 samples with a significantly higher editing

density than prior works, ranging from 5 to 40 concurrent edits per image, with an average of 8.5. Formally, each sample can be defined as $(I^{src}, M, \mathbb{T})$ where $I^{src} \in \mathbb{R}^{H \times W \times 3}$ is the source image to be edited by the model, $M = \{m_n\}_{n=1}^N$ is the set of segmentation masks for each object in the image, and $\mathbb{T}$ is a set of prompts in the format "*Replace the {SRC_OBJ} with {TGT_OBJ}*", where {SRC_OBJ} and {TGT_OBJ} are respectively the name of the object to be edited in the source image and the name of the object that should replace it in the edited image. We provide some samples in the supplementary material.

Data Collection and Annotation. We select 260 samples from the LVIS [11] and COCO [22] validation sets, focusing on having a diverse object count per sample ranging from 5 to 40 instances associated with segmentation masks and source object class names. For each of them, we feed a multimodal LLM (Gemini 3 Flash) the source image, object class names, and normalized bounding boxes $[y_{min}, x_{min}, y_{max}, x_{max}]$ enclosing the original segmentation masks. The model is prompted to suggest creative yet contextually realistic replacements for every masked object. We manually verify these suggestions to ensure target feasibility within the scene (*e.g.*, a "chipmunk" fits where a "mouse" is, whereas an "elephant" does not). We also use Gemini 3 Flash to generate an additional set of instructions explicitly containing textual information to identify each instance within the image (*e.g.*, referring expressions, locations). These are generated both as individual instructions and as a single global prompt. Thanks to the multi-faceted annotation strategy (some samples are reported in the supplementary), MICE-Bench can be used to evaluate a wide range of approaches for multi-instance concurrent editing.

5 Experimental Analysis

Experiment setup. To evaluate our approach for multi-instance concurrent editing, we resort to our proposed MICE-Bench dataset (Sec. 4), and to the recently-proposed LoMOE-Bench [3] dataset, which contains 64 images associated with 2 to 7 editing instructions. To measure the performance, we adopt the following scores proposed in previous approaches tackling the same task [3, 43]:

- Target CLIP Score (**Tgt C**), *i.e.*, the CLIP score [15] computed between the global prompt describing all the edits and the entire edited image [15]. This score captures the overall semantic adherence to the editing instructions.
- Localized CLIP Score (**Loc C**), *i.e.*, the CLIP score [15] computed between the target object description in the editing instruction and the bounding box areas around the edited regions. This score measures more precisely the capability of the model to perform each edit on each intended instance.
- Mean Absolute Error on the Background pixels (**MAE_B**), to estimate the background preservation capabilities of the model under investigation.
- Attempt Rate (**AR_%**), *i.e.*, the percentage of edits that the model has attempted to perform. Specifically, an edit is considered attempted when the pixel-wise difference between the original and edited areas relative to an instance exceeds a fixed threshold.

Implementation Details. We develop our training-free MICE method on the state-of-the-art editing model FLUX.2 [klein] 4B [18] and use it for the majority of our experiments, setting 4 inference steps. For MMDiT-based porting experiments, we choose the FLUX.2 [klein] 9B [18], FLUX.1-Kontext-dev [19] 12B, and FLUX.2-Dev [18] 32B. We use weights and reference code from HuggingFace [27], setting 4 inference steps for FLUX.2 [klein] 9B and 28 steps for FLUX.1-Kontext-dev and FLUX.2-Dev. We use these models end-to-end, prompting them with textual prompts describing all the inputs (an analysis on the most suitable prompting strategy is reported in the supplementary material). For smoothing the segmentation masks of the instances, we set kernel size $k=3$ and, due to differences in the edit types in the two benchmark datasets, temperature $t=1$ for LoMOE-Bench and $t=8$ for MICE-Bench. Please refer to the supplementary for further discussions. To compare MICE against alternative multi-instance editing approaches, we select representative methods for the categories described in Section 2. In particular, we adapt the archetypal multi-branch MultiDiffusion [1], originally proposed for region-based generation, and apply it to FLUX.2 [klein] 4B [18] by using the official code as reference.⁶ We run the multi-branch LoMOE [3], based on Stable Diffusion 2 [31] by using the official implementation⁷ and also port it on FLUX.2 [klein] 4B [18]. Moreover, we run the multi-branch LayerEdit [10], based on Stable Diffusion XL [28] as released by the respective authors.⁸ We also implement the regularization-based multi-instance edit approach IDAttn [43], based on FLUX.1-Kontext-dev [19], and port it on FLUX.2 [klein] 4B. Finally, we compare against the open-source Qwen-Image-Edit-2511 [38] and the closed-source Gemini 3 Pro Image. We run all experiments on NVIDIA A100 40GB, with the exception of the commercial Gemini 3 Pro Image model, which we prompt by using the public Vertex API.

5.1 Results

Effect of the Editing Area Definition. In Table 1, we report an ablation on the effect of the signal used for instance localization. We compare Bounding Boxes (as used by [43] or in the supplementary prompting experiments), strict segmentation masks (as used by [3, 10]), naïvely-smoothed segmentation masks (using the Gaussian kernel), and masks obtained with our proposed instance-aware strategy. From the results, we can see how simply using masks instead of Bounding Boxes allows the model to better localize the edits while preserving the background. When introducing naïve smoothing, the model background preservation capabilities increase at the cost of reduce localized edits. This is mainly due to the overlap between instances and is therefore more detrimental on Mice-Bench. Adding instance-aware smoothing fixes this, and achieves the best performance, as the model is capable of localizing edits well while also harmonizing them in the background, without unwanted inter-instance leakages.

⁶ <https://github.com/omerbt/MultiDiffusion>

⁷ <https://github.com/goirik-chakrabarty/LoMOE>

⁸ <https://github.com/fufy1024/LayerEdit>

Table 1: Ablation analysis of different editing areas definition strategies based on the segmentation masks (Seg. M.). Here, we apply MICE to FLUX.2 [klein] 4B.

	LoMOE				MICE-Bench			
	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$
Bounding Boxes	25.94	29.55	9.38	97.92	25.00	26.57	10.66	98.78
Original Seg. M.	25.78	29.59	7.60	96.35	24.95	26.63	8.15	95.36
Gaussian Seg. M.	25.77	29.55	7.08	96.35	25.00	26.19	7.63	87.30
Our Seg. M.	25.77	29.57	7.07	96.35	25.15	26.85	7.84	92.48

Adaptation Strategies for Multi-Instance Concurrent Editing. In Table 2, we compare different training-free strategies proposed in the literature to obtain a multi-instance concurrent editing behavior from generative models that are trained for single-instance editing. Specifically, we implement the strategies proposed in [1] (**MultiDiffusion**), in [3] (**LoMOE**), and in [43] (**IDAttn**) on top of the same backbone (FLUX.2 [klein] 4B). The combination of multi-branch diffusion paths and gradient-based optimizations of LoMOE results in better performance on the simpler LoMOE-Bench dataset. However, MICE performs better on the more complex MICE-Bench in terms of localized edit quality. We argue that the reason is two fold. First, multi-branch approaches applied to few sampling steps lack enough interaction points, especially with increasing instance number (and, therefore, number of branches to merge). Second, performing latent optimization for multiple instances creates numerical instabilities. This is also shown by the superior performance on MICE-Bench of MultiDiffusion, which has a pipeline similar to LoMOE but without gradient optimization. Regarding AR% and MAE_B, we remark that they are not always correlated with better overall performance in the supplementary. Moreover, we consider 1024² images and compute inference time and VRAM usage with increasing number of edits in Figure 4, using a 40GB A100. As we can see, our method scales similarly to IDAttn. In contrast, MultiDiffusion scales linearly both in VRAM usage and inference time (we extrapolate the scaling behavior on >15 edits, as it does not fit in GPU memory). In summary, even with a few edits at practical resolutions, end-to-end regularized approaches are necessary when using current state-of-the-art editing models. In the supplementary, we further elaborate on computational scaling behavior by presenting an asymptotic analysis of different approaches.


**Fig. 4:** Inference time and VRAM comparison *w.r.t.* the number of edits.

Table 2: Comparison of different training-free strategies applied to FLUX.2 [klein] 4B to achieve multi-instance concurrent editing capabilities.

Strategy	LoMOE-Bench				MICE-Bench			
	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$
MultiDiffusion	26.04	30.12	6.39	95.31	24.79	26.26	6.21	85.25
LoMOE	26.23	30.15	8.32	98.44	24.85	26.04	8.62	91.46
IDAttn	25.29	29.12	7.04	92.19	24.64	25.54	6.42	77.05
MICE	25.77	29.57	7.07	96.35	25.15	26.85	7.84	92.48

Table 3: Performance of various MMDiT-based generative editing backbones with the integration of our approach (+ MICE).

	LoMOE-Bench				MICE-Bench			
	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$
FLUX.1 Kontext	25.77	28.26	9.46	93.23	25.65	25.43	16.75	97.95
+ MICE	25.42	29.13	7.49	94.27	24.54	25.54	6.61	86.43
FLUX.2 [klein] 4B	25.94	28.68	8.76	97.40	25.37	25.97	26.08	99.61
+ MICE	25.77	29.57	7.07	96.35	25.15	26.85	7.84	92.48
FLUX.2 [klein] 9B	26.57	29.45	8.63	99.48	25.66	26.34	27.12	100.00
+ MICE	26.49	30.30	8.88	98.44	25.24	27.79	13.89	98.97
FLUX.2-Dev	27.03	29.00	10.48	97.92	25.93	26.25	19.24	99.61
+ MICE	26.44	30.28	11.50	99.48	25.09	26.91	22.54	98.97

Backbone Adaptability. Our approach can be easily integrated in different MMDiT-based editing models. Note that the k and τ hyperparameters can always be adjusted depending on the type of edit to perform, but our approach does not rely on architecture-related hyperparameters such as the layers where to apply our attention biasing strategy. To evaluate the portability of MICE, we apply it to different state-of-the-art image editing backbones and report the results of this analysis in Table 3. First, we consider FLUX.2 [klein] 9B, whose inputs and outputs are passed through the same blocks as in the 4B version, but whose MMDiT is larger. Moreover, we consider the **FLUX.2-Dev** variant, which has a different preprocessing pipeline compared to the [klein] variants (different text and image encoder) and a larger MMDiT. Finally, we consider **FLUX.1 Kontext**, which has yet another preprocessing pipeline compared to the others. As we can see, our MICE method shows consistent improvements in localized edit quality across all baselines on the two benchmarks, indicated by Loc C. In the remaining experiments, we use the FLUX.2 [klein] 9B porting (dubbed MICE-9B).

Comparison with the State-of-the-Art. Table 4 contains the performance of MICE-9B when compared to end-to-end single pass image editing models, including our baseline model (Flux.2 [klein] 9B), state-of-the-art open (**Qwen-Image-Edit** [38]) and closed (**Gemini 3 Pro Image** [6]) source models provided with textual prompts only. Moreover, we consider recent strategies specifically proposed for multi-instance concurrent editing, on the original backbone proposed by the respective authors. These are **LoMOE** [3] and **LayerEdit** [10]

Table 4: Comparison with State-of-the-Art Editing Models.

Model	LoMOE				MICE-Bench			
	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$	Tgt C $\uparrow$	Loc C $\uparrow$	MAE _B $\downarrow$	AR [%] $\uparrow$
FLUX.2 [klein] 9B	26.57	29.45	8.63	99.48	25.66	26.34	27.12	100.00
Gemini 3 Pro Image	26.82	29.67	7.24	97.37	25.64	27.09	13.08	98.97
IDAttn	25.67	29.26	6.41	98.44	24.68	25.54	5.06	73.01
LayerEdit	25.61	29.07	7.43	97.92	23.77	24.38	9.32	99.51
LoMOE	26.00	29.40	6.66	92.19	24.99	25.46	7.76	99.07
Qwen-Image-Edit	26.34	28.85	7.13	88.02	25.26	25.33	26.47	98.83
MICE	26.49	30.30	8.88	98.44	25.24	27.79	13.89	98.97
+ SAM3	26.22	30.41	8.94	96.88	25.26	27.73	13.36	98.54

(which rely on the segmentation masks of the instances to edit, as MICE does), and **IDAttn** [43] (which takes as input the bounding boxes surrounding the instances). In this comparison, we also include the results obtained by our approach when given instance segmentation masks obtained by prompting the instance segmentation model SAM3 [2] instead of the ground truth ones from the datasets. As we can see, the superior performance of our MICE method in Localized CLIP score indicates that the model successfully performs localized edits, even beating Gemini 3 Pro, which showed better performance than the FLUX.2 [klein] 9B end-to-end baseline. Moreover, our method maintains superior performance also when fed with the segmentation masks from SAM3, showing its usability in real-world scenarios where a user can simply provide points as visual cues instead of precise segmentation masks. Further analysis is reported in the supplementary material.

Qualitative Analysis. We report qualitative results in Figure 5. In both samples, LoMOE shows overall low generation quality, even failing convincingly generate any of the instances in the LoMOE-Bench sample. FLUX.2 [klein] in end-to-end mode fares much better, but fails to preserve the overall composition and misplaces some objects. IDAttn, on the other hand, misses some edits, while retaining the overall composition and object placement. MICE shines in background and layout preservation as well as in prompt following in both samples. We perform an LLM-as-judge evaluation using Gemini 3.1 Pro with high thinking budget, with 30 samples for each dataset, passing to the LLM judge the source image, the edit locations, the edit prompts and the result from two models, asking it to choose a winner among the two in terms of prompt following, background preservation and overall visual quality. We include in this evaluation our model, the FLUX.2 [klein] baseline it is based on, and open-source competitors Qwen-Image-Edit-2511, and LoMOE, LayerEdit and IDAttn, asking the LLM 900 total binary choice questions. We compute ELO scores from these two-way match-ups. Note that [13, 43] showed good correlation between user and LLM responses to questions about evaluating image visual features and

Table 5: LLM-as-judge Results.

	LoMOE-Bench	MICE-Bench
FLUX.2 [klein] 9B	1299.78	1172.16
IDAttn	1223.57	1412.73
LayerEdit	887.98	858.92
LoMOE	1181.02	995.05
Qwen-Image-Edit	1265.12	1192.69
MICE	1342.53	1568.46



Fig. 5: Qualitative Results on LoMOE-Bench and MICE-Bench.

ranking models based on image editing ability. Moreover, to substantiate this correlation in our case, we also conduct a user study, for which we report results in the supplementary. We report the results of the LLM-as-judge evaluation in Table 5. On both datasets, MICE ends up being the best-performing model. The second-best on LoMOE is FLUX.2 [klein], followed by Qwen-Image-Edit and IDAttn, with LoMOE and LayerEdit performing significantly worse in this evaluation. When switching to MICE-Bench, the denser annotation favors the specialized IDAttn model, whereas the end-to-end competitors FLUX.2 [klein] and Qwen-Image-Edit fall much further behind. As highlighted by the authors of [43], the bad generation quality of LayerEdit and LoMOE especially fails to be highlighted by CLIP score, but is made evident by human and LLM-based evaluation.

Edge Cases. Additionally, in Figure 6 we show a qualitative analysis of our proposed method in some edge cases. In particular, we show how MICE does not hinder the base FLUX.2 [klein] model capabilities of performing global edits and harmonizing the resulting image, either when the global harmonization is necessary due to the localized prompt (first image, where illumination changes are necessary after the local edit) or when the edit instructions require combining global style changes with local edits (second image). Finally, in the third image, we show that our method can handle multiple edits affecting the same region. Arguably, this is due to MICE’s smoothly-disentangled attention masking strategy, especially considering the output of IDAttn.

6 Conclusion

In this work, we tackle the problem of multi-instance concurrent image editing by proposing MICE, a training-free and architecture-agnostic method that leverages the joint attention mechanism of MMDiTs by redefining the additive bias applied to attention logits. Specifically, MICE alters token interactions according to instance-specific regions to promote attribute binding within instances while maintaining sufficient interactions across the image to preserve its




Fig. 6: Exemplar output of FLUX.2 [klein], IDAttn, and MICE on the edge cases

global coherence. Moreover, we have introduced MICE-Bench, a new dataset designed to assess performance under more demanding conditions, *i.e.*, a higher number of concurrent edits per sample compared to existing benchmarks for multi-instance image editing. Extensive evaluation across diverse datasets and backbones demonstrates that our method consistently improves instruction adherence, edit localization, and background preservation compared with recent baselines and competitors. These results suggest that attention biasing provides an effective mechanism to scale multi-instance editing to complex scenarios.

Acknowledgment

This paper is based upon work supported by the “AI for Digital Humanities” project funded by “Fondazione di Modena”, the GCP Credit Award, the Google Cloud Research Credits program with the award GCP19980904, and the FARD2025 project (CUP E93C25000370005). We acknowledge EuroHPC Joint Undertaking and ISCRA for awarding us access to LUMI at CSC, Finland, LEONARDO at CINECA, Italy, and MareNostrum5 at BSC, Spain.

References

1. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: MultiDiffusion: Fusing Diffusion Paths for Controlled Image Generation. In: ICML (2023)
2. Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., Alwala, K.V., Khedr, H., Huang, A., et al.: Sam 3: Segment anything with concepts. In: ICLR (2026)
3. Chakrabarty, G., Chandrasekar, A., Hebbalaguppe, R., AP, P.: LoMOE: Localized Multi-Object Editing via Multi-Diffusion. In: ACM MM. p. 3342–3351 (2024)

4. Chen, Y., Ma, Z., Jia, G., Jiang, C., Li, J., Zhou, B.: Context-Aware Autoregressive Models for Multi-Conditional Image Generation. arXiv preprint arXiv:2505.12274 (2025)
5. Chen, Z., Li, Y., Wang, H., Chen, Z., Jiang, Z., Li, J., Wang, Q., Yang, J., Tai, Y.: Region-aware text-to-image generation via hard binding and soft refinement. ICCV (2025)
6. Comanici, G., Bieber, E., Schaekermann, M., Pasupat, I., Sachdeva, N., Dhillon, I., Blistein, M., Ram, O., Zhang, D., Rosen, E., et al.: Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. arXiv preprint arXiv:2507.06261 (2025)
7. Dahary, O., Patashnik, O., Aberman, K., Cohen-Or, D.: Be Yourself: Bounded Attention for Multi-Subject Text-to-Image Generation. In: ECCV. pp. 432–448. Springer (2024)
8. Eijkelboom, F., Zimmermann, H., Vadgama, S., Bekkers, E.J., Welling, M., Naesseth, C.A., van de Meent, J.W.: Controlled Generation with Equivariant Variational Flow Matching. In: ICML (2025)
9. 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: ICML. pp. 12606–12633. PMLR (2024)
10. Fu, F., Huang, M., Zhang, L., Mao, Z.: LayerEdit: Disentangled Multi-Object Editing via Conflict-Aware Multi-Layer Learning. AAAI (2026)
11. Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: CVPR (2019)
12. Han, J., Liu, J., Jiang, Y., Yan, B., Zhang, Y., Yuan, Z., Peng, B., Liu, X.: Infinity: Scaling Bitwise AutoRegressive Modeling for High-Resolution Image Synthesis. In: CVPR (2025)
13. Haraguchi, D., Inoue, N., Shimoda, W., Mitani, H., Uchida, S., Yamaguchi, K.: Can GPTs evaluate graphic design based on design principles? In: SIGGRAPH. pp. 1–4 (2024)
14. He, S., Levine, D., Vrkic, I., Bressana, M.F., Zhang, D., Rizvi, S.A., Zhang, Y., Zappala, E., van Dijk, D.: CaLMFlow: Volterra Flow Matching using Causal Language Models. arXiv preprint arXiv:2410.05292 (2024)
15. Hessel, J., Holtzman, A., Forbes, M., Bras, R.L., Choi, Y.: CLIPScore: a reference-free evaluation metric for image captioning. In: EMNLP (2021)
16. Huang, M., Cai, J., Jia, S., Lokhande, V.S., Lyu, S.: ParallelEdits: Efficient Multi-Aspect Text-Driven Image Editing with Attention Grouping. In: NeurIPS (2024)
17. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: ICCV (2023)
18. Labs, B.F.: FLUX.2: Frontier Visual Intelligence. <https://bfl.ai/blog/flux-2> (2025)
19. Labs, B.F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., et al.: FLUX. 1 Kontext: Flow Matching for In-Context Image Generation and Editing in Latent Space. arXiv preprint arXiv:2506.15742 (2025)
20. Li, X., Qiu, K., Chen, H., Kuen, J., Lin, Z., Singh, R., Raj, B.: ControlVAR: Exploring Controllable Visual Autoregressive Modeling. arXiv preprint arXiv:2406.09750 (2024)
21. Li, Y., Chan, K., Sun, Y., Lam, C., Tong, T., Yu, Z., Fu, K., Liu, X., Tan, T.: MoEdit: On Learning Quantity Perception for Multi-object Image Editing. In: CVPR. pp. 2683–2693 (2025)

22. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft coco: Common objects in context (2015), <https://arxiv.org/abs/1405.0312>
23. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow Matching for Generative Modeling. In: ICLR (2023)
24. Liu, S., Han, Y., Xing, P., Yin, F., Wang, R., Cheng, W., Liao, J., Wang, Y., Fu, H., Han, C., et al.: Step1X-Edit: A Practical Framework for General Image Editing. arXiv preprint arXiv:2504.17761 (2025)
25. Matsuda, H., Togo, R., Maeda, K., Ogawa, T., Haseyama, M.: Multi-object editing in personalized text-to-image diffusion model via segmentation guidance. In: ICASSP. pp. 8140–8144. IEEE (2024)
26. Mun, S., Nam, J., Cho, S., Ok, J.: Addressing text embedding leakage in diffusion-based image editing. In: ICCV (2025)
27. von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Nair, D., Paul, S., Berman, W., Xu, Y., Liu, S., Wolf, T.: Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers> (2022)
28. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In: ICLR (2024)
29. Qu, T., Ke, L., Zhan, X., Tang, L., Liu, Y., Peng, B., Yu, B., Yu, D., Jia, J.: Replan: Reasoning-guided region planning for complex instruction-based image editing (2025), <https://arxiv.org/abs/2512.16864>
30. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. JMLR pp. 1–67 (2020)
31. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution Image Synthesis with Latent Diffusion Models. In: CVPR (2022)
32. Simsar, E., Tonioni, A., Xian, Y., Hofmann, T., Tombari, F.: LIME: Localized Image Editing via Attention Regularization in Diffusion Models. In: WACV. pp. 222–231. IEEE (2025)
33. Tan, Z., Liu, S., Yang, X., Xue, Q., Wang, X.: OminiControl: Minimal and Universal Control for Diffusion Transformer. In: ICCV. pp. 14940–14950 (2025)
34. Team, N., Han, C., Li, G., Wu, J., Sun, Q., Cai, Y., Peng, Y., Ge, Z., Zhou, D., Tang, H., et al.: Nextstep-1: Toward autoregressive image generation with continuous tokens at scale. arXiv preprint arXiv:2508.10711 (2025)
35. Tian, K., Jiang, Y., Yuan, Z., Jia, B., Chen, B.: Visual Autoregressive Modeling: Scalable Image Generation via Next-Scale Prediction. In: NeurIPS (2024)
36. Tong, A., Fatras, K., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Wolf, G., Bengio, Y.: Improving and generalizing flow-based generative models with mini-batch optimal transport. TMLR pp. 1–34 (2024)
37. Wang, P., Shi, Y., Lian, X., Zhai, Z., Xia, X., Xiao, X., Huang, W., Yang, J.: SeedEdit 3.0: Fast and High-Quality Generative Image Editing. arXiv preprint arXiv:2506.05083 (2025)
38. 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)
39. Wu, J., Jiang, Y., Liu, Q., Yuan, Z., Bai, X., Bai, S.: General object foundation model for images and videos at scale. In: CVPR (2024)
40. Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., Zheng, C., Liu, D., Zhou, F., Huang, F., Hu, F., Ge, H., Wei, H., Lin,

- H., Tang, J., Yang, J., Tu, J., Zhang, J., Yang, J., Yang, J., Zhou, J., Zhou, J., Lin, J., Dang, K., Bao, K., Yang, K., Yu, L., Deng, L., Li, M., Xue, M., Li, M., Zhang, P., Wang, P., Zhu, Q., Men, R., Gao, R., Liu, S., Luo, S., Li, T., Tang, T., Yin, W., Ren, X., Wang, X., Zhang, X., Ren, X., Fan, Y., Su, Y., Zhang, Y., Zhang, Y., Wan, Y., Liu, Y., Wang, Z., Cui, Z., Zhang, Z., Zhou, Z., Qiu, Z.: Qwen3 technical report (2025), <https://arxiv.org/abs/2505.09388>
41. Yang, Z., Ding, G., Wang, W., Chen, H., Zhuang, B., Shen, C.: Object-Aware Inversion and Reassembly for Image Editing. In: ICLR (2024)
 42. Ye, Y., He, X., Li, Z., Lin, B., Yuan, S., Yan, Z., Hou, B., Yuan, L.: Imgedit: A unified image editing dataset and benchmark. In: NeurIPS (2025)
 43. Zaccagnino, C., Quattrini, F., Simsar, E., Gazulla, M.T., Cucchiara, R., Tonioni, A., Cascianelli, S.: Shifting the Breaking Point of Flow Matching for Multi-Instance Editing. In: ICML (2026)
 44. Zhang, H., Duan, Z., Wang, X., Chen, Y., Zhang, Y.: EliGen: Entity-Level Controlled Image Generation with Regional Attention. In: ACM-MMAsia (2025)
 45. Zheng, Q., Le, M., Shaul, N., Lipman, Y., Grover, A., Chen, R.T.: Guided flows for generative modeling and decision making. arXiv preprint arXiv:2311.13443 (2023)
 46. Zhou, D., Li, M., Yang, Z., Yang, Y.: DreamRenderer: Taming multi-instance attribute control in large-scale text-to-image models. ICCV (2025)
 47. Zhou, Z., Deng, Y., He, X., Dong, W., Tang, F.: Multi-turn consistent image editing. In: ICCV (2025)
 48. Zhu, H., Liu, H., Fu, B., Wang, Y.: MDE-Edit: Masked Dual-Editing for Multi-Object Image Editing via Diffusion Models. arXiv preprint arXiv:2505.05101 (2025)