

Match-and-Fuse: Consistent Generation from Unstructured Image Sets

Kate Feingold Omri Kaduri Tali Dekel

Weizmann Institute of Science




Figure 1. Given a source set of images (top), depicting shared objects in varied settings (e.g., pose, environment, viewpoint), our method, **Match-and-Fuse**, jointly generates an output set in which the consistency among the shared content is preserved (bottom). The output adheres to the user-provided prompts that describe the target shared content ($\mathcal{P}^{\text{shared}}$), and the scene’s style/theme ($\mathcal{P}^{\text{theme}}$).

Abstract

We present *Match-and-Fuse* – a zero-shot, training-free method for consistent controlled generation of unstructured image sets – collections that share a common visual element, yet differ in viewpoint, time of capture, and surrounding content. Unlike existing methods that operate on individual images or densely sampled videos, our framework performs *set-to-set* generation: given a source set and user prompts, it produces a new set that preserves cross-image consistency of shared content. Our key idea is to model the task as a graph, where each node corresponds to an image and each edge triggers a joint generation of image pairs. This formulation consolidates all pairwise generations into a unified framework, enforcing their local consistency while ensuring global coherence across the entire set. This is achieved by fusing internal features across image pairs, guided by dense input correspondences, without requiring masks or manual supervision. It also allows us to leverage an emergent prior in text-to-image models that

encourages coherent generation when multiple views share a single canvas. *Match-and-Fuse* achieves state-of-the-art consistency and visual quality, and unlocks new capabilities for content creation from image collections. Code and data will be released at our [project page](#).

1. Introduction

Much of our visual experience—how we capture, organize, and interpret the world—is structured not around single images but around sets of them. Photo albums, real-estate listings, product catalogs, and historical archives all offer multiple perspectives on shared content captured at different times or viewpoints. Such sets are richer than individual images yet more flexible than continuous video. Despite major progress in single-image and video generation, Generative AI remains largely underexplored for this fundamental unit of visual communication.

Given a source image set and a user prompt describing the desired shared content, our method produces an output set that adheres to the prompt while preserving the source

layout and the cross-image consistency of shared elements, as illustrated in Fig. 1. These shared elements—semantic regions appearing across multiple images—are kept coherent in identity, appearance, and geometry. Importantly, non-shared regions (e.g., backgrounds) are not forced to align and may vary according to a separate thematic prompt. This capability enables creative workflows of transforming fixed multi-view layouts into coherent edits for product ads, character concept art, film set design, and more.

Achieving consistency across an image set is nontrivial. Unlike videos, which benefit from dense temporal sampling and continuity, image sets challenge this assumption. Without temporal cues such as continuous motion, enforcing visual consistency becomes difficult. The problem is further compounded when the content is deformable – e.g., varying human poses, expressions, or environments – where no stable 3D structure can be inferred. Consequently, existing generative methods lack mechanisms to model shared content without strong spatial or temporal cues.

Our method builds on a pre-trained text-to-image diffusion model in a zero-shot, training-free manner. Our key idea is to model an image set as a graph whose nodes are images and whose edges represent joint pairwise generation. This formulation consolidates all pairwise generations into a joint framework, achieving coherence both locally between image pairs and globally across the entire set. Crucially, this formulation allows us to: (i) harness an emergent prior in text-to-image diffusion models, which exhibit coherent behavior when multiple images are composed within a shared generation canvas; (ii) enhance consistency by incorporating dense pixel correspondences computed from the source images, guiding alignment without requiring masks or manual supervision; (iii) flexibly handle sets of varying sizes without sacrificing resolution. We demonstrate the effectiveness of our framework in producing high-quality results from diverse collections of 2–15 images, ranging from real photographs to hand-drawn sketches.

To summarize, we make the following contributions:

- The first method for unstructured *set-to-set generation*, moving beyond image pairs to entire collections.
- A flexible, automated, training-free mask-free approach, requiring only simple text prompts as input.
- A new metric for evaluating fine-grained cross-image consistency that aligns strongly with human judgments.

2. Related work

Text-to-image controlled generation. T2I models [14, 33, 35, 37, 48] have made a remarkable progress. Numerous methods extended T2I models to go beyond input text and condition the generation on various signals, such as style (e.g., [29, 50]) or various spatial controls: pose, depth, or edge maps [53]. A growing line of work tackles image editing rather than conditioning, including trainable editing models [7, 10, 25] that learn to modify an image according to text or paired examples, as well as training-free editing

approaches [20, 43, 55] that manipulate diffusion trajectories at inference time. However, all these methods operate on individual images and lack the set-level consistency mechanisms required for our task.

Beyond single-image generation.

Storyboard generation. This line of work aims to produce image sequences depicting consistent recurring characters across frames [2, 16, 41, 50], or emphasizing narrative coherence across prompts [19, 22, 27, 54]. Zero-shot methods manipulate attention features through injected correspondences or masks [41, 54], whereas training-based approaches fine-tune diffusion models for improved consistency [19, 22, 27]. Aside from the limited pose control offered by [41], these models rely solely on text and lack spatial control mechanisms necessary for set-to-set generation.

Consistent image-set generation. Related to our work, Edicho [3] addresses pairwise consistent editing, first modifying one image and then transferring the edit to its partner using 2D correspondences to warp intermediate features. Their method is similar to ours in leveraging explicit matches and operating on unstructured image sets, but it remains strictly pairwise: edits are propagated only from a single reference image, causing coherence to degrade for views farther from that anchor. In contrast, [22] enforces consistency by generating all images simultaneously on a single canvas, where a LoRA module [21] is fine-tuned on the multi-image prompt. This design ties the method to a fixed grid, limiting its scalability to larger image sets.

3D and video-based approaches. Distinct from these methods, 3D editing techniques [11, 18, 45, 46] enable manipulation of static multi-view sets through rendered views but operate on 3D representations rather than directly editing 2D images. This imposes restrictive assumptions that limit their applicability to image sets with unknown camera parameters, articulated objects, and varying backgrounds. Similarly, video editing and generation methods (e.g., [17, 44, 51]) are unsuitable for our task, as they assume temporal continuity and rely on motion coherence that does not hold across unordered image sets.

3. Method

The input to our method is an *unstructured* set of N images along with user-provided prompts: $\mathcal{P}^{\text{shared}}$ and $\mathcal{P}^{\text{theme}}$, describing the target shared content and general style or theme, respectively. Our method outputs N images that preserve the source semantic layout while ensuring visual consistency across shared elements.

We build on a pre-trained, frozen, depth-conditioned T2I model. Although designed for single-image generation, these models have been shown to produce image grids when prompted with joint layouts (e.g., “Side-by-side views of...”), establishing cross-image relationships as demonstrated in recent work [22, 38]. However, this emerged capability, which we refer to as the *grid prior*, exhibits several




Figure 2. **Image grid generation vs. our method.** (a) Source image set. (b) Joint generation of two-image grid results in partial consistency, where several regions often remain inconsistent in appearance or semantic meaning (e.g., a dog’s face). (c) Extending this to more images further reduces consistency. (d) Our method leverages this prior yet overcomes its limitations of consistency and scale.

key limitations: (i) it provides only partial consistency in appearance, shape, and semantics (Fig. 2b); (ii) the consistency deteriorates rapidly as more images are composed (Fig. 2b); and (iii) generating a single canvas is bounded by the model’s native resolution, limiting scalability.

Our method leverages the grid prior while overcoming its core limitations (Fig. 2d). Specifically, we model the image set as a **Pairwise Consistency Graph**, comprising all possible two-image grid generations (Sec. 3.2). This allows us to exploit the strong inductive bias of the grid prior, while eliminating its scale limitation. To enhance visual consistency both within each image grid and across grids, we perform joint feature manipulation across all pairwise generations. To this end, we utilize dense 2D correspondences from the source set to automatically identify shared regions – without requiring object masks – and enforce fine-grained alignment. We find that feature-space similarity along these matches correlates strongly with visual coherence, motivating the use of **Multiview Feature Fusion**. We further refine details via **Feature Guidance**, using a feature-matching objective (Sec. 3.4). Our full pipeline is illustrated in Fig. 3.

3.1. Prompt Composition

Given $\mathcal{P}^{\text{shared}}$ and $\mathcal{P}^{\text{theme}}$, we automatically generate detailed *per-image* captions $\mathcal{P}_i^{\text{non-shared}}$ using a Vision-Language Model (VLM). The VLM is instructed to identify shared and non-shared elements, integrate the user-provided ideas while respecting the underlying structure, and describe per-image pose variations. More details in Supplementary Materials (SM).

3.2. Pairwise Consistency Graph

We define a graph $G = (V, E)$, where nodes $V = \{I_i\}_{i=1}^N$ represent images, and edges $E = \{i, j \mid i \neq j, I_i, I_j \in V\}$ connect all distinct image pairs. During the generation, each node is associated with a noisy latent z_i^t , and each edge with latent of a two-image grid $z_{ij}^t = \text{concat}(z_i^t, z_j^t)$. Edges are further assigned concatenated control depth maps

and grid prompts: $\mathcal{P}_{ij} = \text{“Image grid of } [\mathcal{P}^{\text{shared}}]. \text{ Left: } [\mathcal{P}_i^{\text{non-shared}}]. \text{ Right: } [\mathcal{P}_j^{\text{non-shared}}]. \text{”}$

At each generation step, $\{z_{ij}^t\}_{e \in E}$ are constructed from $\{z_i^t\}_{i=1}^N$ and denoised jointly with Multiview Feature Fusion (Sec. 3.3) that encourages consistency across the entire graph by enforcing pre-computed 2D matches M_{ij} :

$$\{z_{ij}^{t-1}\}_{e \in E} = \text{denoise}_{\text{MFF}}(\{z_{ij}^t\}_{e \in E}, \{\mathcal{P}_{ij}\}_{e \in E}, \{M_{ij}\}_{e \in E}) \quad (1)$$

Since every node z_i^t participates in multiple adjacent edges, a graph denoising step yields multiple versions of $z_{i|iij}^{t-1}$. We thus consolidate pairwise latents $\{z_{ij}^{t-1}\}_{e \in E}$ back into image latents $\{z_i^{t-1}\}_{i=1}^N$ by extracting and averaging all $z_{i|iij}^{t-1}$, following [4]. In practice, full graph connectivity is not required, as analyzed in Sec. 4.4.

3.3. Multiview Feature Fusion

The grid prior alone is insufficient for fine-grained alignment within an image pair (Fig. 2). Moreover, denoising separate edges can yield noticeably different appearances, even when a shared image latent is used. To promote pairwise and global consistency, we follow prior work (Sec. 2) and directly manipulate the model’s internal features.

A natural choice for our task is to leverage off-the-shelf 2D correspondences extracted from the source images, which reliably capture shared content through confidence-based filtering. We analyze the model’s feature space and observe that cosine similarity at these matched locations strongly correlates with generation consistency (Fig. 5). Hence, we propose to promote it by increasing the similarity of matched features.

Matches between each image pair are defined as a partial coordinate mapping in feature space, $M_{ij} : \mathcal{C}_i \rightarrow \mathcal{C}_j$, where a coordinate $\mathbf{c} \in \mathcal{C}_i$ corresponds to $\mathbf{c}' = M_{ij}(\mathbf{c}) \in \mathcal{C}_j$ and $\mathcal{C}_i$ contains all matched points (Fig. 4b). We first address pairwise consistency by considering a two-node graph with a single edge. Let $\mathbf{f}_{ij} \in \mathbb{R}^{H \times 2W \times D}$ be a feature map from a model’s forward pass on a grid, interpreted as




Figure 3. **Match-and-Fuse pipeline.** Example for 4 images. In pre-processing, pairwise matches are computed between all inputs, and per-image prompts are generated from the set-level prompts. At each denoising step, noisy image latents form a Pairwise Consistency Graph, whose edges z_{ij}^t are jointly denoised with Multiview Feature Fusion (MFF) and aggregated back into per-image latents z_i^{t-1} by averaging over adjacent edges. The latents are further refined with Feature Guidance via a feature-level matching objective.




Figure 4. **MFF Denoising step.** (a) Two-image grids on all edges are denoised with a frozen DiT. Selected blocks average K, V along adjacent edges into $\bar{f}_i$, which are then fused via source matches (b). Images are fused jointly, illustrated by arrows for $i=1$.

$\mathbf{f}_{ij} = \text{concat}(\mathbf{f}_i, \mathbf{f}_j)$, with $\mathbf{f}_i, \mathbf{f}_j \in \mathbb{R}^{H \times W \times D}$. Given a pooling operator $\mathbf{f}[\mathbf{c}]$ extracting the vector at $\mathbf{c}$, Multiview Feature Fusion (MFF) fuses matched coordinates:

$$\mathbf{f}_i[\mathbf{c}] \leftarrow \frac{1}{2} (\mathbf{f}_i[\mathbf{c}] + \mathbf{f}_j[M_{ij}(\mathbf{c})]), \quad \forall \mathbf{c} \in \mathcal{C}_i. \quad (2)$$

For an N-node graph, Eq. (2) is generalized by aggregating features of each image over its incident edges, followed by joint fusion across all images:

$$\bar{\mathbf{f}}_i = \frac{1}{|\delta(i)|} \sum_{e \in \delta(i)} \mathbf{f}_i^e, \quad \delta(i) = \{e \in E \mid i \in e\} \quad (3)$$

$$\mathbf{f}_i[\mathbf{c}] \leftarrow \frac{1}{N} \left(\bar{\mathbf{f}}_i[\mathbf{c}] + \sum_{j \neq i} \bar{\mathbf{f}}_j[M_{ij}(\mathbf{c})] \right) \quad (4)$$

This stage is illustrated in Fig. 4. See SM for further details.

3.4. Feature Guidance

To align small details, we further perform feature guidance [9, 13, 36], completing each generation step (Fig. 3, right). Specifically, we define an optimization objective over all




Figure 5. **Matched feature similarity vs. visual consistency.** We consider increasingly consistent generation (left to right): random images $\rightarrow$ adding descriptive prompts $\rightarrow$ adding control signals $\rightarrow$ generating in a grid $\rightarrow$ DDIM [40] inversion which reconstructs fully consistent source images. Keys and values differ in scale but follow the same pattern: cosine similarity at matched locations rises with consistency. Dashed lines show the baseline all-to-all similarity of feature maps. Points are averaged over 40 image pairs, all correspondences, blocks, and timesteps.

edges as the distance between matched features, and refine all node latents via gradient descent with $\nabla_{z_i^{t-1}} L_{\text{guide}}$ of:

$$L_{\text{guide}} = \frac{1}{|E|} \sum_{\{i,j\} \in E} \frac{1}{|M_{ij}|} \sum_{\mathbf{c} \in M_{ij}} \|\mathbf{f}_i[\mathbf{c}] - \mathbf{f}_j[M_{ij}(\mathbf{c})]\|_2 \quad (5)$$

The feature maps $\mathbf{f}$ are extracted from an additional DiT forward pass. MFF can be viewed as the analytical solution to this optimization problem operating during inference, while guidance corrects remaining inconsistencies directly in the latent space. Using guidance as the primary driver would be expensive and risk pushing latents off-distribution, so we apply it only as a light refinement. The exact schedule is provided in SM. Because gradients propagate through the model, updates have a wider receptive field than the discrete match locations, improving robustness under sparse correspondences (see Sec. 4.4).




Figure 6. **Quantitative comparison and analysis.** We evaluate subject Consistency (DINO-MatchSim) of Match-and-Fuse as a function (a) of Prompt Adherence (CLIP Score), compared to baselines; (b) Number of Images, compared to baselines and method variant with a fully-connected graph; (c) % of Matches, compared to method variant w/o Feature Guidance. Error bars are SEM.



Figure 7. **Ablation.** Subtraction of each of our method’s components (a-c) degrades consistency, demonstrating their necessity in the full method (d). Each setting is detailed in Sec. 4.4.

	CLIP $\uparrow$	DreamSim $\uparrow$	DINO-MatchSim $\uparrow$
FLUX Kontext	0.65 ± 0.002	0.78 ± 0.004	0.57 ± 0.004
IC-LoRA	0.65 ± 0.001	0.71 ± 0.006	0.65 ± 0.004
FLUX	0.67 ± 0.001	0.76 ± 0.004	0.66 ± 0.004
Edicho	0.65 ± 0.001	0.81 ± 0.004	0.72 ± 0.004
Match-and-Fuse	0.66 ± 0.001	0.85 ± 0.004	0.80 ± 0.003
w/o Guidance	0.66 ± 0.001	0.82 ± 0.004	0.76 ± 0.004
w/o MFF	0.66 ± 0.001	0.83 ± 0.004	0.78 ± 0.003
w/o Pairwise Graph	0.66 ± 0.001	0.82 ± 0.004	0.75 ± 0.004

Table 1. **Quantitative evaluation.** We compare our method to baselines and ablate method components.

4. Experiments

We extensively evaluate Match-and-Fuse. Our implementation builds on FLUX [6] with depth-conditioning [48], and

	Users $\uparrow$	VLM $\uparrow$	Agreement $\uparrow$
Kontext	88%	82%	DreamSim 84.3
IC-LoRA	90%	92%	VLM 84.9
FLUX	92%	94%	DINO-MatchSim 91.4
Edicho	83%	78%	

Table 2. (a) **User Study and VLM evaluation.** 2AFC voting demonstrates the preference of humans and VLM of our method over all baselines. (b) **Agreement between metrics with human judgments.** % of samples for which the winner according to a higher metric agrees with the users’ majority vote. Average across comparisons and baselines.

uses RoMA [12] for matching. Full details are in SM.

4.1. Evaluation Setup

Evaluating our task requires image sets that share content across diverse scenarios, along with a metric that captures fine-grained cross-image consistency. Since no existing benchmark or metric supports this setting, we introduce a setup tailored to consistent image-set generation.

Evaluation set. We curate a benchmark with 400 edits for 149 diverse distinct image sets, 3–15 images in each, combining all sets from [24] and [34] with frames sampled from 3D datasets [5, 18, 26], keyframes of a few publicly sourced videos, and sketched storyboards generated by ChatGPT. Data will be publicly released.

Metrics. We quantitatively evaluate our method, beginning with fine-grained cross-image consistency. Previous work relies on *global* visual-similarity metrics to assess set consistency—for example, [41] uses DreamSim [15] on object-masked regions. However, as our user study (Tab. 2b) shows, such global metrics are not ideal. To address this, we introduce *DINO-MatchSim*. For each image pair, we compute patch-level nearest-neighbor correspondences NN_{ij} between DINOv3 [39] feature maps from the source images, and measure similarity at the corresponding output locations:

$$S_{ij} = \frac{1}{|NN_{ij}|} \sum_{(p,q) \in NN_{ij}} \cos(\tilde{F}_i(p), \tilde{F}_j(q)), \quad (6)$$

where $\tilde{F}_i$ denotes the i^{th} output feature map. The aver-




Figure 8. **Qualitative results.** Match-and-Fuse generates consistent content of rigid and non-rigid shared elements, single and multi-subject, with shared or varying background, preserving fine-grained consistency in textures, small details, and typography. Notably, it can generate consistent long sequences (last row). See SM for full sets of results.

age similarity across all image pairs is *DINO-MatchSim*. For prompt adherence, we report the average CLIP score over the set. To assess human preference, we conduct a large-scale user study on a random subset of the benchmark, comparing Match-and-Fuse to a baseline in a randomized two-alternative forced-choice (2AFC) format. Finally, we include a VLM-based evaluation using GPT-5 [31], which received the same content as the human raters.

4.2. Qualitative results

Sample qualitative results of our method are shown in Figs. 1, 8 and 9, demonstrating that it performs robustly under non-rigid pose variations, diverse viewpoints, and

partial occlusions, on single or multi-object sets (e.g. two girls). The output sets maintain fine-grained consistency in textures (e.g. crochet, balloon and can print), small elements (e.g. kitchen items), and typography (e.g. the can title). Notably, Match-and-Fuse is able to handle big number of images (e.g. $N = 13$ in Fig. 8, last row).

4.3. Comparisons

Since no existing method is designed for our set-to-set generation task, we compare Match-and-Fuse to the closest baselines: (1) IC-LoRA [22], which fine-tunes LoRA modules for specific context types to produce text-driven image sets with shared elements. We use their most relevant pub-




Figure 9. **Qualitative comparisons.** Match-and-Fuse (Ours) generates image sets with the highest consistency compared to our baselines FLUX [6], FLUX Kontext [25], IC-LoRA [22], and Edicho [3]. See Sec. 4.3 for comments on each result and SM for full image sets.

lic checkpoint and add spatial ControlNet conditioning for comparability. (2) FLUX Kontext [25], an image-to-image editing model. While it provides consistent edits over multiple turns on a single image, it lacks an explicit mechanism for achieving consistency across a single-turn image set editing. (3) FLUX [6] with ControlNet [48] conditioning, which serves as a baseline for assessing independent generation using the same base model as our method. (4) Edicho [3], the closest baseline to our method, combines depth-conditioned generation with explicit correspondences, but is limited to an image pair editing through a single reference. Additional implementation details are in Sec. A.2.


Figure 9 shows representative qualitative results. FLUX [6] produces inconsistencies in both coarse and fine details. FLUX Kontext [25] has low prompt adherence (e.g., a toy remains animal-like), and often distorts object structure (e.g. vase, dog). IC-LoRA [22] achieves partial consistency, with coherence often restricted to subsets (vase). Its realism and fidelity are notably lower than FLUX. Edicho [3] performs best among baselines but still shows noticeable inconsistencies (dog), as its one-to-all warping enforces consistency only with the first image, making the choice of an anchor ambiguous and degrading coherence across views. In contrast, Match-and-Fuse (bottom row) maintains high

image quality while substantially improving structural and fine-grained consistency. We additionally compare to the closed-source Nano Banana API [10] and find that it does not solve our task (SM).

We quantitatively evaluate all methods using the metrics in Sec. 4.1. Figure 6a plots Consistency (DINO-MatchSim) versus Prompt adherence (CLIP score). As an upper bound, we report DINO-MatchSim on the source sets, which are consistent by definition. Our method achieves the highest consistency among baselines, approaching the source score, while maintaining a comparable CLIP score. Since DINO-MatchSim evaluates consistency relative to the source structure, methods that distort shape, such as FLUX Kontext, score lowest. Tab. 1 further shows that Match-and-Fuse attains the highest DreamSim score. We complement these results with a user study and VLM-based evaluation. As shown in Tab. 2a, participants and VLM preferred our method over all baselines.


4.4. Analysis & Ablations

Ablation. We ablate each component in Fig. 7 and Tab. 1. *W/o Pairwise Consistency Graph*, pairwise steps become single-image predictions, with MFF and Guidance applied per image. Correspondences alone cannot align appear-



ances, leading to identity drift. In *w/o Multiview Feature Fusion*, pairwise updates use only the grid prior; latent versions aggregated across edges diverge more easily, reducing consistency. *Omitting Feature Guidance* at each step leads to misaligned fine-grained details.

Number of images. A fully connected graph has $O(N^2)$ edges. To maintain scalability, we limit each node degree to 4 random neighbors, yielding full connectivity for $N \leq 5$ and increasing sparsity thereafter. This keeps runtime linear while preserving cross-image communication. We use this setup in all our experiments. Figure 6b shows consistency as a function of N . For each N -image set, we eval-



uate subsets of the first $N' \in [2, N]$ images and report averages over all N' . Our method is slightly below its fully connected variant but remains similar in performance. Although Match-and-Fuse degrades with N it remains more consistent for 9 images than baselines do for 2.

Match sparsity. We evaluate robustness to reduced correspondences by randomly subsampling $x\%$ of matches. As shown in Fig. 6c, DINO-MatchSim remains high even under strong sparsification, whereas the variant without Feature Guidance drops more rapidly—highlighting its role in maintaining consistency with sparse matches.

4.5. Extended Applications

As shown in Fig. 10a, our method generalizes to sketched source sets, enabling consistent visualization of *storyboards*. In Fig. 10b, we further demonstrate *localized editing* by combining our method with FlowEdit [23]. We provide default integration parameters in the SM; however, as is a known limitation of FlowEdit, achieving the desired balance between structure preservation and appearance change often requires mild per-edit hyperparameter tuning. Automating this selection is left for future work.

5. Discussion and Conclusions

We introduced Match-and-Fuse, the first method for controlled generation from unstructured image collections. Extensive experiments demonstrate that it outperforms strong baselines and enables diverse creative applications, advancing a largely underexplored yet fundamental visual modality. Several limitations remain for future work. Our performance depends on the quality and density of pixel correspondences, which may result in inconsistencies in largely unmatched or ambiguous regions (e.g., disocclusions or symmetries; see Fig. 11). It also relies on the ability of the base model to preserve the source conditioning depth maps, which may occasionally deviate from the source layout, depending on the target prompt and generative prior of the T2I model (e.g. body and hand poses of the girls in Fig. 8). We believe these insights lay a path toward future extensions, such as video collections and foundation models for set-to-set generation.

Acknowledgments

We thank Danah Yatim for her valuable comments. We thank Vladimir Kulikov for insightful discussions on FlowEdit. We thank Qingyan Bai and the other authors of Edicho for conducting qualitative comparisons with their method.

References

- [1] Yuval Alaluf, Daniel Garibi, Or Patashnik, Hadar Averbuch-Elor, and Daniel Cohen-Or. Cross-image attention for zero-shot appearance transfer, 2023. 13
- [2] Omri Avrahami, Amir Hertz, Yael Vinker, Moab Arar, Shlomi Fruchter, Ohad Fried, Daniel Cohen-Or, and Dani Lischinski. The chosen one: Consistent characters in text-to-image diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, New York, NY, USA, 2024. Association for Computing Machinery. 2
- [3] Qingyan Bai, Hao Ouyang, Yinghao Xu, Qiuyu Wang, Ceyuan Yang, Ka Leong Cheng, Yujun Shen, and Qifeng Chen. Edicho: Consistent image editing in the wild. In *arXiv preprint arXiv:2412.21079*, 2024. 2, 7, 11
- [4] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. *arXiv preprint arXiv:2302.08113*, 2023. 3
- [5] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. *ICCV*, 2021. 5, 12
- [6] Black-Forest. Flux: Diffusion models for layered image generation. <https://github.com/black-forest-labs/flux>, 2024. Accessed: 2024-09-24. 5, 7, 11, 13
- [7] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023. 2
- [8] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22560–22570, 2023. 13
- [9] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models, 2023. 4
- [10] Google DeepMind. Gemini 2.5 flash image (“nano banana”) model/api, 2025. Accessible via Google Gemini API. 2, 7
- [11] Jiahua Dong and Yu-Xiong Wang. Vica-nerf: View-consistency-aware 3d editing of neural radiance fields. In *NeurIPS*, 2023. 2
- [12] Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. Roma: Robust dense feature matching, 2023. 5, 11, 12
- [13] Dave Epstein, Allan Jabri, Ben Poole, Alexei A. Efros, and Aleksander Holynski. Diffusion self-guidance for controllable image generation. 2023. 4
- [14] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yan-nik Marek, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis, 2024. 2
- [15] Stephanie Fu, Netanel Tamir, Shobhita Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip Isola. Dreamsim: Learning new dimensions of human visual similarity using synthetic data. In *Advances in Neural Information Processing Systems*, pages 50742–50768, 2023. 5
- [16] Daniel Garibi, Shahar Yadin, Roni Paiss, Omer Tov, Shiran Zada, Ariel Ephrat, Tomer Michaeli, Inbar Mosseri, and Tali Dekel. Tokenverse: Versatile multi-concept personalization in token modulation space, 2025. 2
- [17] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. *arXiv preprint arxiv:2307.10373*, 2023. 2, 13
- [18] Ayaan Haque, Matthew Tancik, Alexei Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 2, 5, 12
- [19] Huiguo He, Huan Yang, Zixi Tuo, Yuan Zhou, Qiuyue Wang, Yuhang Zhang, Zeyu Liu, Wenhao Huang, Hongyang Chao, and Jian Yin. Dreamstory: Open-domain story visualization by llm-guided multi-subject consistent diffusion, 2024. 2
- [20] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. 2022. 2
- [21] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 2
- [22] Lianghua Huang, Wei Wang, Zhi-Fan Wu, Yupeng Shi, Huanzhang Dou, Chen Liang, Yutong Feng, Yu Liu, and Jingren Zhou. In-context lora for diffusion transformers. 2024. 2, 6, 7, 11
- [23] Vladimir Kulikov, Matan Kleiner, Inbar Huberman-Spiegelglas, and Tomer Michaeli. Flowedit: Inversion-free text-based editing using pre-trained flow models. *arXiv preprint arXiv:2412.08629*, 2024. 8, 11
- [24] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. 2023. 5, 12
- [25] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. 2, 7, 11
- [26] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, and Zhaoqiang Lv. Neural 3d video synthesis from multi-view video, 2022. 5, 12
- [27] Chang Liu, Haoning Wu, Yujie Zhong, Xiaoyun Zhang, Yanfeng Wang, and Weidi Xie. Intelligent grimm - open-ended visual storytelling via latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6190–6200, 2024. 2

- [28] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jianjun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations, 2022. 11
- [29] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023. 2
- [30] OpenAI. Gpt-4o system card. Technical report, OpenAI, 2024. Technical report. 11
- [31] OpenAI. Gpt-5 system card. Technical report, OpenAI, 2025. Technical report. 6, 12
- [32] Prolific. Prolific. <https://www.prolific.com/>, 2024. 12
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022. 2
- [34] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 5, 12
- [35] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. 2
- [36] Yujun Shi, Chuhui Xue, Jiachun Pan, Wenqing Zhang, Vincent YF Tan, and Song Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. *arXiv preprint arXiv:2306.14435*, 2023. 4
- [37] Zhan Shi, Xu Zhou, Xipeng Qiu, and Xiaodan Zhu. Improving image captioning with better use of captions, 2020. 2
- [38] Chaehun Shin, Jooyoung Choi, Heeseung Kim, and Sungroh Yoon. Large-scale text-to-image model with inpainting is a zero-shot subject-driven image generator. 2024. 2
- [39] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. Dinov3, 2025. 5, 12
- [40] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv:2010.02502*, 2020. 4
- [41] Yoad Tewel, Omri Kaduri, Rinon Gal, Yoni Kasten, Lior Wolf, Gal Chechik, and Yuval Atzmon. Training-free consistent text-to-image generation. *ACM Transactions on Graphics (TOG)*, 43(4):1–18, 2024. 2, 5, 13
- [42] Yoad Tewel, Rinon Gal, Dvir Samuel, Yuval Atzmon, Lior Wolf, and Gal Chechik. Add-it: Training-free object insertion in images with pretrained diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. 13
- [43] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. 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 (CVPR)*, pages 1921–1930, 2023. 2
- [44] Jiangshan Wang, Yue Ma, Jiayi Guo, Yicheng Xiao, Gao Huang, and Xiu Li. Cove: Unleashing the diffusion feature correspondence for consistent video editing. *arXiv preprint arXiv:2406.08850*, 2024. 2
- [45] Yuxuan Wang, Xuanyu Yi, Zike Wu, Na Zhao, Long Chen, and Hanwang Zhang. View-consistent 3d editing with gaussian splatting. *arXiv preprint arXiv:2403.11868*, 2024. 2
- [46] Jing Wu, Jia-Wang Bian, Xinghui Li, Guangrun Wang, Ian Reid, Philip Torr, and Victor Prisacariu. GaussCtrl: Multi-View Consistent Text-Driven 3D Gaussian Splatting Editing. *ECCV*, 2024. 2
- [47] Saining Xie and Zhuowen Tu. Holistically-nested edge detection, 2015. 11
- [48] XLabs-AI. Flux-controlnet collections. <https://huggingface.co/XLabs-AI/flux-controlnet-collections>, 2024. Accessed: 2024-11-25. 2, 5, 7, 11
- [49] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2, 2024. 11
- [50] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. 2023. 2
- [51] Hangjie Yuan, Shiwei Zhang, Xiang Wang, Yujie Wei, Tao Feng, Yining Pan, Yingya Zhang, Ziwei Liu, Samuel Albanie, and Dong Ni. Instructvideo: Instructing video diffusion models with human feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6463–6474, 2024. 2
- [52] zer0int. Longclip-l-diffusers. <https://huggingface.co/zer0int/LongCLIPDiffusers>, 2024. Model on Hugging Face. Accessed: 2025-10-18. 11
- [53] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 2
- [54] Yupeng Zhou, Daquan Zhou, Ming-Ming Cheng, Jia-ashi Feng, and Qibin Hou. Storydiffusion: Consistent self-attention for long-range image and video generation. *NeurIPS 2024*, 2024. 2
- [55] Tianrui Zhu, Shiyi Zhang, Jiawei Shao, and Yansong Tang. Kv-edit: Training-free image editing for precise background preservation. *arXiv preprint arXiv:2502.17363*, 2025. 2

Appendix

A. Implementation details

A.1. Method

Preprocessing. Prompt construction with a VLM (Sec. 3.1) is used as an automated, scalable alternative to manual prompting. We instruct GPT-4o [30] via the system prompt (Fig. A.3) to generate $P_i^{\text{non-shared}}$ and enrich $\mathcal{P}^{\text{shared}}$ with creative artistic details (e.g., colors, textures). Figure A.1 exemplifies VLM input and output. Source 2D correspondences and shared content are identified using RoMA [12] predictions filtered by a confidence threshold $c > 0.05$.

General. All results are generated using FLUX [6] with $t = 25 \dots 0$ steps. We apply ControlNet [48] complementing DepthAnythingV2 [49] depth with HED edge maps [47] with equal weights, and linearly anneal control strength $w_{\text{ctrl}}^t = \text{linear_anneal}(1, 0, t)$. This enforces correct layout while preventing over-constraining control.

In *Pairwise Consistency Graph* (Sec. 3.2), independent Gaussian noise is sampled for each node. Two-image grids are concatenated horizontally or vertically, depending on the image aspect ratio. For $\mathcal{P}_{ij}$, any freely written prefix prompting a side-by-side layout yields similar results. Fig. A.2 shows the concept of image versions aggregated along graph edges and unified with MFF.

Multiview Feature Fusion (Sec. 3.3) and *Feature Guidance* (Sec. 3.4) are applied to keys and values before RoPE. MFF operates on all double-stream blocks for $t_d^{\text{MFF}} > 3$ and the last three single-stream blocks for $t_s^{\text{MFF}} > 15$. Feature Guidance objective is defined on all blocks for $t_a^{\text{guide}} > 10$. Latents are updated using Adam with a linearly annealed learning rate $\alpha^t = \text{linear_anneal}(0.016, 0.002, t)$ and gradient accumulation per edge. In Eq. (4), undefined $M_{ij}(\mathbf{c})$ terms are omitted from averaging.

Localized editing. FlowEdit [23] is an inversion-free method that enables background-preserving editing by traversing a path between source and target image distributions with v -predictions $V = V^{\text{tar}}(z^{\text{tar}}) - V^{\text{src}}(z^{\text{src}})$. To incorporate it into our method, we redefine V_{ij} implied in $\text{denoise}_{\text{MFF}}(\cdot)$ in Eq. (1) as $V_{ij} = V_{ij}^{\text{tar}} - V_{ij}^{\text{src}}$, where

$$V_{ij}^\alpha := V_{ij}^\alpha(\{z_{ij}^\alpha\}_{e \in E}, \{\mathcal{P}_{ij}^\alpha\}_{e \in E}, \{M_{ij}\}_{e \in E}), \quad (\text{A.1})$$

$$\alpha \in \{\text{src}, \text{tar}\} \quad (\text{A.2})$$

Similar to [28], FlowEdit relies on a single starting timestep parameter n_{max} to balance structure preservation and appearance deviation. We ease this trade-off by adopting spatial control with an earlier start step. The used hyperparameters are: $T = 25, n_{\text{max}} = 24, n_{\text{min}} = 10, w_{\text{ctrl}} = 0.5, t_d^{\text{MFF}} > 20$. Feature Guidance is omitted.

Despite the improved balance, some trade-off between structure and appearance remains. We therefore suggest $w_{\text{ctrl}} = 0.8$ when structure is prioritized over editability.



Figure A.1. **Completed captions.** The VLM provides detailed per-image captions aligned with $\mathcal{P}^{\text{theme}}$ and expands on shared content $\mathcal{P}^{\text{shared}}$, offering an alternative to manual prompting.



Figure A.2. **Image versions.** A grid of I_1 with varying neighbors produces pairs that are mutually consistent but exhibit different appearances across pairs. This appearance divergence is prevented every denoising step in Pairwise Consistency Graph by a graph-wide MFF and version averaging.

A.2. Baselines

For *IC-LoRA* [22], we use the most general Film Storyboard checkpoint with a LoRA weight of 0.8. *IC-LoRA* and *FLUX* [6] are run with identical control signals as the method. *FLUX* and *FLUX Kontext* [25] share the same random seed across images, and *FLUX Kontext* is prompted by the instruction "Turn the <object> into ...". For *Edicho* [3], we use the first image as a reference for the rest. Metrics in Tabs. 1 and 2 are computed with its public inpainting implementation available at publication, using the shared region masks from evaluation. The results in Fig. 9 with the ControlNet pipeline were provided by the authors.

A.3. Evaluation Setup

CLIP score is computed with LongCLIP-L [52] and prompts "[$\mathcal{P}^{\text{shared}}$]. [$\mathcal{P}_i^{\text{non-shared}}$].".

The *User Study* evaluating consistency was conducted on 80 random edits from our benchmark using the Prolific

You will receive:

- * A set of images of the source scene with a shared object;
- * A short Object Edit description (e.g., "a two-colored robot");
- * A short Background Edit description (e.g., "autumn", "cyberpunk");

Your task is to provide:

- * Source Object Prompt;
- * Edited Object Prompt;
- * Per Image Source Pose and Background Prompts;
- * Per Image Edit Pose and Background Prompts;

To create the prompts, follow the rules below.

In both Source Object Prompt and Edited Object Prompt:

Focus on the appearance description (colors of elements, shape, texture, details and their position on the object), include as many visual details that fit into the prompt length. Make sure not to include pose or background descriptions.

Source Object Prompt:

Provide a detailed description of the shared source object without considering the edit.

Edited Object Prompt:

Provide a detailed description of the shared edited object that maintains structure and shape but changes the appearance. Use the provided Object Edit as a cue that hints at the new identity or appearance. Be creative: specify new elements, their colors, textures, and other visual details that fit the requested edit. Do not reference the original appearance, make the description stand-alone.

Per Image Source Pose and Background Prompts (rules that apply to each image):

Refer to the common visual element with a consistent and general name: e.g. "the object", "the person", "etc".

Provide a detailed description of the pose of the object (if it's located at the side, rotated, close-up, etc.), especially pose details that are not standard or different than in other images.

Provide a detailed description of the background: list all the surrounding objects and their locations, their appearance, and the environment.

Per Image Edit Pose and Background Prompts (rules that apply to each image):

Provide prompts with a similar sentence patterns to the corresponding Per Image Source Pose and Background Prompts, but focus on editing the background according to the Background Edit description. Don't change the types of any objects and don't add any new objects, just change the appearance description of the existing ones according to theme. Make sure to mention the theme name at least once. Keep the same pose description and a general reference to the common visual element as in the corresponding Per Image Source Pose and Background Prompts.

Rules to follow strictly for all prompts:

- * Each prompt must have between 20-50 words, not more than that!
- * Don't start with - "The scene now..." or "Added to the scene" "Scene has transformed", each prompt must be standalone and understandable without access to other prompts.
- * Don't include any emotion-related details, things that cannot be visualized, and meaningless phrases like: "cozy charm", "playful look", "emphasizing the setting typical of", "evoking a sense of", "sitting «dramatically»", "accentuating its look", "warm setting", etc (don't use any of such phrases!);
- * All listed rules are equally important and must be followed strictly.

Output format - a dictionary with keys:

- * "source_object_prompt" (str, Source Object Caption),
- * "edited_object_prompt" (str, Edited Object Caption),
- * "per_image_pose_background" (list[str], Per Image Source Pose and Background Prompts, the number of items corresponds to the number of provided images).
- * "per_image_edit_pose_background" (list[str], Per Image Edit Pose and Background Prompts, the number of items corresponds to the number of provided images).

Figure A.3. **VLM instructions for prompt construction.** System prompt. Source prompts are used for localized editing with FlowEdit.

[32] platform. Participants were shown the source image set, $\mathcal{P}^{shared}$, outputs from Match-and-Fuse and a baseline in a randomized 2AFC setup (see the interface in Fig. A.4). They answered: *"The source image set contains shared foreground object(s). In which image set, A or B, are these objects edited more consistently across images?"*. We collected 1,600 judgments from 40 participants, with 5 respondents per comparison.

The automatic *VLM evaluation* was conducted on the full benchmark using GPT-5 [31] with the system prompt: *"You will receive an image to evaluate along with an evaluation question. Output A or B."* Each comparison used the same inputs as human participants: the compiled image (source set with randomized generations A and B), $\mathcal{P}^{shared}$, and the user-study question.

Our *DINO-MatchSim* metric quantifies fine-grained cross-image consistency in appearance and structure. Generated images are compared in the DINOv3 [39] feature space at fixed correspondences NN_{ij} —nearest neighbors from background-removed source images. DINOv3 is used for its strong and versatile representation and to ensure independence from the 2D matches used by our method. In practice, pairwise cosine similarities S_{ij} (see Eq. (6)) are averaged into $\bar{S}$, then transformed as: $DINO-MatchSim = \exp((\bar{S} - 1)/\tau)$ with $\tau = 0.6$. This mapping amplifies consistency differences by compressing negative similarities into small scores ($\approx 0-0.2$) and expanding positive similarities into a wider interval ($\approx 0.2 - 1$).

Table A.1 shows the distribution of the evaluation set.

Shared Content	# Sets	# Edits	Source Datasets
Subject-driven	131	357	CustomConcepts101 [24]
			DreamBooth [34]
			DyNeRF [26] Instruct-NeRF2NeRF [18] Mip-NeRF [5]
Subject & Background	11	27	royalty-free videos
Storyboard sketch	7	16	ChatGPT-generated

Table A.1. **Benchmark composition.** A distribution of sets and edits by the types of shared content.

B. Runtime & Memory

During the extraction of 2D matches in preprocessing, RoMA has a runtime of 198.8 ms per image pair on an RTX6000 GPU [12]. In Fig. B.5, we report the per-image runtime of our full method as a function of N , progressively adding its components. As also discussed in Sec. 4.4, Match-and-Fuse maintains linear runtime for $N \geq 5$ by constraining graph connectivity to a node degree of 4. Since Feature Guidance noticeably increases runtime, it can be omitted when efficiency is critical, trading fine-grained consistency for speed.


In terms of memory constraints, we can generate up to 20 images using the full method, or 22 without Feature Guidance. Figure B.6 illustrates how increasing the set size to the maximum affects consistency.

All measurements are conducted on an NVIDIA A100 GPU with an image resolution of 512×512 .


Definitions:

Consistency — same object identity and its appearance details across views.

Inconsistency — identity and its appearance details mismatch between views.



Edit prompt: a book titled 'Enchanted Waters'



The Source image set (top row) contains **shared foreground object(s)**.

In which image set, A or B, are these shared objects edited **more consistently across the images?**

Assess only the shared foreground object(s), **ignore the background**.*

☐ A (above)

☐ B (below)

Submit Answers and Continue

Figure A.4. **User study interface.** Example of an interface with a question comparing Match-and-Fuse to FLUX [6].




Figure B.5. **Runtime per Image vs. Set Size.** Match-and-Fuse maintains approximately constant runtime per image for $N \geq 5$.

C. Adapting Extended Attention for Our Task

We explore the additional qualitative baseline of prompting consistency through extended attention, widely used in UNets for cross-sample appearance sharing [1, 8, 17, 41], and

observe that in FLUX, this approach leads to artifacts and lacks semantic reasoning. We share our findings to inform future research.

In DiTs, Rotary Position Embeddings (RoPE) are applied before self-attention in each block, thus extended keys of the image modality are:

$$K_{img} = [PE_{self}(K_{I_{self}}), PE_2(K_{I_2}), \dots, PE_N(K_{I_N})] \quad (C.3)$$

The positional embeddings (PEs) with a 2D token coordinate must be defined for all keys. [42] showed that in FLUX they strongly influence the spatial object placement. However, since extra keys are not truly present in the generated image plane, the choice of their positions is ambiguous. We test three different strategies: (i) Re-using PE_{self} . This variant yields identical generations that completely disregard the control signal (Fig. C.7a); (ii) Continuing positions analogous to PEs of spatially concatenated images in the




Figure B.6. **Consistency vs. Set Size.** The same 3 images when generated as part of increasingly bigger sets of size N . Low and middle frequencies remain consistent, while high frequencies disagree more with higher N .

grid generation. Despite this similarity, without a prompt indicating a joint canvas, this setting leads to artifacts and fails to achieve consistent outputs (Fig. C.7b); (iii) We extend only with keys that have a source 2D match and warp PEs of K_{Iself} such that matched keys share the same position. Additionally, extra keys are scaled by k to control their temperature in softmax (Fig. C.7c). With $k = 1$, generations approach consistency at the cost of severe quality degradation. Reducing k alleviates the artifacts, but com-

promises consistency before image quality is fully restored. In contrast, our method achieves consistency while avoiding artifacts (Fig. C.7d). These findings demonstrate a strong connection between features that share the same position in RoPE. Exploring this behavior may open future research directions.




Figure C.7. **Extended Attention in DiTs.** (a-c) Settings varying by PEs of extra keys, illustrated in 1D for simplicity below each setting, (d) Ours. (a) Re-using PEs leads to duplicated generations. (b) Extending PEs leads to low consistency and artifacts. (c) Warping PEs by source matches demonstrates a severe quality-consistency tradeoff.