

TALK2MOVE: Reinforcement Learning for Text-Instructed Object-Level Geometric Transformation in Scenes

Jing Tan^{1,4*} Zhaoyang Zhang¹ Yantao Shen¹ Jiarui Cai²
 Shuo Yang¹ Jiajun Wu³ Wei Xia¹ Zhuowen Tu¹ Stefano Soatto¹
¹AWS Agentic AI ²Amazon Web Services ³Amazon Robotics ⁴CUHK

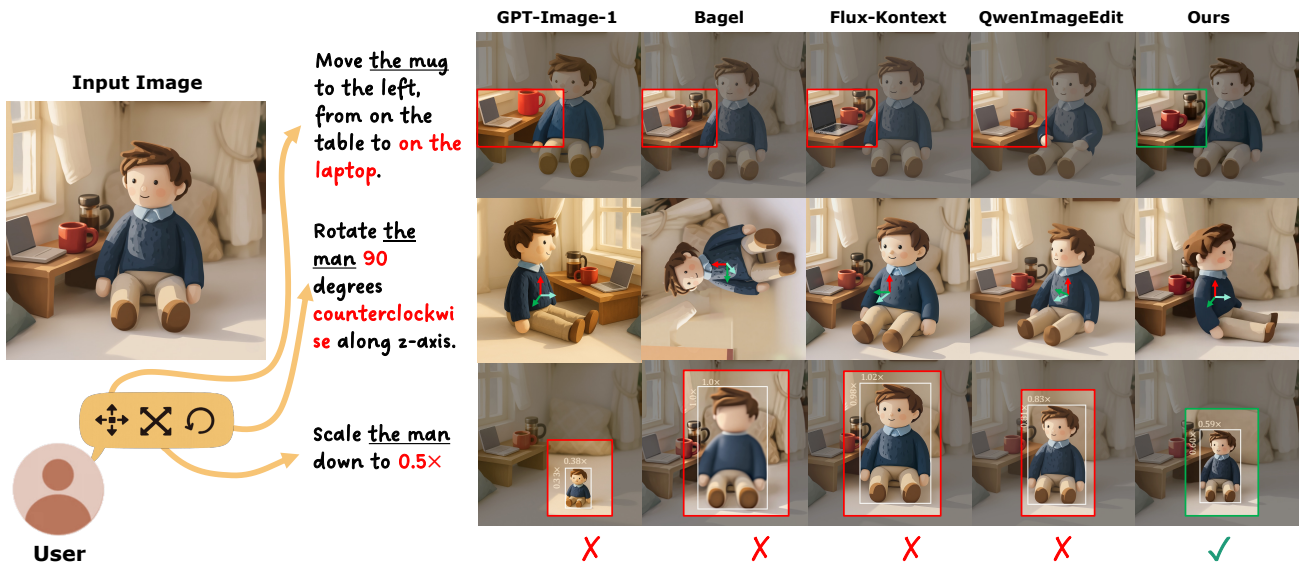


Figure 1. We introduce **TALK2MOVE**, a text-guided scene editing model for object-level geometric transformation, focusing on object translation, rotation and resizing, achieving superior results over current SOTA image editing models.

Abstract

We introduce *Talk2Move*, a reinforcement learning (RL) based diffusion framework for text-instructed spatial transformation of objects within scenes. Spatially manipulating objects in a scene through natural language poses a challenge for multimodal generation systems. While existing text-based manipulation methods can adjust appearance or style, they struggle to perform object-level geometric transformations—such as translating, rotating, or resizing objects—due to scarce paired supervision and pixel-level optimization limits. *Talk2Move* employs Group Relative Policy Optimization (GRPO) to explore geometric actions through diverse rollouts generated from input images and lightweight textual

variations, removing the need for costly paired data. A spatial reward guided model aligns geometric transformations with linguistic description, while off-policy step evaluation and active step sampling improve learning efficiency by focusing on informative transformation stages. Furthermore, we design object-centric spatial rewards that evaluate displacement, rotation, and scaling behaviors directly, enabling interpretable and coherent transformations. Experiments on curated benchmarks demonstrate that *Talk2Move* achieves precise, consistent, and semantically faithful object transformations, outperforming existing text-guided editing approaches in both spatial accuracy and scene coherence.

1. Introduction

Object transformation is essential in scene image editing, enabling precise control over spatial layout and object inter-

* The work was done during an internship at AWS Agentic AI. Jing Tan is now with the Chinese University of Hong Kong. Correspondence to ozhaozha@amazon.com

actions. These geometric manipulations let users adjust content to match their intent, such as “*move the mug to the left*” or “*rotate the man 90 degrees counterclockwise.*” Guiding these transformations through natural-language instructions is particularly valuable, as text provides an intuitive and low-barrier interface that makes spatial editing accessible to a broad range of users.

Current graphics-based editing methods [19, 31, 33, 48] often rely on 2D/3D primitives as control signals to assist scene editing. This requires human intervention and domain expertise in 2D/3D generation, making it **not user-friendly** nor in natural form of interaction for general users. Moreover, these approaches typically **demands large amounts of high-quality paired data** for SFT training, yet such data are rare and expensive to obtain. Spatial-manipulation pairs are particularly scarce: video datasets or 3D simulations provide only limited examples and are costly to scale. Top-down modeling of image content exists [7, 37] which requires small training data but their performances are not meeting the practical requirement. Moreover, SFT-based methods rely on pixel-level MSE loss that struggles to disentangle objects from scenes, leading to limited spatial control.

In this work, we introduce TALK2MOVE, the first RL-based framework that leverages **text guidance** to perform geometric object-level transformations for scene editing. Built upon a flow-based GRPO paradigm, TALK2MOVE perturbs diffusion trajectories through stochastic noise injection to explore diverse spatial transformations during training. It scales efficiently by generating diverse rollouts with simple prompt variations, eliminating the need for costly paired data. To disentangle objects from scene backgrounds, TALK2MOVE employs a spatially grounded reward model that directly evaluates object displacement, rotation, and scaling behaviors beyond pixel-level similarity, leading to interpretable and geometry-aware optimization objectives.

However, conventional dense rollouts in GRPO training are computationally expensive, and not all denoising steps contribute equally to learning. Therefore, we introduce a step-importance measure that quantifies how each perturbed denoising step influences the task reward. Building on this, we propose step-wise active sampling, which adaptively selects informative denoising steps through lightweight off-policy reward evaluation and skips redundant ones via skip connections. This design substantially improves training efficiency by $2\times$ while preserving the reward robustness.

To summarize, our contributions are four-fold.

- To the best of our knowledge, we are the first to formulate the text-guided geometric object transformation problem within a reinforcement learning framework, enabling the editing model to follow textual instructions for precise spatial manipulation.
- Our pipeline is data-efficient compared to SFT-based methods, reducing dependence on costly paired annotations.

We design a spatial-aware reward model that separates target objects from the scene to improve geometric transformation learning, and introduce an early-exit mechanism for flow-based GRPO to focus training on the most informative steps and lower sampling cost.

- We establish an effective data collection pipeline that provides high-quality data for geometric object transformations, and we demonstrate that this pipeline can be scaled to produce large quantities of input pairs suitable for GRPO online training.
- Experiments demonstrate that TALK2MOVE significantly improves spatial accuracy and scene coherence under limited budgets, achieving state-of-the-art spatial transformation performance across both open-source and proprietary editing systems.

2. Related Work

Geometry-conditioned Image Editing. This line of research manipulates images through explicit spatial or structural constraints. Point-based methods [23, 26, 31, 33, 39, 40] enable interactive control via handle–target interactions, yet they typically require manual point specification and struggle with high-level semantic instructions. To incorporate stronger geometric priors, 3D-aware approaches [3, 47, 48] lift 2D images into explicit 3D representations for transformation. However, these multi-stage lifting and rendering pipelines often introduce cumulative errors and compromise image fidelity. Alternatively, layout-based editing [5, 28] utilizes explicit 3D layout to guide scene structure manipulation, though it necessitates the often-costly acquisition of precise 3D layouts. ORIGIN [30] leverages a 3D orientation condition and iteratively refines the latent via “test-time reflection” to satisfy the rotation reward. Overall, these geometry-dependent methods rely heavily on explicit spatial primitives and manual intervention, limiting their scalability for natural scene-level object manipulation.

Text-based Image Editing. Recent image generation and editing methods fall broadly into two families: diffusion-based models and VLM-connected frameworks. Diffusion-based methods [36, 38] generate images via denoising, with recent flow-matching models such as Flux and Flux.1 Kontext [15, 16] improving convergence and edit consistency through in-context learning. However, relying solely on text tokenizers limits their ability to follow fine-grained spatial instructions.

Another line of work integrates pre-trained VLMs with diffusion decoders via lightweight adapters. Methods such as MetaQueries [34], Seed-X [9], and Emu2 [41], as well as unified transformer frameworks like MoT and Bagel [4, 22], enhance prompt following but often struggle to preserve visual coherence when only ViT-based encoders are used. QwenImageEdit [44] addresses this by combining VLM semantic features with VAE reconstructive features for stronger

alignment between instruction following and visual fidelity. We adopt QwenImageEdit as our backbone for spatial manipulation tasks.

Reinforcement Learning for Visual Generation. Reinforcement learning [6, 11, 29, 45] for diffusion-based generation typically formulates the denoising process as a Markov Decision Process (MDP). Reward-Weighted Regression [35] models generation as a one-step MDP and shifts the image distribution toward high-reward samples. DDPO [2] extends this to a multi-step MDP by treating each denoising step as an action, making it well-suited for deterministic samplers without requiring repeated responses from the same query. Recent works adapt GRPO to diffusion models. Flow-GRPO [24] and DanceGRPO [46] introduce stochastic drift into DDPM and Rectified Flow samplers, enabling multiple rollout paths around the deterministic trajectory. Subsequent improvements [17, 20, 49] reduce training cost via sliding-window optimization or tree-based trajectory reorganization with reward backpropagation. Although effective, these methods remain computationally expensive during sampling. In this work, we further improve time efficiency by introducing an early-exit mechanism to accelerate rollout generation. TempFlow-GRPO [12] shows that temporal reweighting improves general GRPO stability. Our work instead focuses on identifying task-dependent transformation stages for object-level geometry. RL-based image editing methods [18, 21, 27] exist but they are not focused on the object transformation tasks.

3. Geometric Transformation for Object-Level Scene Editing

Problem Formulation. We study text-guided geometric transformation for object-level scene editing, where the goal is to modify an object’s position, orientation, and scale in an image according to a text instruction, while keeping the scene static and unedited regions consistent. We focus on three fundamental spatial operations: object translation, rotation, and resizing, which together span the core dimensions of geometric object transformation.

Formally, we define a spatial editing model as a mapping $f_\theta : (I, T) \rightarrow I'$, where I is the input scene image, T is a text instruction, and I' is the edited output. To regularize the instruction–action space, we use a set of predefined transformation templates t_k for translation, rotation, and resizing. Each template specifies a target object and a spatial relation (e.g., “move the mug to the left”), providing a consistent basis for supervision and evaluation.

Standardized Transformation Templates. For object translation, the template defines both the movement direction (left, right, up, down) and the target spatial reference (e.g., “on top of the sofa”, “near the violin”). For object rotation, it specifies the rotation axis (x, y, z), direction (clockwise or counterclockwise), and angle ($45^\circ, 90^\circ, 135^\circ, 180^\circ$). The ro-

tation axes are defined in the object’s local coordinate frame, which aligns with the object’s inherent orientation in the scene. For object resizing, the template defines the scaling ratio (e.g., $1.25\times, 1.5\times, 2\times, 3\times, 4\times$). Together, these standardized templates constrain the open object transformation problem into a well-defined and verifiable space suitable for reproducible learning.

4. Geometric Transformation Data Curation

GRPO-style RL training requires only input samples consisting of a reference image and a text prompt. We take the object translation task as an example to illustrate the data curation process.

Reference Image Generation. We prompt a large language model (LLM) to generate textual scene descriptions that specify realistic scenes with a background and multiple sparsely distributed movable objects, and then employ open-source text-to-image models [16, 44] to synthesize the corresponding reference images.

Instruction Generation. For each generated image, we use a vision-language model (VLM) to annotate spatial editing instructions with predefined templates. These annotations are parsed into a structured format that preserves geometric attributes such as object location, movement direction, rotation angle, and resizing ratio for reward computation. Given a reference image, we can expand the dataset by varying textual prompts to generate many training samples and extend this process to other object transformation tasks, substantially reducing data creation costs compared to paired supervision. In total, we construct 3,200 samples over 800 unique images for object translation. We leave large-scale data expansion to future work, as the current dataset is sufficient for training.

Target Image Synthesis for SFT Cold Start. Current image editing backbones struggle with precise spatial manipulation, so we first apply an SFT cold-start to learn basic spatial-editing priors from a small amount of high-quality ground-truth data. However, constructing paired spatial-editing data is challenging, as real-world datasets rarely provide matched before–after examples for specific geometric transformations. To address this, we synthesize supervision with tailored strategies for translation, rotation, and resizing. For object translation and rotation, we use the API-based video generation model [8] to simulate physically plausible object motion in diverse scenes. Conditioned on a reference image as the first frame and a manipulation prompt, the model generates short videos depicting the desired transformation. After filtering by spatial accuracy and visual consistency, we obtain 800 valid translation pairs and 43 high-quality rotation pairs. For object resizing, where current video generation models perform poorly, we instead rely on open-source image editing models [16, 44] to synthesize coarse up/down scaling. The edited images are paired with their references

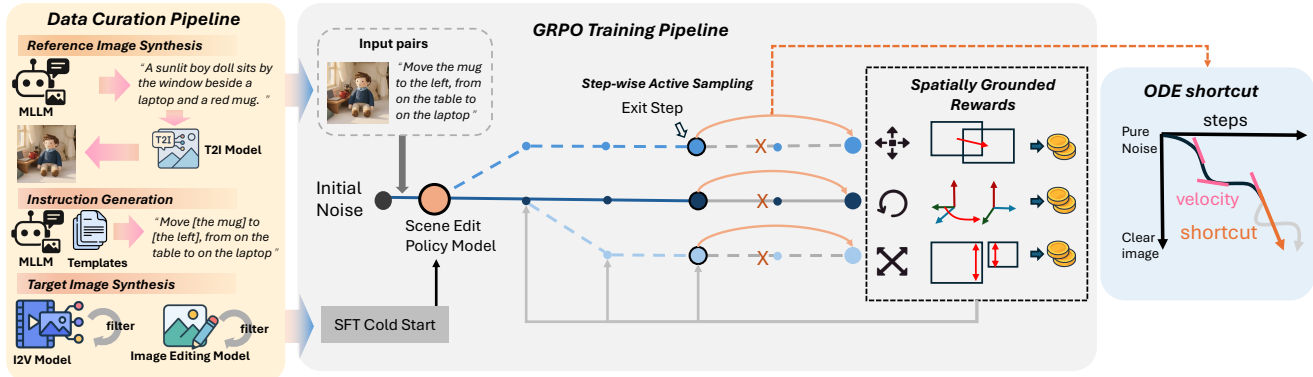


Figure 2. **The pipeline of TALK2MOVE.** TALK2MOVE streamlines a GRPO-style reinforcement learning pipeline tailored for flow-based image editing. Starting from an initial noise sample, stochastic perturbations are injected at each diffusion step to generate diverse sampling trajectories. Spatially grounded rewards from specialist models, which explicitly evaluate object-level geometric changes, are then used to compute group-relative advantages for policy gradient updates.

and filtered via perceptual matching, yielding 110 diverse resizing pairs.

5. Method

5.1. Preliminary

Diffusion as a Markov Decision Process. The reverse denoising process of a diffusion model can be naturally formulated as a Markov Decision Process (MDP). We define the state at time step t as $s_t = (c, \mathbf{x}_t)$, where $\mathbf{x}_t \in \mathbb{R}^d$ is the latent representation and c is the conditioning context (e.g., text or image features). Each denoising step corresponds to an RL action a_t , representing the predicted noise or velocity vector that transitions the current state \mathbf{x}_t to the subsequent latent \mathbf{x}_{t-1} . The state transition probability $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, c)$ characterizes the likelihood of this incremental reconstruction. In this framework, the entire generation trajectory from $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ to the clean image \mathbf{x}_0 constitutes a single RL rollout.

From ODE to Stochastic Rollouts. Standard flow-matching models typically solve a Probability Flow Ordinary Differential Equation (ODE): $dx = v_\theta(\mathbf{x}, t, c)dt$, which yields a deterministic trajectory for any initial noise \mathbf{x}_T . However, Group Relative Policy Optimization (GRPO) [?] relies on contrasting multiple samples within a group to estimate relative advantages. To facilitate the necessary exploration and variance, [24, 46] inject stochastic perturbations into the velocity field, effectively transforming the deterministic ODE into a Stochastic Differential Equation (SDE): $dx = v_\theta(\mathbf{x}, t, c)dt + \sigma(t)d\mathbf{w}$, where $d\mathbf{w}$ denotes a Wiener process. This stochastic formulation enables sampling a group of G distinct rollouts for the same input c , allowing the model to iteratively refine the velocity field toward trajectories that satisfy verifiable geometric rewards.

5.2. Training Pipeline

In this section, we describe the overall training pipeline of TALK2MOVE. We first use an SFT cold start to equip the editing model with basic geometric transformation ability, and then present the RL stage in detail.

Cold Start. We use lightweight LoRA fine-tuning with limited ground-truth data to equip the diffusion backbone with basic spatial editing capabilities. Annotated image pairs from all three tasks are combined to train a unified SFT checkpoint. LoRA adapters are applied to attention projections, normalization, and linear layers, while the text encoder, VAE, and ViT remain frozen. This initialization introduces coarse spatial priors, improving rollout stability and sample quality during GRPO training.

RL implementation. To achieve data-efficient and spatially accurate scene manipulation, TALK2MOVE adopts a GRPO-style reinforcement learning framework tailored for diffusion policies. Starting from an initial noise sample, we follow the flow-based GRPO paradigm [24, 46], injecting stochastic perturbations at each diffusion step to form groups of diverse sampling trajectories. These rollouts capture a rich set of spatial transformations conditioned on the input prompt and diffusion dynamics. Spatially grounded rewards are provided by off-the-shelf specialist models that disentangles objects explicitly from the scene and evaluates object-level transformations, from which we compute group-relative advantages to update model parameters via policy gradients.

5.3. Efficient GRPO with Early Exit

Dense rollouts in GRPO incur significant computational overhead, yet not all denoising steps contribute equally, and some even harm sampling performance, highlighting the need to prioritize steps with stronger learning signals.

Off-policy Step Evaluation. Inspired by reasoning LLMs [42], which emphasize tokens with high intrinsic




Figure 3. **Early-exit step sampling:** (a) is the full trajectory sampling [24, 46]; (b) [17, 20] use sliding windows to reduce the training steps; (c) we introduce early-exit active sampling that selects the informative steps (red) and uses shortcuts to bypass other steps, significantly reducing the sampling time.

uncertainty, we propose to focus on diffusion steps exhibiting high *extrinsic* uncertainty. Since the noise variance in diffusion is fixed by the scheduler, we measure uncertainty through the variance of rollout rewards at each step. High reward variance indicates strong exploration potential and thus higher learning utility, analogous to high-entropy tokens in language RL. In addition, according to MixGRPO [17] and our empirical findings, early denoising steps largely determine global layout, while later steps refine fine-grained details. Therefore, the step importance is defined as the reward variance of rollouts when exiting from each denoising step. In practice, starting from step 0, we incrementally perturb small sets of denoising steps on a few input images (2–4) and compute rollout rewards. This efficient calibration, requiring only a single GPU, yields a task-specific, step-aware reward distribution. The optimal exit step is then selected as the last perturbed step achieving the maximum reward variance.

Active Step Sampling. With the exit step identified, we introduce *step-wise active step sampling*, which takes ODE shortcuts at the exit step to bypass redundant later steps, thereby improving both rollout and training efficiency. Flow-based GRPO approaches apply ODE updates at every denoising step, resulting in a time complexity of $T(t_{\text{sample}} + t_{\text{optim}})$ for T steps. The sliding-window mechanism in [17] alleviates this by optimizing only a subset of steps per iteration. We further enhance efficiency by directly denoising from the exit step K to the final step T using model predictions (see Fig. 3), effectively shortening the trajectory. This reduces total computation to approximately $K(t_{\text{sample}} + t_{\text{optim}})$, where $K \leq T$. As long as $K < T$, our method achieves consistently higher training efficiency while maintaining comparable task performance.

5.4. Spatially Grounded Reward Design

The choice of reward model is one of the most important factors in GRPO training, as it provides critical information

for the policy update direction. Previous works [24, 27] employ holistic, image-level metrics such as aesthetic rewards, CLIP-based alignment, or VLMs fine-tuned for common editing tasks. These rewards serve as a coarse measure of overall image quality or human preference. We extend the reward model to fine-grained, spatially grounded rewards tailored for spatial manipulation editing.

To evaluate the correctness of geometric object transformation, we first localize the prompted object in both the reference and edited images using text-driven segmentation [13, 25] to obtain segmentation masks and bounding boxes. Each manipulation task defines a specific spatial reward. For object translation, we compute the relative displacement between the reference and edited object centers following the GenEval [10] protocol. For rotation, we use Orient-Anything [43] to estimate orientations for object crops from the scene image and evaluate alignment along the prompted axis, direction, and angle. For resizing, we compare the normalized size ratio differences of reference and edited object bounding boxes. Although each reward differs in formulation, all share a unified goal, which is to measure the alignment between the prompted and achieved spatial transformations in normalized coordinate space. These task-specific rewards provide fine-grained feedback to guide policy updates toward spatially accurate editing behavior.

6. Experiments

6.1. Implementation Details

For cold start, we train Qwen-Image-Edit [44] with rank-64 LoRA layers for 3,000 iterations and with a learning rate of $1e-4$. For GRPO training, we build upon the FlowGRPO [24] baseline, with a sample noise level of 1.0 and a clip range of $2e-4$, on 16 H200 GPUs. The total training time for each subtask is ~ 160 GPU hours.

6.2. Evaluation Metrics

To evaluate the spatial accuracy for scene editing, we introduce task-specific metrics.

Translation. We report the Translation Distance (Trans. Dist.) that measures the relative distance between object center points in the reference and edited image; and the translation Accuracy (Acc.) that measures the success rate. An edit is considered successful only if it passes all four criteria: 1) correct movement: following GenEval [10], we verify whether the detected object in the edited image has moved in the specified direction relative to its position in the reference image; 2) object identity preservation: the edited object must remain the same object (CLIP similarity > 0.75); 3) no duplication: the object should no longer appear at its original location (similarity < 0.95); 4) scene preservation: the overall scene should remain consistent (background L1 distance ≤ 0.2).

Table 1. **Quantitative comparison** of object transformation tasks on our curated **synthetic** test benchmark in terms of editing accuracy, average translation distance, and editing errors. The results are obtained from unified SFT followed by task-specific GRPO training. We also report **human evaluation** results in terms of winning rate. We note that win rate does not capture second-best or near-tie preferences, and thus may underestimate methods that produce competitive but less visually salient results, such as GPT-based models.

Methods	Translation			Rotation			Resize		
	Trans. Dist. \uparrow	Acc. \uparrow	Human Win Rate \uparrow	Rot. Err. \downarrow	Acc. \uparrow	Human Win Rate \uparrow	Scale Err. \downarrow	Acc. \uparrow	Human Win Rate \uparrow
GPT-Image-1 [32]	0.5416	64.29%	26.25%	0.4293	2.33%	6.25%	0.3501	5.08%	1.39%
Flux-Kontext [16]	0.0499	4.41%	1.25%	0.4259	6.82%	16.25%	0.4318	1.67%	15.28%
Bagel [4]	0.1705	14.49%	2.50%	0.3240	13.64%	1.25%	0.4790	0.00%	8.33%
QwenImageEdit [44]	0.2551	32.86%	12.50%	0.4129	9.30%	7.50%	0.4203	7.50%	11.11%
Ours	0.6667	76.67%	57.50%	0.2861	29.55%	68.75%	0.3894	9.17%	63.89%

Table 2. **Quantitative comparison** of object transformation tasks on curated **real** images from OpenImages-V6 [14] in terms of editing accuracy, average translation distance and editing errors.

Methods	Translation		Rotation		Resize	
	Trans. Dist. \uparrow	Acc. \uparrow	Rot. Err. \downarrow	Acc. \uparrow	Scale Err. \downarrow	Acc. \uparrow
GPT-Image-1 [32]	0.3237	23.08%	0.1774	25.00%	0.6898	7.14%
Flux-Kontext [16]	0.1992	26.92%	0.2567	6.67%	0.7120	1.79%
Bagel [4]	0.2258	15.38%	0.3700	6.25%	0.7159	3.57%
QwenImageEdit [44]	<u>0.3725</u>	<u>42.31%</u>	0.2101	25.00%	<u>0.6048</u>	1.79%
Ours	0.5196	53.85%	<u>0.1997</u>	31.25%	0.5947	7.14%

Rotation. We compute rotation accuracy by verifying whether the edited rotation angle lies within $\pm 20^\circ$ of the target angle. We also report rotation error, defined as the average normalized difference between the target (prompted) angle and the predicted (edited) angle.

Resizing. We measure resizing accuracy by checking whether the edited size falls within a $\pm 10\%$ tolerance of the target scaling ratio. We also report the scaling error, defined as the average difference between the prompted scaling ratio and the actual edited ratio.

6.3. Quantitative Evaluation

We show the quantitative comparison with state-of-the-art image editing methods, including diffusion-based editing [16], unified model [4], our baseline [44] and GPT-Image-1. Tab. 1 reports TALK2MOVE’s result on our synthetic evaluation benchmark for three geometric transformation tasks, featuring 100 input samples under each transformation type. Our model outperforms open-sourced baselines among most metrics. We also sample 85 real images from OpenImagesV6 [14] for evaluation in Tab. 2 and results prove our model’s effectiveness on real images.

Moreover, a user study is also conducted to further evaluate the editing accuracy and scene coherence. We invite 15 users with 3-year + expertise in multi-modality generation to select a winner among editing candidates from each method and report the winning rate over 30 questions. Results in Tab. 1 demonstrates that TALK2MOVE achieves the highest winning rate with a sweeping edge.

Table 3. **Ablation study on SFT and RL** under object translation task in terms of translation distance, accuracy and image L1 distance.

Method	Trans. Dist. \uparrow	Acc. \uparrow	L1 \downarrow
Flux-Kontext	0.0499	4.41%	0.5704
Flux-Kontext+SFT	0.3692	38.81%	0.5403
Ours	0.4798	47.14%	0.5302
QwenImageEdit	0.2551	32.86%	0.5834
QwenImageEdit+SFT	0.5953	67.14%	0.2562
Ours	0.6667	73.13%	0.2012
QwenImageEdit+SFT [†]	0.2257	26.67%	0.5928
Ours[†]	0.6507	73.33%	0.2629

[†] Trained with 1/10 of the original training samples.

Table 4. **Ablation on active step sampling** under translation task, efficiency measured in seconds(s).


Sampling	NFE $_{\pi_{old}}$	NFE $_{\pi_{\theta}}$	Sample Train	Total	Trans. Dist.	Acc.
Full	10	10	45.78	126.25	172.32	0.6602 69.12%
Sliding window	10	4	45.63	55.69	101.61	0.5983 67.14%
Ours	4	4	30.95	56.02	87.27	0.6667 76.67%

6.4. Qualitative Evaluation


In Fig. 5, we provide multiple qualitative examples including both real and synthetic inputs to showcase TALK2MOVE’s ability in object transformation compared to other baseline methods. We find that GPT-Image-1 performs transformation edits well, but its result do not preserve the scene details from the reference image, usually exhibits a warmer tone and brighter lighting. Compared to the baseline methods, TALK2MOVE executes the instruction more accurately and better preserves the original scene details.

6.5. Ablation Studies


SFT vs RL. Tab. 3 compares how SFT and RL contribute to spatial manipulation performance. We evaluate both approaches on two backbone models: the widely used editing backbone Flux-Kontext [16] and the more advanced Qwen-ImageEdit model [44], which integrates a VLM with a diffusion generator. To ensure a fair comparison, all experiments are trained solely on object translation data.



(a) Off-policy reward variance distribution for translation, rotation and resizing.



(b) Reward curves for different sampling strategies under the translation task.



(c) Ablative comparison of VLM reward vs. spatial reward under the rotation task.

Figure 4. Reward behavior across tasks: (a) reward variance distribution, (b) GRPO sampling strategies, and (c) reward model ablations.

Table 5. **Background ID Consistency** of three object transformation tasks in terms of image-level CLIP and L1 distance. The best and second best results are highlighted with **bold** and underline.

Method	Translation		Rotation		Resize	
	CLIP-im \uparrow	L1 \downarrow	CLIP-im \uparrow	L1 \downarrow	CLIP-im \uparrow	L1 \downarrow
GPT-Image-1 [32]	0.9318	0.4351	0.9453	<u>0.4136</u>	0.9405	0.4278
Flux-Kontext [16]	0.9660	0.5704	0.9821	0.5390	<u>0.9739</u>	0.5145
Bagel [4]	0.9570	<u>0.3308</u>	0.9356	0.4138	0.9790	0.3027
QwenImageEdit [44]	0.9728	0.5834	<u>0.9777</u>	0.6000	0.9724	0.5869
Ours	<u>0.9699</u>	0.2012	0.9717	0.2774	0.9705	<u>0.3841</u>

With a reasonably large training set (800 images), SFT provides a strong initial boost in editing accuracy. However, once SFT converges, applying RL on top of the SFT checkpoint further increases both the average translation distance and the final accuracy for both backbones.

To further assess the data efficiency of RL, we reduce the training data to one tenth of the original size (~ 80 annotated pairs and ~ 400 input-only pairs). As shown in the third block of Tab. 3, SFT fails to achieve meaningful gains under such limited data, whereas RL built upon the small-data SFT checkpoint still reaches performance comparable to the full-data setting. This demonstrates that even with very limited input data, a well-designed reward function allows RL to effectively learn spatial manipulation behaviors and achieve competitive editing accuracy.

Background Scene Preservation. A good editing model should preserve the background identity while manipulating the foreground. We evaluate scene consistency using the image-level L1 distance between the input and edited images. As shown in Tab. 5, our method achieves an L1 distance comparable to other state-of-the-art editing models, indicating that it does not compromise scene consistency during training.

The CLIP metric is relatively coarse and not sensitive to subtle background changes; only large deviations meaningfully indicate identity loss. This is also evident from our qualitative results, where methods with similar CLIP-im scores can still show almost the same level of background preservation. In contrast, GPT-Image-1 exhibits noticeably

lower CLIP similarity than the other methods. This is because GPT-Image-1 tends to heavily refine or alter the entire image, which is also evident in our qualitative results.

Off-policy Step Evaluation Results. We present the off-policy evaluation results for the three object transformation tasks. As shown in Fig. 4a, the maximum reward variance occurs at step 4 (out of 10) for both translation and resizing, while for rotation it appears at step 10. These different exit steps demonstrate that in flow-based diffusion, injecting the same Gaussian noise at different steps leads to varying reward sensitivity and exploration space across tasks.

Ablation on Active Step Sampling. We further study the time efficiency of our method. Tab. 4 reports the sampling and optimization time for a batch of 16 rollouts. Based on the off-policy step analysis, we adopt step 4 as the exit step for the translation task, resulting in only 4 sampled steps per rollout. We observe that our method reduces the overall iteration time by 14% compared to the sliding-window baseline (translation-sliding window), and by 49% compared to full sampling (translation-full). At the same time, despite using fewer sampled steps, our shortcut strategy still achieves higher editing correctness, outperforming both baselines in terms of Translation Distance and Accuracy.

We also plot the evaluation reward curves in Fig. 4b for both baselines and our method. Interestingly, taking shortcuts not only reduces computation but also accelerates convergence, further suggesting that focusing on the steps that provide larger geometric transformation space is beneficial for exploration and overall performance.

Ablation on Reward Model. We conduct an ablation comparing a VLM-based reward model with our spatial reward derived from a specialist model (Orient-Anything [43] + rule-based evaluation). Editing quality is assessed using Qwen2.5-VL-Instruct [1] to measure alignment between the edited image and the prompt. As shown in Fig. 4c, both reward models improve evaluation reward during GRPO training. However, the VLM-based reward often produces overly optimistic and unstable scores, yielding a rotation error of 0.3294 and an accuracy of 11.63% according to evaluation. In contrast, our spatial reward leads to a lower

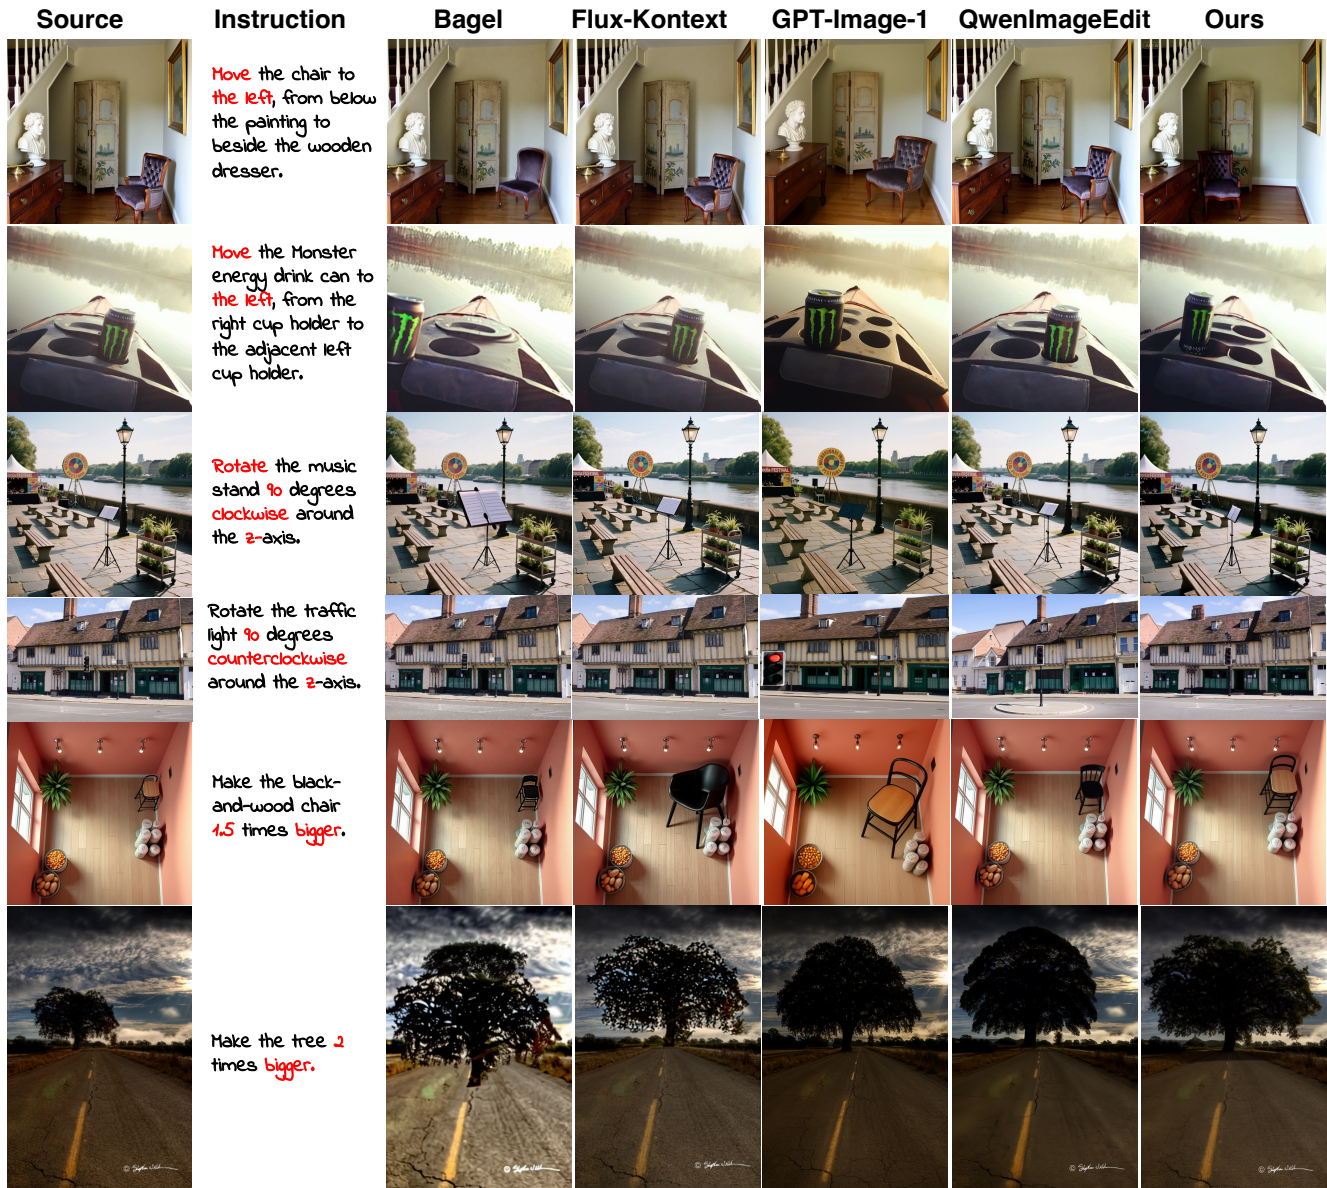


Figure 5. **Qualitative results** on object translation, rotation and resize over state-of-the-art image editing models. For each task, we provide one real image editing result (source from OpenImagesV6 [14]) and one synthetic image editing result to showcase the generalization ability of TALK2MOVE.

error of 0.2861 and a much higher accuracy of 29.55%, indicating that it provides a more reliable supervision signal for GRPO-based spatial manipulation.

7. Conclusion and Discussion

In this work, we presented TALK2MOVE, the first reinforcement learning framework for text-guided geometric object transformation in scene editing. Built on GRPO-style training, TALK2MOVE enables data-efficient learning with reduced reliance on expensive paired data and uses

spatially grounded rewards to disentangle objects from the scene and improve fine-grained transformations. We further enhance computational efficiency through off-policy step evaluation and active step sampling, prioritizing informative steps with higher reward variance and larger transformation space. Experiments show that TALK2MOVE outperforms existing open- and closed-source models in spatial accuracy and visual coherence. Our RL-based recipe can be extended to other generative frameworks (e.g., GANs, autoregressive models) for verifiable, controllable visual generation, which we leave for future work.

References

- [1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. [7](#)
- [2] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. *arXiv preprint arXiv:2305.13301*, 2023. [3](#)
- [3] Yen-Chi Cheng, Krishna Kumar Singh, Jae Shin Yoon, Alexander Schwing, Liang-Yan Gui, Matheus Gadelha, Paul Guerrero, and Nanxuan Zhao. 3d-fixup: Advancing photo editing with 3d priors. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–10, 2025. [2](#)
- [4] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. [2](#), [6](#), [7](#)
- [5] Abdelrahman Eldesokey and Peter Wonka. Build-a-scene: Interactive 3d layout control for diffusion-based image generation. *arXiv preprint arXiv:2408.14819*, 2024. [2](#)
- [6] Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, Tianshuo Peng, Junfei Wu, Xiaoying Zhang, Benyou Wang, and Xiangyu Yue. Video-r1: Reinforcing video reasoning in mllms. *arXiv preprint arXiv:2503.21776*, 2025. [3](#)
- [7] Patrick W Gallagher, Shuai Tang, and Zhuowen Tu. What happened to my dog in that network: Unraveling top-down generators in convolutional neural networks. *arXiv preprint arXiv:1511.07125*, 2015. [2](#)
- [8] Yu Gao, Lixue Gong, Qiushan Guo, Xiaoxia Hou, Zhichao Lai, Fanshi Li, Liang Li, Xiaochen Lian, Chao Liao, Liyang Liu, et al. Seedream 3.0 technical report. *arXiv preprint arXiv:2504.11346*, 2025. [3](#)
- [9] Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. Seed-x: Multi-modal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*, 2024. [2](#)
- [10] Dhruva Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:52132–52152, 2023. [5](#)
- [11] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. [3](#)
- [12] Xiaoxuan He, Siming Fu, Yuke Zhao, Wanli Li, Jian Yang, Dacheng Yin, Fengyun Rao, and Bo Zhang. Tempflow-grpo: When timing matters for grpo in flow models. *arXiv preprint arXiv:2508.04324*, 2025. [3](#)
- [13] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026, 2023. [5](#)
- [14] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Hajja, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Shahab Kamali, Matteo Mallocci, Jordi Pont-Tuset, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://storage.googleapis.com/openimages/web/index.html>*, 2017. [6](#), [8](#)
- [15] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2023. [2](#)
- [16] 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](#), [3](#), [6](#), [7](#)
- [17] Junzhe Li, Yutao Cui, Tao Huang, Yinping Ma, Chun Fan, Miles Yang, and Zhao Zhong. Mixgrpo: Unlocking flow-based grpo efficiency with mixed ode-sde. *arXiv preprint arXiv:2507.21802*, 2025. [3](#), [5](#)
- [18] Tiancheng Li, Jinxiu Liu, Huajun Chen, and Qi Liu. Instruct4pix: Training diffusion for image editing by reinforcement learning. *arXiv preprint arXiv:2406.09973*, 2024. [3](#)
- [19] Weiqi Li, Shijie Zhao, Chong Mou, Xuhan Sheng, Zhenyu Zhang, Qian Wang, Junlin Li, Li Zhang, and Jian Zhang. Omnidrag: Enabling motion control for omnidirectional image-to-video generation. *arXiv preprint arXiv:2412.09623*, 2024. [2](#)
- [20] Yuming Li, Yikai Wang, Yuying Zhu, Zhongyu Zhao, Ming Lu, Qi She, and Shanghang Zhang. Branchgrpo: Stable and efficient grpo with structured branching in diffusion models. *arXiv preprint arXiv:2509.06040*, 2025. [3](#), [5](#)
- [21] Zongjian Li, Zheyuan Liu, Qihui Zhang, Bin Lin, Shenghai Yuan, Zhiyuan Yan, Yang Ye, Wangbo Yu, Yuwei Niu, and Li Yuan. Uniworld-v2: Reinforce image editing with diffusion negative-aware finetuning and mllm implicit feedback. *arXiv preprint arXiv:2510.16888*, 2025. [3](#)
- [22] Weixin Liang, Lili Yu, Liang Luo, Srinivasan Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen-tau Yih, Luke Zettlemoyer, et al. Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *arXiv preprint arXiv:2411.04996*, 2024. [2](#)
- [23] Haofeng Liu, Chenshu Xu, Yifei Yang, Lihua Zeng, and Shengfeng He. Drag your noise: Interactive point-based editing via diffusion semantic propagation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6743–6752, 2024. [2](#)
- [24] Jie Liu, Gongye Liu, Jiajun Liang, Yangguang Li, Jiaheng Liu, Xintao Wang, Pengfei Wan, Di Zhang, and Wanli Ouyang.

- Flow-grpo: Training flow matching models via online rl. *arXiv preprint arXiv:2505.05470*, 2025. 3, 4, 5
- [25] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, pages 38–55. Springer, 2024. 5
- [26] Jingyi Lu, Xinghui Li, and Kai Han. Regiondrag: Fast region-based image editing with diffusion models. In *European Conference on Computer Vision*, pages 231–246. Springer, 2024. 2
- [27] Xin Luo, Jiahao Wang, Chenyuan Wu, Shitao Xiao, Xiyan Jiang, Defu Lian, Jiajun Zhang, Dong Liu, et al. Editscore: Unlocking online rl for image editing via high-fidelity reward modeling. *arXiv preprint arXiv:2509.23909*, 2025. 3, 5
- [28] Léopold Maillard, Tom Durand, Adrien Ramanana Rahary, and Maks Ovsjanikov. Laconic: A 3d layout adapter for controllable image creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18046–18057, 2025. 2
- [29] Xiaoyue Mi, Wenqing Yu, Jiesong Lian, Shibo Jie, Ruizhe Zhong, Zijun Liu, Guozhen Zhang, Zixiang Zhou, Zhiyong Xu, Yuan Zhou, et al. Video generation models are good latent reward models. *arXiv preprint arXiv:2511.21541*, 2025. 3
- [30] Yunhong Min, Daehyeon Choi, Kyeongmin Yeo, Jihyun Lee, and Minhyuk Sung. Origen: Zero-shot 3d orientation grounding in text-to-image generation. *arXiv preprint arXiv:2503.22194*, 2025. 2
- [31] Chong Mou, Xintao Wang, Jiechong Song, Ying Shan, and Jian Zhang. Dragondiffusion: Enabling drag-style manipulation on diffusion models. *arXiv preprint arXiv:2307.02421*, 2023. 2
- [32] OpenAI. Gpt-image-1: Openai’s image generation model. <https://platform.openai.com/docs/models/gpt-image-1>, 2024. Accessed: 2025-11-11. 6, 7
- [33] Xingang Pan, Ayush Tewari, Thomas Leimkühler, Lingjie Liu, Abhimitra Meka, and Christian Theobalt. Drag your gan: Interactive point-based manipulation on the generative image manifold. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–11, 2023. 2
- [34] Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Jiu hai Chen, Kunpeng Li, Felix Juefei-Xu, et al. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*, 2025. 2
- [35] Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pages 745–750, 2007. 3
- [36] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021. 2
- [37] Scott E Reed, Yi Zhang, Yuting Zhang, and Honglak Lee. Deep visual analogy-making. *Advances in neural information processing systems*, 28, 2015. 2
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2
- [39] Yujun Shi, Chuhui Xue, Jun Hao Liew, Jiachun Pan, Hanshu Yan, Wenqing Zhang, Vincent YF Tan, and Song Bai. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8839–8849, 2024. 2
- [40] Joonghyuk Shin, Daehyeon Choi, and Jaesik Park. Instantdrag: Improving interactivity in drag-based image editing. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–10, 2024. 2
- [41] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiyang Yu, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14398–14409, 2024. 2
- [42] Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025. 4
- [43] Zehan Wang, Ziang Zhang, Tianyu Pang, Chao Du, Hengshuang Zhao, and Zhou Zhao. Orient anything: Learning robust object orientation estimation from rendering 3d models. *arXiv preprint arXiv:2412.18605*, 2024. 5, 7
- [44] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, Yuxiang Chen, Zecheng Tang, Zekai Zhang, Zhengyi Wang, An Yang, Bowen Yu, Chen Cheng, Dayiheng Liu, Deqing Li, Hang Zhang, Hao Meng, Hu Wei, Jingyuan Ni, Kai Chen, Kuan Cao, Liang Peng, Lin Qu, Minggang Wu, Peng Wang, Shuting Yu, Tingkun Wen, Wensen Feng, Xiaoxiao Xu, Yi Wang, Yichang Zhang, Yongqiang Zhu, Yujia Wu, Yuxuan Cai, and Zenan Liu. Qwen-image technical report, 2025. 2, 3, 5, 6, 7
- [45] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:15903–15935, 2023. 3
- [46] Zeyue Xue, Jie Wu, Yu Gao, Fangyuan Kong, Lingting Zhu, Mengzhao Chen, Zhiheng Liu, Wei Liu, Qiushan Guo, Weilin Huang, et al. Dancegrpo: Unleashing grpo on visual generation. *arXiv preprint arXiv:2505.07818*, 2025. 3, 4, 5
- [47] Jiraphon Yenphraphai, Xichen Pan, Sainan Liu, Daniele Panofzo, and Saining Xie. Image sculpting: Precise object editing with 3d geometry control. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4241–4251, 2024. 2
- [48] Qihang Zhang, Yinghao Xu, Chaoyang Wang, Hsin-Ying Lee, Gordon Wetzstein, Bolei Zhou, and Ceyuan Yang. 3ditscene: Editing any scene via language-guided disentangled gaussian splatting. *arXiv preprint arXiv:2405.18424*, 2024. 2

- [49] Yujie Zhou, Pengyang Ling, Jiazi Bu, Yibin Wang, Yuhang Zang, Jiaqi Wang, Li Niu, and Guangtao Zhai. G2rpo: Granular grpo for precise reward in flow models. *arXiv preprint arXiv:2510.01982*, 2025. [3](#)