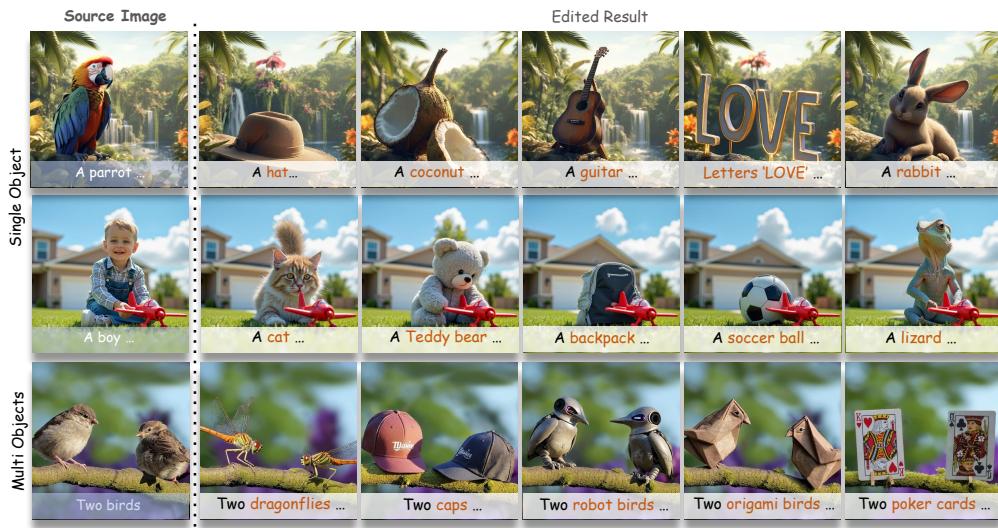


EDITANYSHAPE: SHAPE-AWARE IMAGE EDITING VIA TRAJECTORY-GUIDED REGION CONTROL

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review



024 Figure 1: We propose **EditAnyShape**, a training- and mask-free image editing framework that
025 excels at prompt-driven shape transformation. Our method enables flexible modification of arbitrary
026 object shapes while strictly maintaining non-target content. The examples demonstrate both single-
027 object and multi-object cases involving significant shape transformation.

ABSTRACT

032 While recent flow-based image editing models demonstrate general-purpose cap-
033 abilities across diverse tasks, they often struggle to specialize in challenging
034 scenarios—particularly those involving large-scale shape transformations. When
035 performing such structural edits, these methods either fail to achieve the intended
036 shape change or inadvertently alter non-target regions, resulting in degraded back-
037 ground quality. We propose **EditAnyShape**, a training-free and mask-free frame-
038 work that supports precise and controllable editing of object shapes while strictly
039 preserving non-target content. Motivated by the divergence between inversion and
040 editing trajectories, we compute a **Trajectory Divergence Map (TDM)** by com-
041 paring token-wise velocity differences between the inversion and denoising paths.
042 The TDM enables precise localization of editable regions and guides a **Scheduled**
043 **KV Injection** mechanism that ensures stable and faithful editing. To facilitate a
044 rigorous evaluation, we introduce **ReShapeBench**, a new benchmark comprising
045 120 new images and enriched prompt pairs specifically curated for shape-aware
046 editing. Experiments demonstrate that our method achieves superior editability
047 and visual fidelity, particularly in tasks requiring large-scale shape replacement.

1 INTRODUCTION

048 Recent advances in generative models have greatly expanded the scope of image editing, enabling
049 more controllable and realistic modifications across diverse scenarios. Image editing methods based
050 on diffusion (Cao et al., 2023; Tumanyan et al., 2023; Feng et al., 2025) and flow models (Lipman
051 et al., 2022; Labs, 2024; Labs et al., 2025; Kulikov et al., 2025) have demonstrated considerable

success in general tasks, yet they often fail when faced with complex, large-scale shape transformations. These models can struggle to modify an object’s structure as intended or may inadvertently alter background regions, which degrades the overall image quality. This limitation indicates a critical gap in their ability to perform precise structural edits while maintaining the integrity of unedited content.

The primary cause for this limitation lies in the inadequacy of existing region control strategies (Zhu et al., 2025; Cao et al., 2023). Methods that rely on external binary masks are often too rigid and struggle with the fine details of object boundaries. Alternatively, strategies that use cross-attention maps to infer editable regions are frequently unreliable, as these maps can be noisy and inconsistent. While unconditional Key-Value (KV) injection can preserve background structure, it lacks selectivity and tends to suppress the intended edits (Avrahami et al., 2025; Wang et al., 2024). We argue that a breakthrough requires a new approach: one that derives the editable region dynamically from the editing process itself by analyzing how the model’s behavior shifts between the source and target conditions.

To address these challenges, we propose **EditAnyShape**, a training- and mask-free framework for precise and controllable shape editing. As illustrated in Figure 2, the core innovation of our pipeline is the **Trajectory Divergence Map (TDM)**. The TDM is generated by computing the token-wise difference between the denoising velocity fields of the source and target prompts. This map accurately localizes the regions intended for editing, which in turn guides a selective KV injection mechanism to ensure that modifications are applied precisely where needed while preserving the background.

However, directly applying TDM-guided injection across all denoising timesteps is suboptimal because the TDM can be unstable in the early, high-noise stages of the process. We therefore introduce a **Scheduled KV Injection** strategy that adapts its guidance throughout the denoising process. As visualized in Figure 2, this staged approach first performs unconditional KV injection to stabilize the initial trajectory, and only then applies TDM-guided editing once a coherent latent structure has formed. This scheduling and staged editing pipeline ensures a more robust and faithful editing outcome compared to a direct application.

To validate our approach, we introduce **ReShapeBench**, a new benchmark with paired images and refined text prompts specifically designed for evaluating large-scale shape modifications. Beyond this new dataset, we further evaluate EditAnyShape on the public *PIE-Bench* (Ju et al., 2023) to ensure generalizability. EditAnyShape achieves state-of-the-art performance on both benchmarks, demonstrating superior background preservation, text–image alignment, and overall visual quality, confirming its effectiveness in both shape-aware and general editing tasks.

Our primary contributions are summarized as follows:

- A novel and training-free editing framework, **EditAnyShape**, that utilizes a **Trajectory Divergence Map (TDM)** to achieve precise, large-scale shape transformations while preserving background content.
- A trajectory-guided **scheduled injection strategy** that improves editing stability by adapting the guidance mechanism throughout the denoising process.
- A new benchmark, **ReShapeBench**, designed for the systematic evaluation of shape-aware image editing methods.

2 RELATED WORK

Region-Specific Image Editing. A central challenge in image editing is localizing modifications to specific regions (Barnes et al., 2009; Zhang et al., 2024; Liu et al., 2025; Huang et al., 2025). Early methods often relied on explicit user-provided masks to delineate editable areas (Lugmayr et al., 2022; Avrahami et al., 2023; Chen et al., 2023; Xiong et al., 2025; Wan et al., 2024). While effective for certain tasks, this approach requires manual annotation, limiting its applicability. To address this, subsequent work explored techniques to infer editable regions directly from text prompts. Methods such as Prompt-to-Prompt (Hertz et al., 2022) and Plug-and-Play (Tumanyan et al., 2023) manipulate cross-attention maps to associate textual tokens with spatial areas, enabling localized edits without explicit masks. Other approaches, such as DiffEdit (Couairon et al., 2022), generate

108 a mask by computing differences between diffusion model predictions conditioned on source and
 109 target prompts. However, attention-based localization can be imprecise and unstable, especially
 110 during large-scale shape transformations where object boundaries change significantly (Pang et al.,
 111 2024; Cao et al., 2023). In contrast, EditAnyShape provides a training-free and mask-free method
 112 for identifying editable regions directly from the model’s behavior, avoiding the need for external
 113 masks or noisy attention maps.

114
 115 **Structure Preservation via Inversion and Feature Reuse.** Preserving non-target regions is
 116 equally critical for high-fidelity editing, and this is closely tied to the quality of the model’s in-
 117 version process. For diffusion models, significant research has focused on improving DDIM in-
 118 version (Song et al., 2020) to better reconstruct a source image from noise. Previous works like
 119 null-text inversion (Mokady et al., 2023) and optimization-based methods (Wallace et al., 2023) aim
 120 to reduce the discrepancy between the reconstruction and editing trajectories. With the shift toward
 121 flow-based models, inversion fidelity has become even more important due to their deterministic
 122 nature. RF-Inversion (Rout et al., 2024) formulates the inversion process as a dynamic optimal control
 123 problem, while RF-Solver (Wang et al., 2024) achieve more accurate reconstructions by incorpo-
 124 rating higher-order derivative information. Beyond improving inversion, another line of work focuses
 125 on explicitly reusing modules or features from the source image’s generation process (Zheng et al.,
 126 2024; Ma et al., 2025; Yan et al., 2025). Techniques based on Key-Value (KV) caching (Zhu et al.,
 127 2025; Avrahami et al., 2025) or feature injection (Wang et al., 2024; Feng et al., 2025) enforce
 128 structural consistency by propagating source-image features into the new generation process. In
 129 contrast to prior methods that rely on simple heuristics, EditAnyShape employs a trajectory-guided
 130 scheduled injection strategy to achieve more precise, content-aware control.

131 3 METHODOLOGY

132 Our goal is to enable precise object shape-aware editing while strictly preserving the background.
 133 Motivated by the limitations of existing region control strategies and the need for a more adaptive
 134 mechanism, we introduce Trajectory Divergence Map (TDM) that quantifies token-wise semantic
 135 deviation between inversion and editing trajectories, as shown in Figure 2. The overall pipeline of
 136 EditAnyShape is shown in Figure 3.

137 3.1 MOTIVATION

138 Effective image editing requires a precise balance between introducing new content and preserv-
 139 ing the original structure. As illustrated in Figure 2 (left), traditional structure-preserving editing
 140 approaches often produce unstable denoising trajectories that deviate significantly from the stable
 141 reconstruction path, leading to severe structural degradation and undesired artifacts. Moreover, prior
 142 methods for localizing edits have notable drawbacks:

- 143 • **Binary Segmentation Masks:** Rely on external tools (Kirillov et al., 2023; Ronneberger
 144 et al., 2015), introducing overhead and a dependency on mask quality. Their rigid bound-
 145aries hinder large-scale shape changes and often produce artifacts.
- 146 • **Cross-Attention Masks:** Inferred from model’s cross attention during the diffusion pro-
 147 cess, these maps are often noisy and inconsistent, proving unreliable for localizing edits,
 148 especially during significant shape transformations.
- 149 • **Unconditional Feature Injection:** This strategy preserves structure by globally injecting
 150 source features, but its lack of selectivity suppresses intentional edits, creating a conflict
 151 between editability and consistency.

152 To address these limitations, we propose a new approach from a dynamical systems perspective.
 153 We posit that the semantic difference between the source and target concepts can be measured by
 154 the divergence between their respective denoising trajectories. Based on this, we achieved a precise
 155 and mask-free method (shown in Figure 3) to stabilize the editing trajectory and perform targeted,
 156 shape-aware modifications without relying on external masks or rigid heuristics.

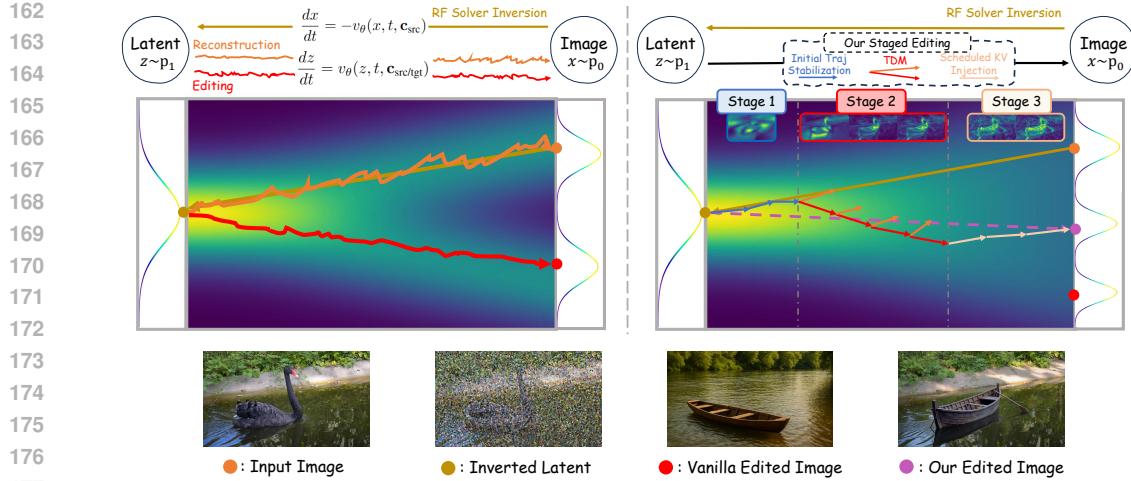


Figure 2: **Motivation for Trajectory Divergence Map (TDM) Guided Editing.** **Left:** Vanilla editing methods (red) often produce unstable trajectories compared to the stable reconstruction path (orange). **Right:** Our staged editing approach better resembles the ideal editing path. The TDM visualizes the dynamically localized editing region across different timesteps, with different border colors corresponding to different stages. [Update the figure](#).

3.2 EDITANYSHAPE

We perform shape-aware editing through a staged editing process that combines scheduled Key-Value (KV) injection with structural guidance, where the edit is localized by the Trajectory Divergence Map (TDM).

3.2.1 TRAJECTORY DIVERGENCE MAP

Our approach is grounded in the perspective of flow trajectories within the latent space, extending concepts from flow-matching frameworks to the inference setting. As illustrated in Figure 2 (left), a standard reconstruction follows a stable denoising trajectory guided by the source prompt \mathbf{c}_{src} . In an editing task, conditioning on a target prompt \mathbf{c}_{tgt} alters the velocity field, causing the denoising trajectory to deviate from this initial path. We posit that the magnitude of this deviation spatially localizes the semantic difference between the two prompts. Regions intended for modification will exhibit significant divergence, while background areas will follow nearly identical trajectories. To formalize this, let $\{\mathbf{x}_t\}_{t=0}^T$ be the latent sequence from the source image inversion, and let $\{\mathbf{z}_t\}_{t=0}^T$ be the corresponding sequence during the editing (denoising) process. We define the token-wise **Trajectory Divergence Map (TDM)** δ_t at timestep t as the L_2 norm of the difference between the velocity vectors predicted under the two prompts:

$$\delta_t^{(i)} = \left\| v_{\theta}(\mathbf{z}_t^{(i)}, t, \mathbf{c}_{\text{tgt}}) - v_{\theta}(\mathbf{x}_t^{(i)}, t, \mathbf{c}_{\text{src}}) \right\|_2, \quad (1)$$

where the velocity fields are evaluated at their respective trajectory latents, \mathbf{z}_t and \mathbf{x}_t . To enhance interpretability and prepare the map for temporal aggregation, we apply min-max normalization across all spatial tokens i at each timestep:

$$\tilde{\delta}_t^{(i)} = \frac{\delta_t^{(i)} - \min_j \delta_t^{(j)}}{\max_j \delta_t^{(j)} - \min_j \delta_t^{(j)}}. \quad (2)$$

As shown in Figure 2 (right), this produces a normalized TDM, $\{\tilde{\delta}_t^{(i)}\}$, which quantifies the localized editing strength on a scale of $[0, 1]$.

3.2.2 STAGED EDITING AND STRUCTURAL GUIDANCE

Directly applying TDM-guided injection across all timesteps is suboptimal due to the instability of the TDM in early, high-noise regimes (Figure 2 right). Early latents provide weak and noisy

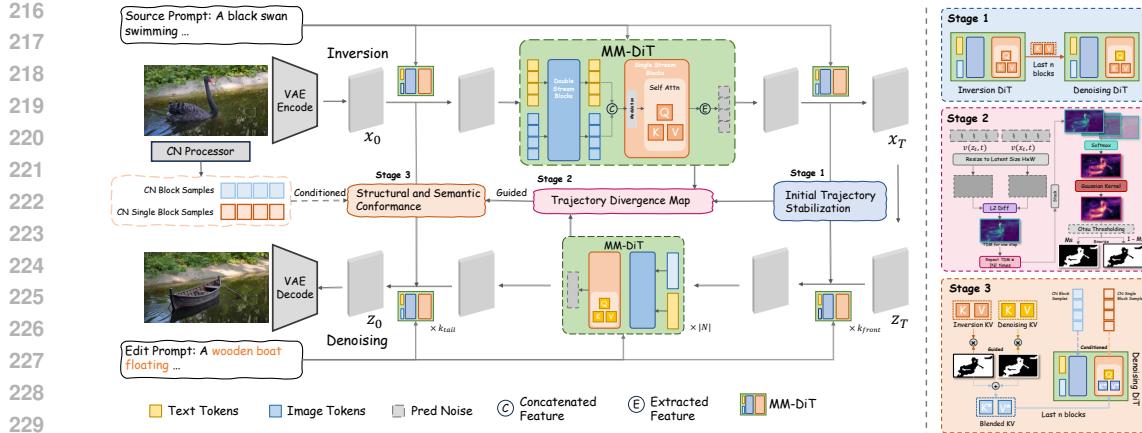


Figure 3: **Overview of our proposed pipeline.** Given a source image and the corresponding prompt, we first perform inversion to obtain the initial noisy latent code x_T . The editing process is then divided into three stages. **In Stage 1**, we stabilize the initial denoising trajectory by injecting key-value (KV) features from the inversion path into the denoising model during its initial steps. **In Stage 2**, we compute a Trajectory Divergence Map (TDM) by comparing the denoising trajectories generated from the source and edit prompts. This map is then processed to precisely identify the regions intended for editing. **In Stage 3**, guided by the TDM, blended KV features are injected into the final attention blocks of the denoising model to introduce the new semantics. Simultaneously, ControlNet conditions are supplied to ensure the edited regions conform to the original structure.

spatial signals, which can mislocalize edits if aggressive guidance is applied too soon. To address this, we introduce a scheduled injection strategy that partitions the N denoising steps into three distinct phases and adapts the guidance mechanism to the latent state: the first phase emphasizes stabilization, the second collects and aggregates TDM evidence while allowing exploration, and the third enforces structural and semantic conformance.

Stage 1: Initial Trajectory Stabilization. For an initial set of k_{front} timesteps, we perform unconditional KV injection from the source inversion path across all spatial tokens. This operation enforces a global reconstruction objective, equivalent to setting the edit mask $M_S = \mathbf{0}$, which stabilizes the trajectory and prevents semantic drift while the latent representation \mathbf{z}_t is still dominated by noise. Intuitively, the model first anchors to a faithful reconstruction manifold before any region-specific modification is attempted, reducing the risk of spurious changes to background layout or texture.

Stage 2: Editing and TDM Aggregation. Once a stable latent structure has emerged, we begin the editing phase over a predefined window of timesteps N . During this window, we perform editing by setting the edit mask $M_S = 1$ at every step, allowing the model to explore target-guided generation path. Simultaneously, we compute and store the normalized TDMs $\tilde{\delta}_t$ at each timestep within N , capturing the trajectory divergence guided by the source and target prompts. After this editing window concludes, we aggregate the stored TDMs $\{\tilde{\delta}_t\}$ across time to construct a temporally consistent and spatially coherent edit mask. Specifically, throughout the denoising process, a token that appears unchanged at an individual timestep may still experience evolution at subsequent steps. Therefore, to ensure that the aggregation faithfully captures such temporal dynamics, we employ a softmax-weighted temporal fusion for each token i :

$$\hat{\delta}^{(i)} = \sum_{t \in N} \alpha_t^{(i)} \cdot \tilde{\delta}_t^{(i)}, \quad \text{where} \quad \alpha_t^{(i)} = \frac{\exp(\tilde{\delta}_t^{(i)})}{\sum_{t' \in N} \exp(\tilde{\delta}_{t'}^{(i)})}. \quad (3)$$

To ensure spatial coherence and suppress noisy edges, the resulting map $\hat{\delta}$ is further refined via convolution with a Gaussian kernel \mathcal{G}_σ to obtain $\tilde{M}_S \in [0, 1]^{H \times W}$:

$$\tilde{M}_S = \mathcal{G}_\sigma * \hat{\delta}. \quad (4)$$

We observe that the distribution of values in \tilde{M}_S typically exhibits a skewed unimodal shape (as shown in Figure 4), characterized by a dominant background mode and a long-tailed foreground response. Such a distribution is well suited for Otsu’s method (Otsu et al., 1975), which selects the threshold τ that maximizes the inter-class variance of values. Formally, for a candidate threshold τ , let ω_0, μ_0 and ω_1, μ_1 denote the class probabilities and means of the background ($\tilde{M}_S \leq \tau$) and foreground ($\tilde{M}_S > \tau$), respectively. The between-class variance is defined as:

$$\sigma_b^2(\tau) = \omega_0(\tau) \omega_1(\tau) (\mu_0(\tau) - \mu_1(\tau))^2. \quad (5)$$

The optimal threshold is then given by:

$$\tau^* = \arg \max_{\tau} \sigma_b^2(\tau). \quad (6)$$

The final binary mask M_S is thus obtained by applying this threshold:

$$M_S = \mathbf{1}[\tilde{M}_S > \tau^*], \quad M_S \in \{0, 1\}^{H \times W}, \quad (7)$$

where $\mathbf{1}[\cdot]$ denotes the indicator function.

Stage3: Structural and Semantic Conformance. Our framework enforces structural conformance by jointly leveraging TDM-guided feature injection for background preservation and ControlNet residual conditioning for stabilizing structural patterns. The mask M_S obtained in Stage 2 modulates the fusion of Key-Value features, activating the target features ($K_{\text{tgt}}, V_{\text{tgt}}$) in edited regions and reverting to the source features ($K^{\text{inv}}, V^{\text{inv}}$) elsewhere. This feature-blending operation is formulated as:

$$\{K^*, V^*\} \leftarrow M_S \odot \{K^{\text{tgt}}, V^{\text{tgt}}\} + (1 - M_S) \odot \{K^{\text{inv}}, V^{\text{inv}}\}. \quad (8)$$

For structural guidance, ControlNet conditions the process on structural information \mathbf{c}_{cond} by injecting a residual stream into each block of the denoising model v_θ . For a latent representation \mathbf{z}_t at a given block, the output \mathbf{z}'_t is computed as:

$$\mathbf{z}'_t = \text{Block}(\mathbf{z}_t) + \beta \cdot \text{ControlNetBlock}(\mathbf{z}_t, \mathbf{c}_{\text{cond}}), \quad (9)$$

where β controls the guidance strength. Concurrently, our feature injection mechanism builds on RF-Edit’s background preservation by replacing the standard self-attention with a TDM-guided variant. The modified attention output \mathbf{F}'_{out} is computed using the blended key-value pairs from Eq. 8:

$$\mathbf{F}'_{\text{out}} = \text{Attention}(Q^{\text{tgt}}, K^*, V^*). \quad (10)$$

This synergy between ControlNet’s geometric enforcement and our TDM-guided semantic preservation enables precise, high-fidelity edits. The algorithmic implementation can be found in Algorithm 1 in Appendix.

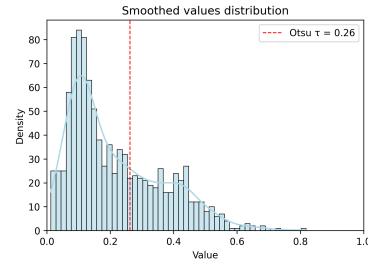


Figure 4: Distribution of values of \tilde{M}_S . The red dashed line indicates the Otsu threshold τ .

4 RESHAPEBENCH: A BENCHMARK FOR LARGE-SCALE SHAPE TRANSFORMATIONS

Overview. Existing benchmarks for image editing (Wang et al., 2023; Ju et al., 2023; Zhang et al., 2023) are not tailored to the demands of shape-aware editing, where the goal is to change object geometry while preserving the surrounding background. In particular, *PIE-Bench* (Ju et al., 2023) (700 images) uses concise prompts that often lack spatial or structural detail, and it aggregates heterogeneous tasks (object replacement, stylization, background modification) rather than isolating shape transformation as a first-class target. These properties make it difficult to diagnose whether a method truly performs structural change or relies on side effects such as texture shifts or background re-synthesis. We therefore introduce *ReShapeBench*, a benchmark that centers on mask-free, prompt-driven shape transformation with paired prompts and controlled background settings. This design isolates the factors relevant to structural change and reduces confounds from style or background alterations, enabling a targeted and reproducible evaluation protocol. It also serves as a targeted complement to existing evaluation suites, providing shape-focused test cases that fill the gap left by current general-purpose editing benchmarks.

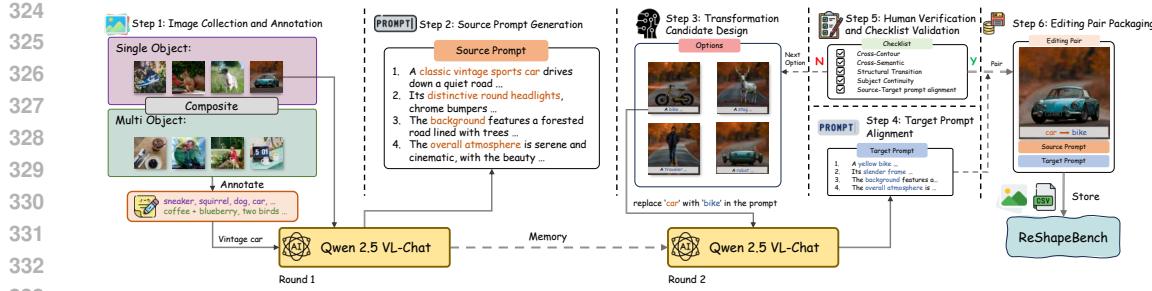


Figure 5: **Construction Process of ReShapeBench.** Note that images in Step 3 are generated after benchmark construction to serve as visual references. Checklist validation is performed on prompt.

Benchmark Construction. *ReShapeBench* contains 120 newly collected images split into three subsets: 70 single-object scenes for precise shape editing, 50 multi-object scenes for targeted mask-free edits, and a general evaluation set of 50 images that combines samples from both subsets with curated *PIE-Bench* cases to assess generalization. All images are standardized to 512×512 to normalize spatial scale and reduce variability across methods and backbones. Each new image is paired with two distinct shape transformations, yielding 240 editing cases across the single- and multi-object subsets, plus 50 cases in the general set, for a total of 290 shape-aware editing cases; this pairing increases task coverage and controls difficulty by varying the magnitude of structural change. Source-target prompts follow a structured template and differ only in the foreground object description, which stabilizes text-to-image alignment while holding background fixed. All prompt pairs are generated by Qwen-2.5-VL (Bai et al., 2025) and validated by human raters to ensure alignment and that the transformation satisfies the predefined shape criteria; ambiguous cases are double-checked to maintain consistency. Figure 5 and the Appendix B present the construction procedure and sample cases, including the selection checklist and prompt schema used during curation.

5 EXPERIMENT

5.1 EXPERIMENTAL SETUP

We use the open-source FLUX.1-[dev] model (Labs, 2024) as the base and run all experiments in PyTorch on an NVIDIA A100 (40 GB). We set the number of denoising steps to 14, guidance scale to 2.0, and k_{front} to 2. We evaluate both ControlNet-free and ControlNet-enabled variants of our method, since the core shape-editing mechanism operates without any external conditioning. If enabled, we apply multi-ControlNet conditioning with depth and Canny branches over the normalized denoising interval $[0.1, 0.3]$, with respective strengths 2.5 and 3.5. Unless otherwise stated, we keep the same inference scheduler and tokenizer as the official release and fix random seeds for reproducibility; additional implementation details and runtime/memory profiles are provided in Appendix C.

5.2 COMPARISON WITH BASELINES

5.2.1 QUALITATIVE COMPARISON

We compare EditAnyShape with *diffusion-based* and *flow-based* methods. Diffusion-based baselines include MasaCtrl (Cao et al., 2023), PnPInversion (Ju et al., 2023), and Dit4Edit (Feng et al., 2025), which modulate attention and conditions during the diffusion process. *Flow-based* baselines include RF-Edit (Wang et al., 2024), FlowEdit (Kulikov et al., 2025), KV-Edit (Zhu et al., 2025), FLUX.1 Fill (Labs, 2024), and FLUX.1 Kontext (Labs et al., 2025), which build on Rectified Flow. FLUX.1 Fill is designed for prompt-based masked image completion, while FLUX.1 Kontext leverages context-token concatenation for in-context editing. Figure 6 shows that EditAnyShape achieves stronger shape-aware editing and background preservation. Diffusion-based methods tend to degrade the background under structural edits and may fail on high-magnitude shape changes, while flow-based methods produce higher-quality images but still exhibit detail jitter, ghosting, or

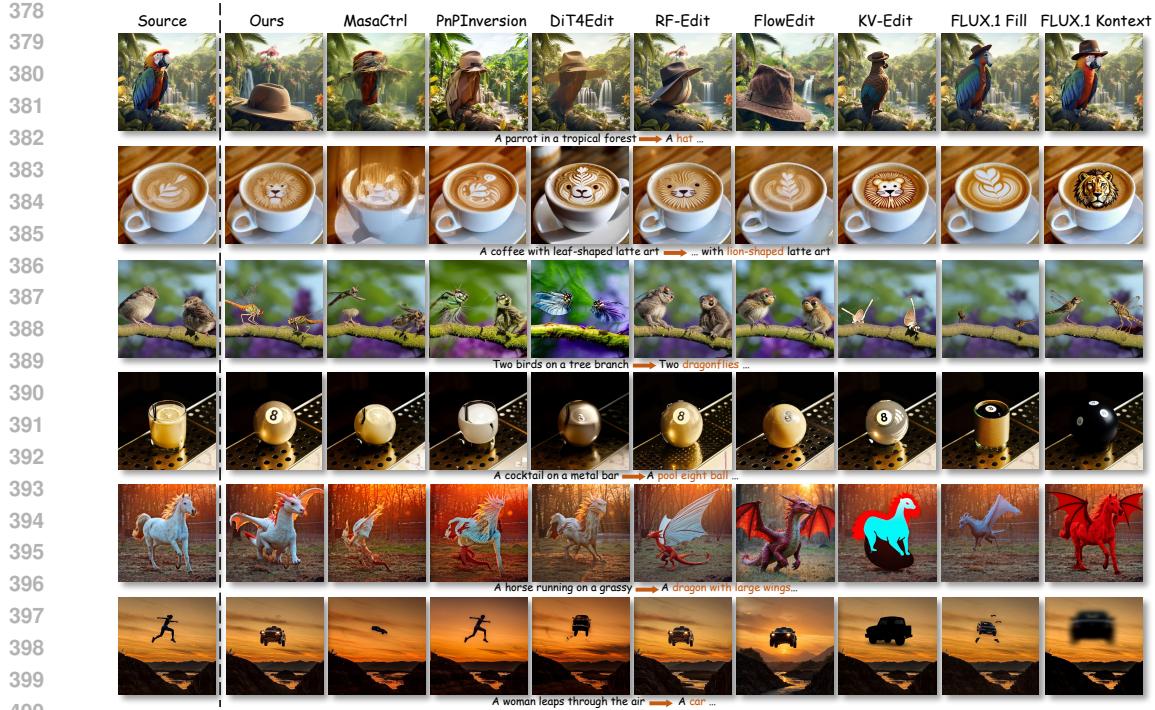


Figure 6: **Qualitative comparisons on various shape-aware editing cases.** EditAnyShape successfully performs large-scale shape transformations while preserving the background, demonstrating advantages in both editing ability and visual consistency over existing baselines. **We add the qualitative comparison with FLUX.1 Fill.**

Table 1: **Quantitative comparison with state-of-the-art methods on *ReShapeBench* and *PIE-Bench*.** **We add the quantitative comparison with FLUX.1 Fill, and our method without ControlNet.**

Datasets		<i>ReShapeBench</i>				<i>PIE-Bench</i>			
Metrics	Image Quality	Background Preservation		Text Align	Image Quality	Background Preservation		Text Align	
		AS \uparrow	PSNR \uparrow			AS \uparrow	PSNR \uparrow		
MasaCtrl(Cao et al., 2023)	5.83	23.54	125.36	20.84	5.61	21.58	130.71	19.53	
PnPInversion(Ju et al., 2023)	6.11	24.77	102.91	19.23	5.94	22.69	108.43	24.62	
Dit4Edit(Feng et al., 2025)	6.14	24.36	83.75	22.66	6.03	22.74	97.65	23.87	
RF>Edit(Wang et al., 2024)	6.52	33.28	17.53	30.41	6.49	31.97	15.34	29.67	
Flow>Edit(Kulikov et al., 2025)	6.42	32.46	18.92	28.94	6.37	32.68	16.42	28.93	
KV>Edit(Zhu et al., 2025)	6.51	34.73	16.42	26.97	6.47	33.45	13.72	28.14	
FLUX.1Fill(Labs, 2024)	6.32	31.57	19.04	28.75	6.33	32.76	17.43	26.59	
FLUX.1Kontext(Labs et al., 2025)	6.53	32.91	18.35	28.53	6.47	34.91	14.62	28.79	
Ours (w/o ControlNet)	6.52	34.85	9.04	32.97	6.49	35.62	9.74	32.47	
Ours (Full Model)	6.57	35.79	8.23	33.71	6.55	36.02	8.34	33.51	

incomplete transformations in difficult cases. EditAnyShape performs large-scale shape transformations while preserving non-target regions.

5.2.2 QUANTITATIVE COMPARISON

We conduct quantitative evaluations on both *ReShapeBench* and *PIE-Bench* against diffusion- and flow-based baselines to assess both shape-aware editing and general editing performances. To ensure fairness, we use identical source and target prompts and the same number of denoising steps across methods. Because we follow RF-Solver with a second-order scheme, we double the number of steps for methods without a second-order update to match the number of function evaluations (NFE). We disable the ControlNet modules to isolate the effect of TDM-guided editing.

432 As shown in Table 1, we evaluate background consistency with PSNR (Huynh-Thu & Ghanbari,
 433 2008) and LPIPS (Zhang et al., 2018), image quality with the LAION Aesthetic Score (Schuhmann
 434 et al., 2022), and text alignment with CLIP similarity (Radford et al., 2021). Appendix B.3 demon-
 435 strates the implementation details, including the preprocessing and metric computation settings used
 436 for all methods. Our method outperforms all baselines across metrics. Additionally, without the
 437 ControlNet module does not lead to a significant degradation in editing performance, indicating the
 438 improvements of our model are independent of this module. The region-controlled editing strategy
 439 improves fine-grained shape-aware editing while the mask M_S preserves background content.

440

441 5.3 ABLATION STUDY

442

443 Initial trajectory stabilization and the timing and strength of ControlNet conditioning have the largest
 444 impact on editing performance. We therefore ablate these two components: the former regulates
 445 early trajectory stability, while the latter controls structural guidance during mid–late steps.

446

447 5.3.1 EFFECTIVENESS OF INITIAL TRAJECTORY STABILIZATION

448

449 To assess the role of initial trajectory stabilization, we
 450 vary the number of stabilization steps k_{front} from 0 to
 451 4. As shown in Figure 7(i), small k_{front} leads to drift
 452 and structural deviation, while large k_{front} restricts the in-
 453 tended shape change. Table 2 shows that larger k_{front} im-
 454 proves background preservation but reduces CLIP sim-
 455 ilarity, indicating a trade-off between stability and ed-
 456 itability. $k_{\text{front}} = 2$ provides the best balance, yielding
 457 stable trajectories while maintaining sufficient freedom
 458 for large shape transitions.

459

460 5.3.2 EFFECTIVENESS OF CONTROLNET 461 CONDITIONING Timestep and Strength

462

463 To explore the effect of ControlNet conditioning timestep,
 464 we vary the injection interval within the normalized de-
 465 noising range $[0, 1]$. Figure 7(ii.a) shows that earlier in-
 466 jection yields better results, as latent features are less
 467 noisy and more receptive to structural guidance. We also
 468 vary the depth and Canny strengths. As shown in Figure 7(ii.b), moderate values (e.g., $(2.5, 3.5)$)
 469 best balance structure preservation and editability, while overly weak or strong signals under- or
 470 over-constrain the edit. These results suggest that early, moderate guidance best stabilizes geometry
 471 without suppressing desired semantic changes in the edited regions.

472

473

474 Table 2: **Ablation study on different k_{front} .**

475

476

k_{front}	Image Quality		Background Preservation		Text Align	
	Aesthetic Score \uparrow	PSNR \uparrow	LPIPS $\times 10^3 \downarrow$	CLIP Sim \uparrow		
0	6.51	32.79	10.04	31.05		
1	6.55	34.38	9.88	32.56		
2	6.57	35.79	8.23	33.71		
3	6.52	31.25	10.52	29.41		
4	6.48	30.41	12.37	27.66		

477

478

479

480

481

482

483

484

485



Figure 7: (i) Ablation study on k_{front} .
 (ii.a) ControlNet conditioning applied within five subranges of the normalized denoising interval $[0, 1]$. (ii.b) ControlNet conditioning strength of the depth and canny branches, denoted as (depth, canny).

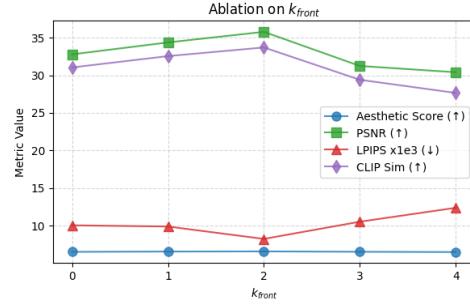


Figure 8: **Visualization of ablation on k_{front} .**
 While the aesthetic score remains relatively stable, PSNR, LPIPS, and CLIP Similarity reveal a clear trade-off between editing strength and background preservation.

486

6 CONCLUSION

488 We introduce EditAnyShape, a framework that enables large-scale object shape transformation by
 489 using a novel trajectory-based region control mechanism. Our method achieves precise, mask-free
 490 edits while preserving background integrity by dynamically localizing modifications through a Tra-
 491 jectory Divergence Map with scheduled injection. To properly evaluate this task, we developed
 492 *ReShapeBench*, a new benchmark tailored for complex shape-aware editing. To the best of our
 493 knowledge, EditAnyShape is the first work to systematically address prompt-driven shape editing.
 494 Extensive qualitative and quantitative experiments validate its state-of-the-art performance on the
 495 proposed benchmark. Our work thus opens promising new avenues for controllable generation.

496

497 ETHICS STATEMENT

498 Our method is designed for image editing tasks involving large shape transformations. To mitigate
 499 potential misuse such as malicious editing, the final model output incorporates an NSFW filtering
 500 component. The benchmark dataset used in this work is entirely collected from publicly available
 501 sources (<https://www.pexels.com/>), which explicitly permit free usage and modification of images
 502 and videos. A processed version of the dataset will be released solely for research purposes af-
 503 ter further optimization. Beyond these considerations, our work does not involve any sensitive or
 504 personally identifiable data.

505

506 REPRODUCIBILITY STATEMENT

- 509 1. Implementation details of the EditAnyShape are described in Section 3; the corresponding
 510 pseudocode is provided in Algorithm 1.
- 511 2. Dataset and benchmark constructions are described in Section 3.2.2 and Appendix B. A
 512 complete version of *ReShapeBench* will be released after further verification. A sample
 513 subset of images and prompts is provided in Appendix B.2.
- 514 3. Source code is provided in the supplementary files. Details of hyperparameters is provided
 515 in Appendix C.1.

516

517 REFERENCES

518 Omri Avrahami, Ohad Fried, and Dani Lischinski. Blended latent diffusion. *ACM transactions on*
 519 *graphics (TOG)*, 42(4):1–11, 2023.

520 Omri Avrahami, Or Patashnik, Ohad Fried, Egor Nemchinov, Kfir Aberman, Dani Lischinski, and
 521 Daniel Cohen-Or. Stable flow: Vital layers for training-free image editing. In *Proceedings of the*
 522 *Computer Vision and Pattern Recognition Conference*, pp. 7877–7888, 2025.

523 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
 524 Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan,
 525 Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng,
 526 Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025.
 527 URL <https://arxiv.org/abs/2502.13923>.

528 Connally Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A ran-
 529 domized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24,
 530 2009.

531 Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Mas-
 532 actrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In
 533 *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 22560–22570,
 534 2023.

535 Siran Chen, Qinglin Xu, Yue Ma, Yu Qiao, and Yali Wang. Attentive snippet prompting for video
 536 retrieval. *IEEE Transactions on Multimedia*, 26:4348–4359, 2023.

540 Siran Chen, Yuxiao Luo, Yue Ma, Yu Qiao, and Yali Wang. H-mba: Hierarchical mamba adap-
 541 tation for multi-modal video understanding in autonomous driving. In *Proceedings of the AAAI*
 542 *Conference on Artificial Intelligence*, volume 39, pp. 2212–2220, 2025.

543

544 Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. Diffedit: Diffusion-
 545 based semantic image editing with mask guidance. *arXiv preprint arXiv:2210.11427*, 2022.

546

547 Kunyu Feng, Yue Ma, Bingyuan Wang, Chenyang Qi, Haozhe Chen, Qifeng Chen, and Zeyu Wang.
 548 Dit4edit: Diffusion transformer for image editing. In *Proceedings of the AAAI Conference on*
 549 *Artificial Intelligence*, volume 39, pp. 2969–2977, 2025.

550

551 Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or.
 552 Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*,
 553 2022.

554

555 Yi Huang, Jiancheng Huang, Yifan Liu, Mingfu Yan, Jiaxi Lv, Jianzhuang Liu, Wei Xiong,
 556 He Zhang, Liangliang Cao, and Shifeng Chen. Diffusion model-based image editing: A survey.
 557 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

558

559 Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of psnr in image/video quality as-
 560 sessment. *Electronics letters*, 44(13):800–801, 2008.

561

562 Xuan Ju, Ailing Zeng, Yuxuan Bian, Shaoteng Liu, and Qiang Xu. Direct inversion: Boosting
 563 diffusion-based editing with 3 lines of code. *arXiv preprint arXiv:2310.01506*, 2023.

564

565 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete
 566 Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceed-
 567 ings of the IEEE/CVF international conference on computer vision*, pp. 4015–4026, 2023.

568

569 Vladimir Kulikov, Matan Kleiner, Inbar Huberman-Spiegelglas, and Tomer Michaeli. Flowedit:
 570 Inversion-free text-based editing using pre-trained flow models. In *Proceedings of the IEEE/CVF*
 571 *International Conference on Computer Vision*, pp. 19721–19730, 2025.

572

573 Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024.

574

575 Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril
 576 Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey,
 577 Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini,
 578 Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and
 579 editing in latent space, 2025. URL <https://arxiv.org/abs/2506.15742>.

580

581 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 582 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

583

584 Hongyu Liu, Xuan Wang, Ziyu Wan, Yue Ma, Jingye Chen, Yanbo Fan, Yujun Shen, Yibing Song,
 585 and Qifeng Chen. Avatarartist: Open-domain 4d avatarization. In *Proceedings of the Computer*
 586 *Vision and Pattern Recognition Conference*, pp. 10758–10769, 2025.

587

588 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 589 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

590

591 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast
 592 ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural*
 593 *information processing systems*, 35:5775–5787, 2022.

594

595 Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool.
 596 Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the*
 597 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 11461–11471, 2022.

598

599 Xiang Lv, Mingwen Shao, Yecong Wan, Yue Ma, Yuanshuo Cheng, and Lingzhuang Meng. Bm-
 600 edit: Background retention and motion consistency for zero-shot video editing. *Knowledge-Based*
 601 *Systems*, pp. 113784, 2025.

594 Xuran Ma, Yexin Liu, Yaofu Liu, Xianfeng Wu, Mingzhe Zheng, Zihao Wang, Ser-Nam Lim, and
 595 Harry Yang. Model reveals what to cache: Profiling-based feature reuse for video diffusion
 596 models. *arXiv preprint arXiv:2504.03140*, 2025.

597 Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for
 598 editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF conference*
 599 *on computer vision and pattern recognition*, pp. 6038–6047, 2023.

600 Nobuyuki Otsu et al. A threshold selection method from gray-level histograms. *Automatica*, 11
 601 (285–296):23–27, 1975.

602 Yatian Pang, Bin Zhu, Bin Lin, Mingzhe Zheng, Francis EH Tay, Ser-Nam Lim, Harry Yang, and
 603 Li Yuan. Dreamdance: Animating human images by enriching 3d geometry cues from 2d poses.
 604 *arXiv preprint arXiv:2412.00397*, 2024.

605 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*
 606 *the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

607 Reiner Pope, Sholto Douglas, Aakanksha Chowdhery, Jacob Devlin, James Bradbury, Jonathan
 608 Heek, Kefan Xiao, Shivani Agrawal, and Jeff Dean. Efficiently scaling transformer inference.
 609 *Proceedings of machine learning and systems*, 5:606–624, 2023.

610 Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng
 611 Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. In *Proceedings of the*
 612 *IEEE/CVF International Conference on Computer Vision*, pp. 15932–15942, 2023.

613 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 614 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 615 models from natural language supervision. In *International conference on machine learning*, pp.
 616 8748–8763. PMLR, 2021.

617 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
 618 image segmentation. In *International Conference on Medical image computing and computer-assisted*
 619 *intervention*, pp. 234–241. Springer, 2015.

620 Litu Rout, Yujia Chen, Nataniel Ruiz, Constantine Caramanis, Sanjay Shakkottai, and Wen-Sheng
 621 Chu. Semantic image inversion and editing using rectified stochastic differential equations. *arXiv*
 622 *preprint arXiv:2410.10792*, 2024.

623 Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi
 624 Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An
 625 open large-scale dataset for training next generation image-text models. *Advances in neural information*
 626 *processing systems*, 35:25278–25294, 2022.

627 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*
 628 *preprint arXiv:2010.02502*, 2020.

629 Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for
 630 text-driven image-to-image translation. In *Proceedings of the IEEE/CVF conference on computer*
 631 *vision and pattern recognition*, pp. 1921–1930, 2023.

632 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 633 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information*
 634 *processing systems*, 30, 2017.

635 Bram Wallace, Akash Gokul, and Nikhil Naik. Edict: Exact diffusion inversion via coupled trans-
 636 formations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recog-*
 637 *nition*, pp. 22532–22541, 2023.

638 Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu,
 639 Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai
 640 Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi
 641 Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang,
 642 Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang,
 643 Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang,
 644 Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang,
 645 Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang,
 646 Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang,
 647 Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang,

648 Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng
 649 Zhou, Wente Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan
 650 Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You
 651 Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen
 652 Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models,
 653 2025. URL <https://arxiv.org/abs/2503.20314>.

654 Zhen Wan, Yue Ma, Chenyang Qi, Zhiheng Liu, and Tao Gui. Unipaint: Unified space-time video
 655 inpainting via mixture-of-experts. *arXiv preprint arXiv:2412.06340*, 2024.

656

657 Jiangshan Wang, Junfu Pu, Zhongang Qi, Jiayi Guo, Yue Ma, Nisha Huang, Yuxin Chen, Xiu Li,
 658 and Ying Shan. Taming rectified flow for inversion and editing. *arXiv preprint arXiv:2411.04746*,
 659 2024.

660 Su Wang, Chitwan Saharia, Ceslee Montgomery, Jordi Pont-Tuset, Shai Noy, Stefano Pellegrini,
 661 Yasumasa Onoe, Sarah Laszlo, David J Fleet, Radu Soricut, et al. Imagen editor and editbench:
 662 Advancing and evaluating text-guided image inpainting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18359–18369, 2023.

663

664 Zhen Xiong, Yuqi Li, Chuanguang Yang, Tiao Tan, Zhihong Zhu, Siyuan Li, and Yue Ma. Enhancing
 665 image generation fidelity via progressive prompts. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2025.

666

667 Zexuan Yan, Yue Ma, Chang Zou, Wenteng Chen, Qifeng Chen, and Linfeng Zhang. Edit:
 668 Rethinking the spatial and temporal redundancy for efficient image editing. *arXiv preprint arXiv:2503.10270*, 2025.

669

670

671 Kai Zhang, Lingbo Mo, Wenhui Chen, Huan Sun, and Yu Su. Magicbrush: A manually annotated
 672 dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*,
 673 36:31428–31449, 2023.

674

675 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 676 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

677

678 Zhongping Zhang, Jian Zheng, Zhiyuan Fang, and Bryan A Plummer. Text-to-image editing by
 679 image information removal. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 5232–5241, 2024.

680

681

682 Mingzhe Zheng, Yongqi Xu, Haojian Huang, Xuran Ma, Yexin Liu, Wenjie Shu, Yatian Pang, Fei-
 683 long Tang, Qifeng Chen, Harry Yang, et al. Videogen-of-thought: A collaborative framework for
 684 multi-shot video generation. *arXiv e-prints*, pp. arXiv–2412, 2024.

685

686 Tianrui Zhu, Shiyi Zhang, Jiawei Shao, and Yansong Tang. Kv-edit: Training-free image editing for
 687 precise background preservation. *arXiv preprint arXiv:2502.17363*, 2025.

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

A PRELIMINARIES

703

A.1 RECTIFIED FLOW (RF)

706 Let p_0 and p_1 denote the source and target distributions, respectively. Flow Matching (Lipman et al.,
 707 2022) models the transport between them by learning a time-dependent velocity field $v(t, x)$ that
 708 defines a continuous transformation $\psi_t(x)$ via the ordinary differential equation:

$$709 \quad \frac{d\psi_t(x)}{dt} = v(t, \psi_t(x)), \quad \psi_0(x) \sim p_0, \quad \psi_1(x) \sim p_1. \quad (11)$$

712 Rectified Flow (RF) (Liu et al., 2022) simplifies this by assuming a linear trajectory between $X_0 \sim$
 713 p_0 and $X_1 \sim p_1$:

$$714 \quad X_t = (1 - t)X_0 + tX_1, \quad t \in [0, 1], \quad (12)$$

715 with the associated velocity field becomes:

$$717 \quad v(X_t, t) = X_1 - X_0. \quad (13)$$

719 The model is trained by minimizing the conditional flow matching loss:

$$720 \quad \mathcal{L}_{\text{CFM}} = \mathbb{E}_{X_0, X_1, t} [\|v(X_t, t) - (X_1 - X_0)\|^2]. \quad (14)$$

722 During inference, the learned velocity field is used to generate new samples by solving the reverse-
 723 time ODE:

$$725 \quad \frac{dX_t}{dt} = -v(X_t, t), \quad (15)$$

726 starting from a sample $X_1 \sim \mathcal{N}(0, I)$. Since a closed-form solution is not available in general, we
 727 perform numerical integration over a discretized set of timesteps $\{t_i\}_{i=0}^N$. A standard choice uses
 728 first-order solvers such as Euler or Heun’s method to approximate the trajectory:

$$730 \quad X_{t_{i-1}} = X_{t_i} - h \cdot v(X_{t_i}, t_i), \quad (16)$$

731 where $h = t_i - t_{i-1}$ is the integration step size.

733 However, first-order solvers can suffer from numerical instability and truncation error, especially in
 734 high-dimensional generation tasks. Several recent works (Lu et al., 2022; Rout et al., 2024; Lv et al.,
 735 2025; Chen et al., 2025; Wang et al., 2024) explore higher-order integration strategies or adaptive
 736 solvers to improve generation fidelity. Specifically, RF-Solver (Wang et al., 2024) introduces a
 737 second-order update derived from a Taylor expansion of the velocity field:

$$738 \quad X_{t_{i-1}} = X_{t_i} - h \cdot v(X_{t_i}, t_i) + \frac{1}{2}h^2 \cdot \partial_t v(X_{t_i}, t_i), \quad (17)$$

740 where $\partial_t v(X_{t_i}, t_i)$ is the time derivative of the learned velocity field. This correction term reduces
 741 local integration error and leads to more accurate inversion and sampling trajectories, which is par-
 742 ticularly important for downstream editing tasks that require high structural fidelity.

744

A.2 KV INJECTION

746 Key-Value (KV) injection is adapted from the KV caching mechanism originally used in Trans-
 747 formers (Vaswani et al., 2017) to accelerate autoregressive inference (Pope et al., 2023). In large
 748 language models, cached key and value tensors allow reuse of past attention computations, enabling
 749 efficient decoding without recomputing earlier tokens.

750 When extended from language to vision models, KV reuse often generalizes beyond strict token
 751 caching. In U-Net based models (Cao et al., 2023; Qi et al., 2023), a common practice is to reuse
 752 intermediate attention maps or inject features derived from the inverted source image into self-
 753 attention layers. This feature-level injection plays a role similar to KV caching in LLMs by enforc-
 754 ing spatial consistency and anchoring the generative process to the source structure. In DiT-based
 755 architectures (Peebles & Xie, 2023), this idea extends to reusing value (V) matrices or full KV pairs,
 providing finer-grained control over how structural information is preserved during denoising.

756 To reduce memory cost and avoid limiting foreground flexibility, recent work such as Stable-
757 Flow (Avrahami et al., 2025) explores the vital layers within DiT crucial for image formation.
758 Therefore, by only reusing KV pairs in a subset of layers, it can balance structural fidelity and
759 editability while effectively reducing memory usage.

760 In our work, we demonstrate that KV injection provides a modular and interpretable mechanism
761 for controllable image editing and particularly effective in shape-aware tasks where edits must stay
762 localized without affecting the broader scene.

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

B ADDITIONAL DETAILS ON RESHAPEBENCH CONSTRUCTION

B.1 SHAPE TRANSFORMATION

While recent image editing models exhibit strong general-purpose editing capabilities, the concept of shape transformation remains ambiguous in the literature. In practice, object modifications usually include detail adjustments, color changes, or limited geometric variations, often relying on masks or ControlNet images; however, such operations cannot be explicitly framed as shape transformations. When constructing *ReShapeBench*, we need a clear definition and categorization of shape transformation to guide data curation and enable meaningful evaluation.

From a geometric perspective, shape transformation is a structural change beyond local affine operations such as scaling, rotation, or minor warping. It reconfigures the object's global contour and part topology, and may involve a shift in semantic class. At the same time, the transformed object must remain spatially coherent in the scene, occupying a similar anchor position and continuing to serve as the subject in context. Guided by these principles, we propose four criteria—cross-contour, cross-semantic, structural transition, and subject continuity—that together capture the essential properties of shape transformation.

- **Cross-contour:** The object's boundary undergoes a substantial change, exceeding local warping or affine resizing. This captures large-scale alterations to the external shape.
- **Cross-semantic:** The transformation shifts the object into a different semantic class, indicating a categorical rather than attributive change, while preserving overall scene coherence.
- **Structural transition:** The internal part topology is reconfigured, requiring modifications across multiple components instead of only simple attributes such as color or texture.
- **Subject continuity:** The transformed object retains its spatial anchor and role in the scene, remaining contextually consistent despite the change in shape and semantic.

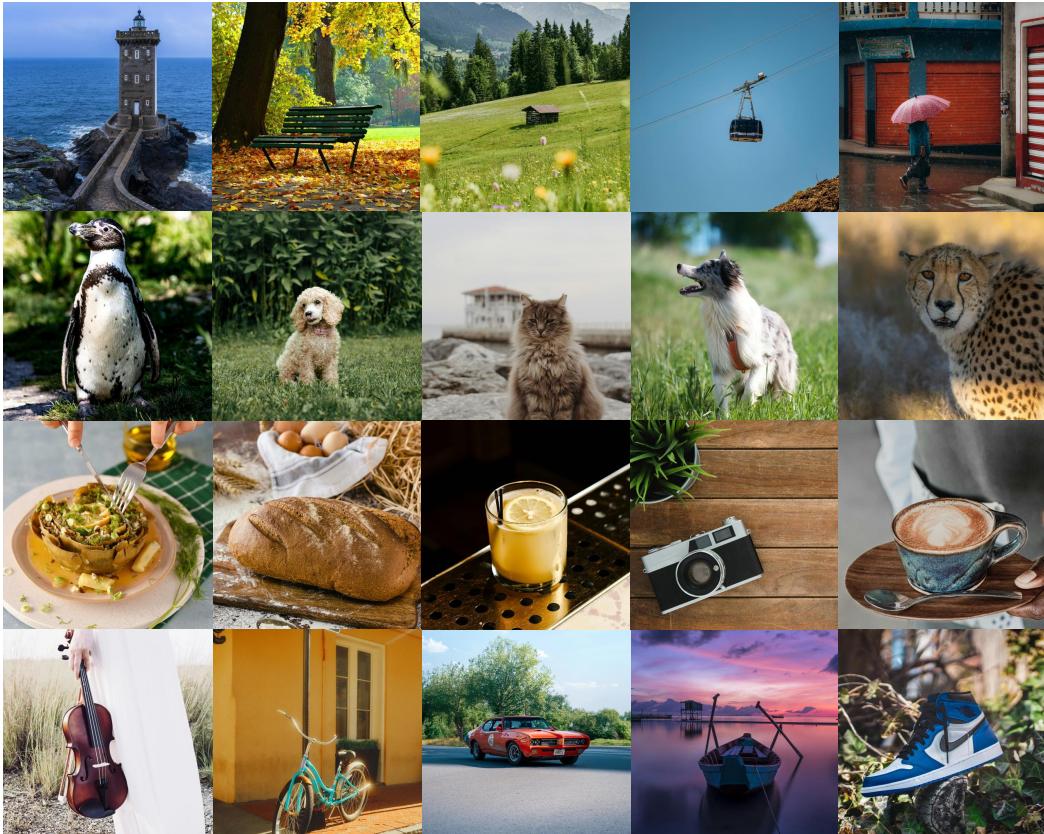
Almost every editing case in the paper can be classified as shape transformation. Figure 9 provides an additional illustrative example. Note that we exclude posture or viewpoint changes (standing → sitting), as these involve articulation or perspective variation rather than structural transformation. The most challenging part in shape transformation is that it requires the model to localize and reinterpret object shape while maintaining consistency in background and composition. Unlike existing editing benchmarks that cover diverse editing tasks, our benchmark emphasizes large-scale shape transformation under prompt guidance, without relying on masks or external conditioning.



Figure 9: **Visualization of Shape Transformation.** The object's contour, semantic, structure changed while ensuring its subject continuity.

864 B.2 IMAGE AND PROMPT EXAMPLES
865

866 As described in Section 3.2.2, the benchmark is constructed from collected images divided into
867 single-object and multi-object categories. Specifically, single-object cases broadly cover four cat-
868 egories—nature, animals, indoor, and outdoor scenes; multi-object cases can also be categorized
869 into indoor and outdoor scenes. Figures 10 and 11 provide representative samples. For each image,
870 source and target prompts follow the four-sentence template, which is illustrated in Table 5, 6, and
871 7.

890 Figure 10: Single-Object Cases in *ReShapeBench*
891916 Figure 11: Multi-Object Cases in *ReShapeBench*
917

918 B.3 EVALUATION METRICS
919920 We use four metrics grouped under three aspects: image quality, background preservation, and text-
921 image alignment.922 Aesthetic Score (AS) measures the **perceptual quality** of the generated image by indicating how
923 well the visual content conforms to natural image statistics. We compute AS with the LAION aes-
924 thetic predictor (Schuhmann et al., 2022), which applies a linear estimator on top of CLIP embed-
925 dings¹. AS helps detect unnatural boundaries or blending artifacts that may occur during large-scale
926 shape transitions and indicates how well the new shape integrates into the scene.927 For **background preservation**, we adopt two widely used metrics that capture different aspects
928 of visual similarity. Peak Signal-to-Noise Ratio (PSNR) measures low-level pixel fidelity, while
929 Learned Perceptual Image Patch Similarity (LPIPS) evaluates perceptual similarity based on deep
930 feature representations. **Since we do not assume access to ground-truth masks and the edited shape
931 can vary across models, we localize the edited region by applying a fixed-size box centered on the
932 subject to occlude the foreground and compute similarity over the remaining background area.** This
933 heuristic enables fair evaluation of how well the unedited content is preserved.934 Finally, to assess **text-image alignment**, we compute CLIP similarity between the generated image
935 and the target prompt as an embedding-based measure of semantic consistency.
936937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

1^{https://github.com/LAION-AI/aesthetic-predictor}

972 C ADDITIONAL IMPLEMENTATION DETAILS
973974 C.1 PSEUDOCODE FOR EDITANYSHAPE
975976 **Algorithm 1** Region-Controlled Editing
977

978 **Input:** Inference steps T , Predicted velocities $\{v_{\text{src}}(\mathbf{x}_t^{(i)}), v_{\text{tgt}}(\mathbf{z}_t^{(i)})\}_{i=1}^T$, Source inversion features
979 $\{K_t^{\text{inv}}, V_t^{\text{inv}}\}_{t=0}^T$, target prompt \mathbf{c}_{tgt} , schedule phase durations $\{k_{\text{front}}, k_{\text{tail}}\}$

```

980 1:  $N \leftarrow \{\}$  ▷ Editing window set
981 2: for  $t = T$  down to 1 do
982 3:   if  $t > T - k_{\text{front}}$  then
983 4:      $M_S \leftarrow \mathbf{0}$ 
984 5:   else if  $T - k_{\text{front}} \geq t > k_{\text{tail}}$  then
985 6:      $\delta_t^{(i)} \leftarrow \left\| v_{\text{tgt}}(\mathbf{z}_t^{(i)}, t) - v_{\text{src}}(\mathbf{x}_t^{(i)}, t, \mathbf{c}_{\text{src}}) \right\|_2$ 
986 7:      $\tilde{\delta}_t^{(i)} \leftarrow \frac{\delta_t^{(i)} - \min_j \delta_t^{(j)}}{\max_j \delta_t^{(j)} - \min_j \delta_t^{(j)}}$  ▷ TDM Computation
987 8:      $M_S \leftarrow \mathbf{1}$ 
988 9:      $N \leftarrow N \cup \{t\}$ 
989 10:  else
990 11:     $\hat{\delta}^{(i)} \leftarrow \sum_{t' \in T} \frac{\exp(\tilde{\delta}_{t'}^{(i)})}{\sum_{t'' \in T} \exp(\tilde{\delta}_{t''}^{(i)})} \cdot \tilde{\delta}_{t'}^{(i)}$ 
991 12:     $\tilde{M}_S \leftarrow \mathcal{G}_\sigma * \hat{\delta}$  ▷ TDM Aggregation
992 13:     $\tau \leftarrow \arg \max_{\tau'} \mathbb{P}[\tilde{M}_S \leq \tau'] \mathbb{P}[\tilde{M}_S > \tau'] (\mathbb{E}[\tilde{M}_S | \tilde{M}_S \leq \tau'] - \mathbb{E}[\tilde{M}_S | \tilde{M}_S > \tau'])^2$ 
993 14:     $M_S \leftarrow \mathbf{1}[\tilde{M}_S > \tau]$ 
994 15:  end if
995 16:   $K^* \leftarrow M_S \odot K_t^{\text{tgt}} + (1 - M_S) \odot K_t^{\text{inv}}$ 
996 17:   $V^* \leftarrow M_S \odot V_t^{\text{tgt}} + (1 - M_S) \odot V_t^{\text{inv}}$ 
997 18: end for
998 19: return  $K^*, V^*$ 

```

1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

1026

1027

1028

C.2 HYPERPARAMETERS

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

Table 3: **Hyperparameters**

Hyperparameter	Value
<i>General</i>	
Inference Step	15
Guidance	2
<i>Model Specific</i>	
k_{front}	2
k_{tail}	3
Softmax scale (temperature)	5
Gaussian smoothing σ	0.7
Injecting DiT block (start idx)	19
ControlNet Timing	[0.1, 0.3]
ControlNet Strength	(2.5, 3.5)

Table 4: **Hyperparameters used for the results displayed in the paper.**

Task	k_{front}	k_{tail}	CN type	CN timing	CN strength
parrot → hat, coconut, guitar, LOVE, rabbit	2	3	None	NA	NA
boy → cat, Teddy, backpack, soccer, lizard	1	3	Depth & Canny	[0.1, 0.3]	[2.5, 3.5]
bird → dragonflies, caps, robots, pokers	2	3	None	NA	NA
swan → boat, flamingo, crocodile, turtle	3	3	Depth	[0.1, 0.3]	0.6
leaf latte → lion, horse → dragon, cocktail → ball	2	3	None	NA	NA

C.3 RUNNING TIME AND MEMORY USAGE

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

Running Time Analysis. The method requires an additional diffusion pass each step to compute the second-order prediction, resulting in a total of 28 NFEs. As introduced in Section 5.1, we conduct our experiment on an NVIDIA A100 (40 GB), and the average running time for one image (averaged across multiple trials) is approximately 65.3 seconds. The computational cost is comparable to existing methods such as FlowEdit and RF-Solver-Edit. Since our work focuses on controllability and structural fidelity rather than acceleration, no additional optimization loops are introduced. Speed optimization is orthogonal to our contributions.

Memory Usage. During inference, the method stores a set of KV features and TDM maps on CPU memory, which amounts to approximately 12 GB in total. This cost is incurred only once

1080 during inversion and does not grow with the number of denoising steps. The GPU memory us-
1081 age remains stable at around 25 GB throughout editing, comparable to existing flow-based editing
1082 pipelines. Since the method does not introduce any additional training procedures, this overhead is
1083 not a limiting factor in practice.

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

1126

1127

1128

1129

1130

1131

1132

1133

1134

1135

1136

C.4 POST-HOC ANALYSIS OF STRUCTURAL LOCALIZATION SIGNALS

1137

1138 To better understand whether the proposed Trajectory Divergence Map (TDM) provides a mean-
 1139 ingful structural signal, we conduct a post-hoc analysis from two complementary perspectives: (1)
 1140 mask-level comparison and (2) comparison against cross-attention maps.

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

Mask-level comparison. Our first goal is to verify that TDM indeed identifies the correct structural region to be edited. Using the crocodile case from Figure 9, we construct a pseudo ground-truth region by taking the union of the source mask and the target mask (generated by SAM (Kirillov et al., 2023)), and then downsampling it to the latent resolution. This union mask reflects the intuitive region of change: areas covered by either the original foreground or the edited foreground are exactly those that should be modified, while the background should remain unchanged. We compare this downsampled pseudo mask with our mask M_S . As shown in the top row of Fig 12, the two maps exhibit similarity in editable areas, providing an intuitive signal that TDM captures the correct semantic region for editing without relying on external supervision. This supports our claim that TDM provides an effective and interpretable estimate of “where” the model intends to apply shape transformation.

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

TDM vs. cross-attention. We further compare TDM with cross-attention maps in FLUX DiT. We visualize representative timesteps from the most responsive attention block. Because cross-attention activation is tightly tied to prompt tokens, its localization quality depends heavily on which word has the strongest response. In practice, identifying a single “correct” token is not feasible—responses vary significantly across tokens, heads, and layers—so we use a softmax-normalized variant over all text tokens. Even under this stabilized setting, cross-attention remains spatially noisy and useless with the true editing region. These attention maps are extracted when using TDM-guided region control. Maps without guided control are substantially noisier and therefore omitted for clarity. In contrast, TDM offers a much cleaner and more direct indication of structural change, and it is obtained in a far simpler manner by relying solely on trajectory differences.

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

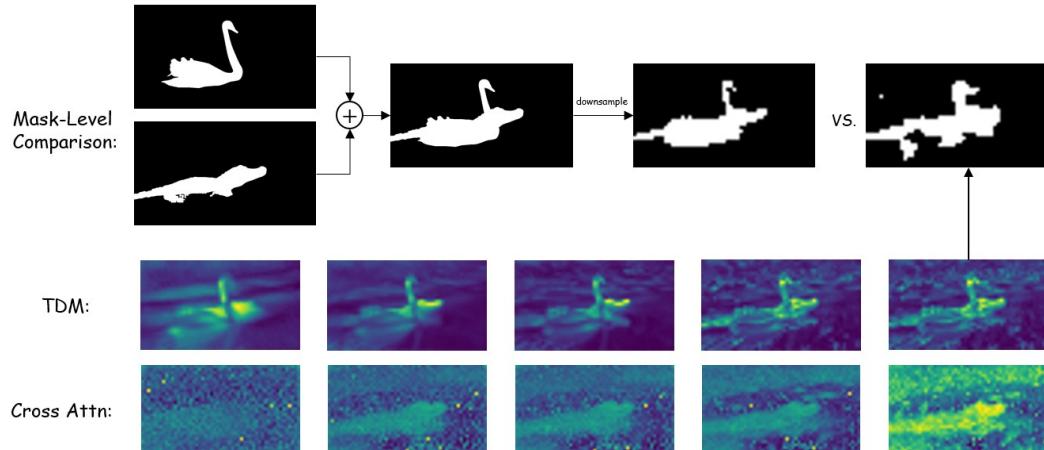


Figure 12: Post-hoc analysis of structural localization signals.

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188

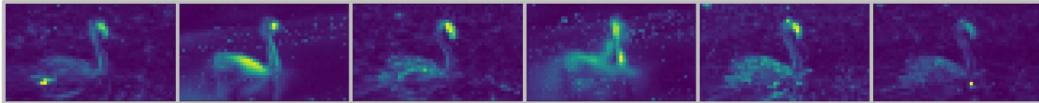
D LIMITATIONS AND FUTURE WORK

1189
1190 While our method demonstrates strong performance in shape-aware image editing, it also comes
1191 with certain limitations that suggest directions for future work.
11921193
1194

D.1 FAILURE CASES

1195
1196 Our method can be sensitive to prompt ambiguity and imprecise editing instructions. Since the edit-
1197 ing behavior is driven entirely by prompt-guided inversion and denoising trajectories, the quality
1198 of the editing outcome depends on how clearly the intended modification is specified in the text
1199 prompt. When editing instructions are vague, lack clear semantic targets, or have low discriminative
1200 specificity, the model may struggle to determine where and how strongly to apply the modification.
1201 This often leads to weak, diffuse, or inconsistent edits, particularly in cases where the intended
1202 change is not explicitly specified by the prompt. For example, prompts that describe abstract trans-
1203 formations or rely on implicit assumptions about the editing target may result in edits that do not
1204 match user expectations (see Figure 13). Clear and well-defined editing descriptions that explicitly
1205 identify the object to be modified and the desired transformation are therefore important for reliable
1206 performance.
12071208
1209 Source Prompt: A small stack of smooth black stones stands on the right side of a calm
1210 reflective surface, with a single round stone placed beside it on the left. The sky glows with
1211 warm pastel colors at sunset, and the water reflects the tranquil scene.
12121213
1214 Target Prompt: A small stack of smooth [white] stones stands on the right side of a calm
1215 reflective surface, with a single round stone placed beside it on the left. The sky glows with
1216 warm pastel colors at sunset, and the water reflects the tranquil scene.
12171218
1219 **Figure 13: Failure Case**
12201221
1222

D.2 EXTENDING TO VIDEO EDITING

1223
1224 We also explore extending our shape-aware editing framework to the video domain using Wan
1225 2.1 (Wan et al., 2025), an open-source video generation model that uses Rectified Flow. While
1226 our method can in principle be applied to all frames, we find that the temporal dimension intro-
1227 duces a major challenge, where the TDM becomes much less stable and effective when extended
1228 across time, as shown in Figure 14. In particular, the spatial editing regions indicated by TDM often
1229 fluctuate across frames, leading to inconsistent or incomplete transformations in the resulting video.
1230 Since a well-defined and temporally consistent TDM is crucial for successful editing, future work
1231 may consider strategies such as temporally-aware TDM construction, or explicit disentanglement of
1232 spatial and temporal components in the denoising trajectory.
12331234
1235 **Figure 14: TDM of Wan2.1 video editing at a single timestep across different frames.**
1236

1242 E MORE EDITING RESULTS
12431244 We present additional shape-aware editing results in Figure 15 and Figure 16. We also present
1245 general task editing results in Figure 17.
12461247 F THE USAGE OF LARGE LANGUAGE MODELS
12481249 In this paper, the usage of the LLM mainly falls into the following aspects:
12501251

- 1252 • **Grammar checking and format optimization:** In the paragraphs of the paper, LLMs are
1253 used for grammar error checking and format checking of charts and figures.
1254
- 1255 • **Language polishing:** The text description part of the paper uses LLMs to polish and opti-
1256 mize the language expression.
1257
- 1258 • **Prompt Generation:** We use Qwen-2.5-VL to generate paired source and edit prompts for
1259 images in our constructed benchmark.”
1260
- 1261 • All authors are responsible for the content generated by the LLMs.
1262

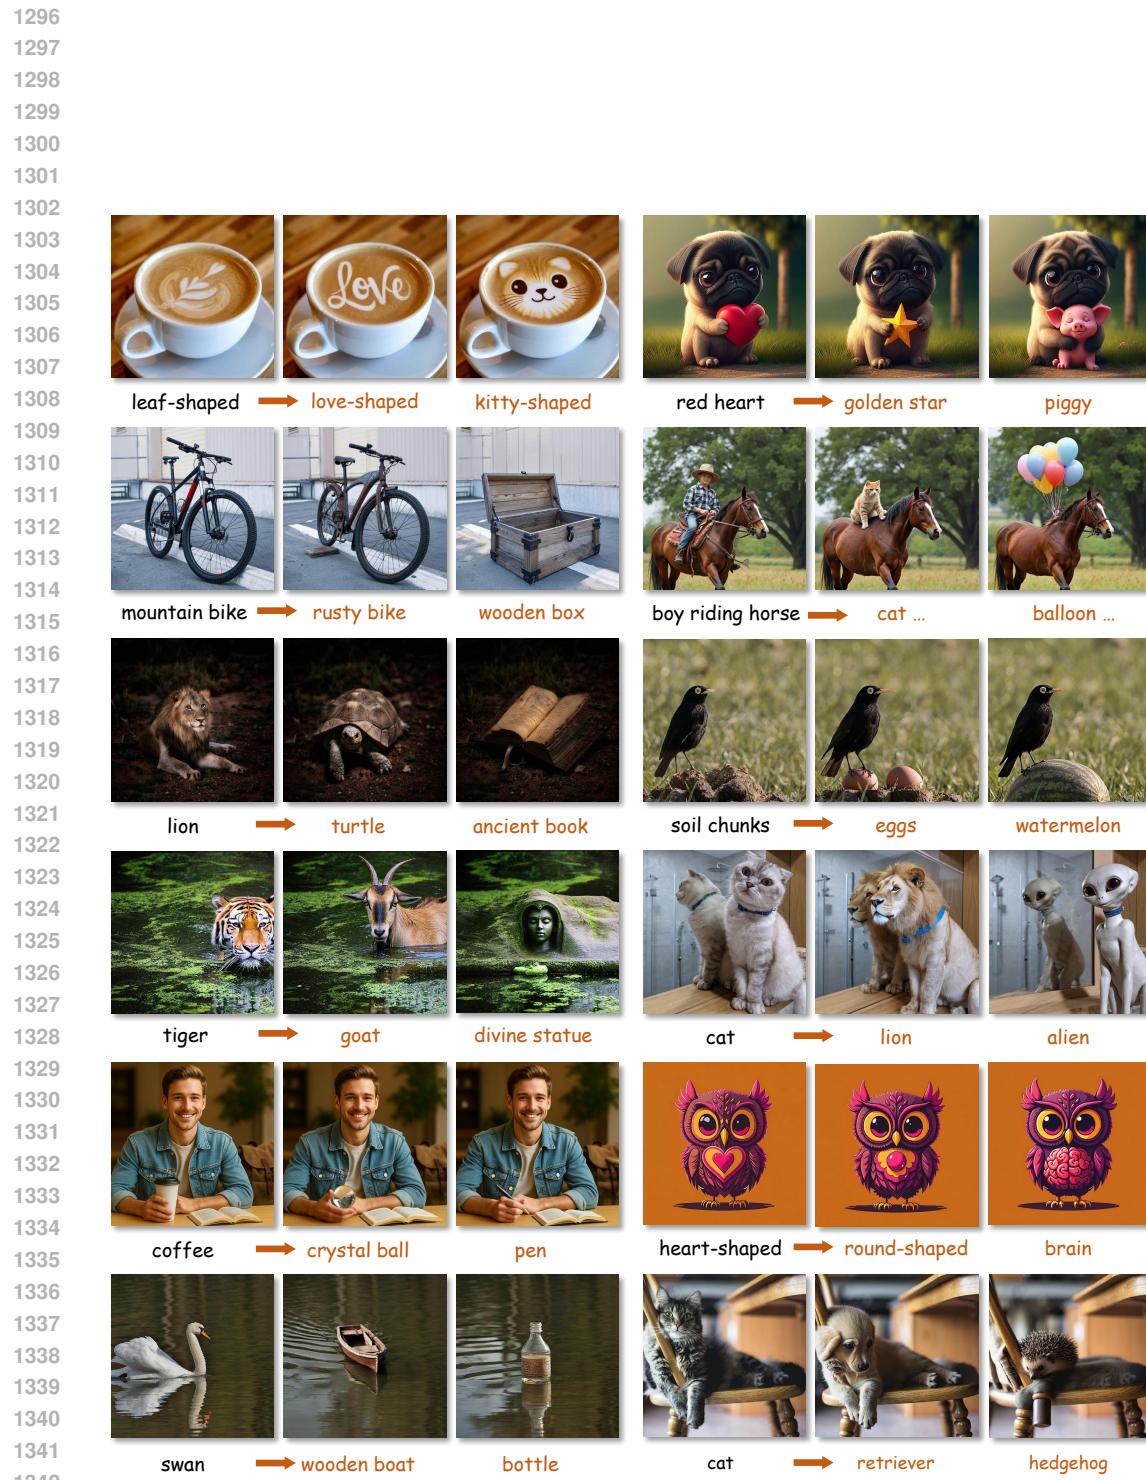


Figure 15: Additional Editing Results

1343
1344
1345
1346
1347
1348
1349

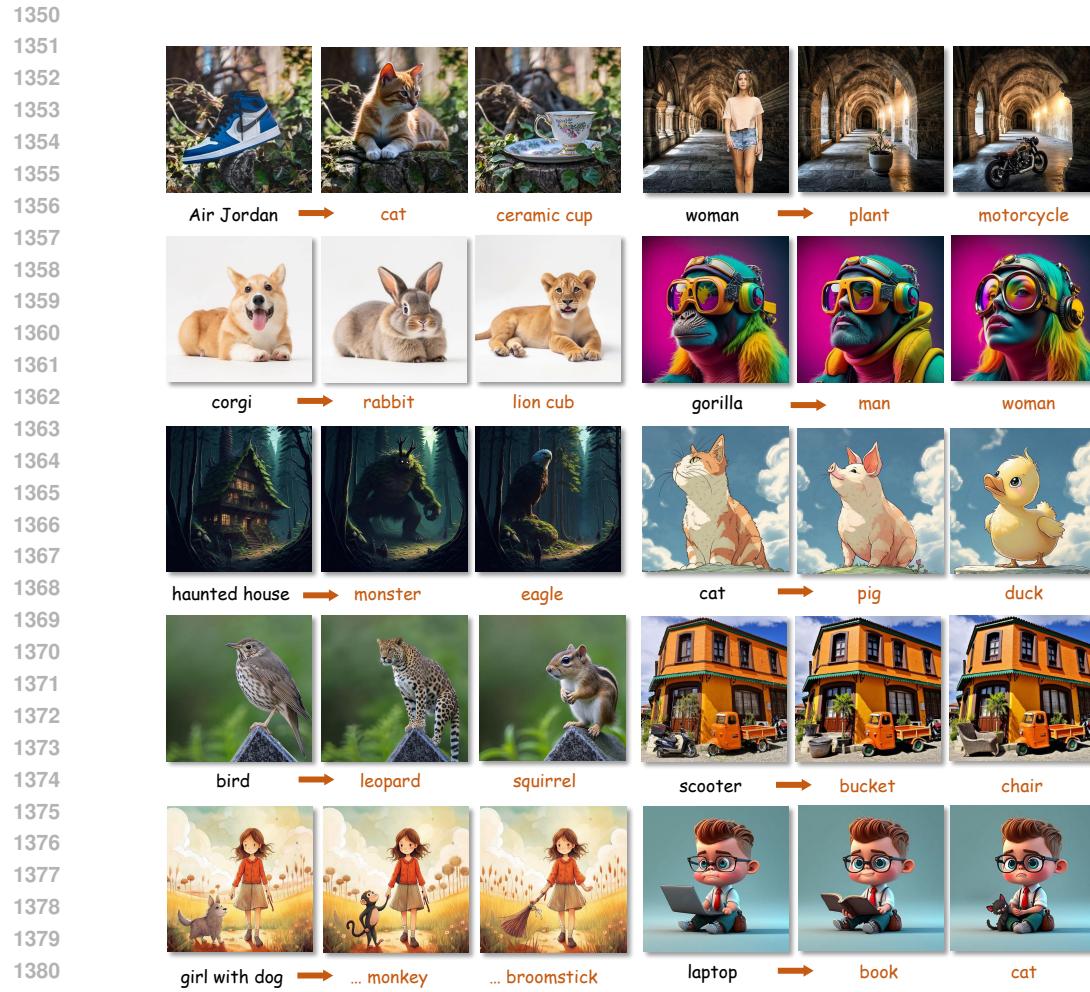


Figure 16: Additional Editing Results

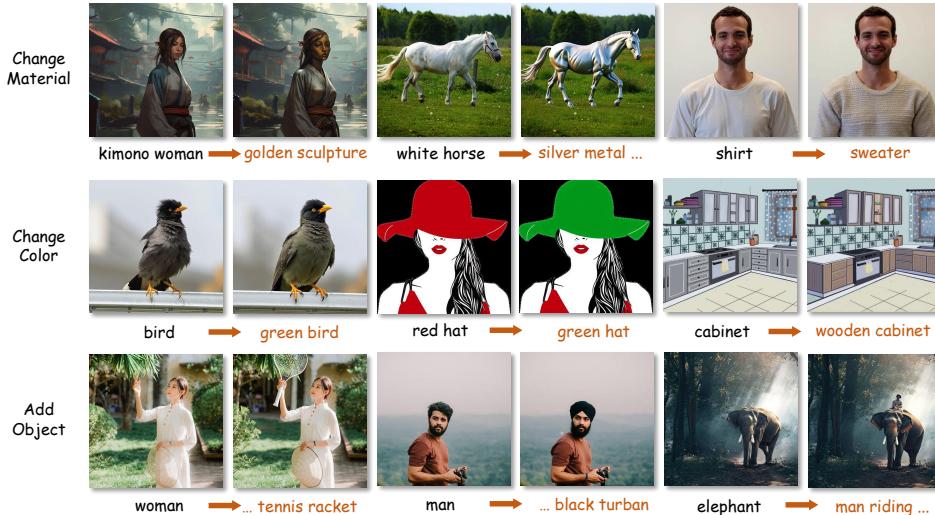


Figure 17: Additional Editing Results

Table 5: Example 1: Image–prompt pairs in *ReShapeBench*

1404	Image	Prompts
1405		
1406		
1407		
1408		
1409		
1410		
1411		
1412		
1413		
1414		
1415		
1416		
1417		
1418		
1419		
1420		
1421		
1422		
1423		
1424		
1425		
1426		
1427		
1428		
1429		
1430		
1431		
1432		
1433		
1434		
1435		
1436		
1437		
1438		
1439		
1440		
1441		
1442		
1443		
1444		
1445		
1446		
1447		
1448		
1449		
1450		
1451		
1452		
1453		
1454		
1455		
1456		
1457		

Table 6: Example 2: Image–prompt pairs in *ReShapeBench*

1427	Image	Prompts
1428		
1429		
1430		
1431		
1432		
1433		
1434		
1435		
1436		
1437		
1438		
1439		
1440		
1441		
1442		
1443		
1444		
1445		
1446		
1447		
1448		
1449		
1450		
1451		
1452		
1453		
1454		
1455		
1456		
1457		

1458
 1459
 1460
 1461
 1462
 1463
 1464
 1465
 1466
 1467
 1468
 1469
 1470
 1471
 1472
 1473
 1474
 1475

Table 7: Example 3: Image–prompt pairs in *ReShapeBench*

	Image	Prompts
1476		
1477		Source Prompt: A young boy crouches low on a quiet street, one hand resting gently on a worn soccer ball, his gaze steady and thoughtful under the soft glow of late afternoon light. His teal striped sweater, red shorts with white stripes, scuffed sneakers, and the cracked leather surface of the ball stand out sharply against the rough asphalt beneath, every thread, crease, and stitch rendered with photorealistic clarity. Behind him, a gentle blur of weathered brick walls and distant buildings melts into warm bokeh, hinting at a humble neighborhood bathed in golden-hour haze. The entire scene feels quietly powerful — where childhood dreams meet everyday resilience, serene yet deeply human, captured with cinematic depth and organic warmth.
1478		Edit Prompt 1: A young boy crouches low on a quiet street, one hand resting gently on a worn leather backpack beside his knee, his gaze steady and thoughtful under the soft glow of late afternoon light. His teal striped sweater, red shorts with white stripes, scuffed sneakers, and the frayed straps and faded stitching of the backpack all stand out sharply against the rough asphalt beneath, every texture rendered with photorealistic clarity. Behind him, a gentle blur of weathered brick walls and distant buildings melts into warm bokeh, hinting at a humble neighborhood bathed in golden-hour haze. The entire scene feels quietly nostalgic — where school days meet quiet contemplation, serene yet deeply human, captured with cinematic depth and organic warmth.
1479		
1480		
1481		
1482		
1483		
1484		
1485		
1486		
1487		
1488		
1489		
1490		
1491		
1492		
1493		
1494		
1495		
1496		
1497		
1498		
1499		
1500		
1501		
1502		
1503		
1504		
1505		
1506		
1507		
1508		
1509		
1510		
1511		