

OSOR: One-Step Diffusion Inpainting for Effect-Aware Object Removal

Qinming Zhou^{1,2*}, Chenxi Sun^{1,3*}, Deyang Kong^{1,3}, Junhao He¹, Xiangheng Tang^{1,4}, Peike Yu^{1,5}, Haotian Wu¹, Leilei Cao^{6†}, and Linfeng Zhang^{1‡}

¹ Shanghai Jiao Tong University, Shanghai, China
zhanglinfeng@sjtu.edu.cn

² Tsinghua University, Beijing, China

³ University of Electronic Science and Technology of China, Chengdu, China

⁴ Xidian University, Xi'an, China

⁵ Tongji University, Shanghai, China

⁶ Transsion, Shenzhen, China

*Equal contribution. †Project Leader. ‡Corresponding Author.




Fig. 1: Comparison between OSOR and other methods. OSOR effectively removes object-associated effects (e.g., shadows) while running $10.6\times$ faster than ObjectClear, the second-best method. A 1024×1024 image can be processed in under one second on a single A100 GPU. The ranking is averaged across six benchmarks.

Abstract. Real-world object removal is challenging due to two key difficulties: the target object’s non-local effects (e.g., shadows and reflections) that are hard to model, and the fact that user-provided masks are often inaccurate or incomplete. With billions of parameters and tens of denoising steps, diffusion-based models achieve this goal at the expense of massive computational cost, limiting their use in interactive applications and edge devices. To solve this problem, we present **OSOR** (One-Step Object Removal), which achieves efficient, effect-aware, and mask-robust object removal at the same time. Concretely, OSOR introduces (1) an occupancy-guided discriminator for precise boundary supervision, enabling stable single-step diffusion training, (2) an alpha head, which leverages the knowledge in diffusion models to predict correct removal regions with minimal overhead, handling imperfect masks, and (3) a semantic-anchored verification pipeline (SAVP) that filters noisy instruction-based triplets to produce effect-aware supervision at scale. Using SAVP, we curate **CORNE** with 280K verified removal pairs and further annotate AnimeEraseBench and TextEraseBench to evaluate model performance in more complex tasks. Experiments show that OSOR surpasses strong multi-step diffusion baselines in perceptual quality while achieving $4\times$ to $30\times$ faster inference. Code and resources are available at <https://github.com/Zhouqm-Git/osor>.

Keywords: Object removal · Image inpainting · Efficient diffusion

1 Introduction

Object removal is a fundamental image editing task that aims to eliminate a target object and its visual effects from a photograph, restoring a natural background in the affected region. Generally speaking, real-world object removal presents three major challenges as follows.

- (I) **Effect-awareness.** Removing an object is not simply a matter of erasing the pixels within a specified mask. The target object often leaves behind persistent visual effects in the scene, including cast shadows, reflected appearances, and other environmental interactions. Eliminating these effects requires a strong semantic understanding of the scene geometry and lighting, as the model must infer what the background should look like without the object and its influence.
- (II) **Mask-robustness.** In practical usage scenarios, the removal mask is typically provided by the user through interactive selection tools. These user-generated masks are frequently imperfect: they may be too large and cover unrelated content, or too small and fail to encompass the full extent of the object and its effects. This masks robustness issue is critical for deployment in real applications where users expect reasonable results even with careless input.
- (III) **Efficiency.** Object removal is often performed on mobile devices in an interactive manner, introducing requirements on extremely low latency.

Early approaches to object removal predominantly employed Generative Adversarial Networks (GANs) [7, 10, 25, 36]. These methods leveraged adversarial learning where ground-truth removal results served as real samples and model

outputs as fake samples. While GANs offered computational advantages through single forward passes in lightweight networks, their limited representational capacity constrained generation quality and prevented achieving satisfactory effect-aware and mask-robust results. Recent advances in diffusion models have fundamentally transformed the landscape of image generation and editing tasks. Modern diffusion models leverage hundreds of billions of parameters and dozens of iterative denoising steps to achieve unprecedented generative capabilities. These powerful models have been successfully adapted to image inpainting and object removal tasks [2, 8, 21, 26, 30], demonstrating superior perceptual quality compared to traditional GAN-based approaches. Despite these improvements, diffusion-based methods face a critical bottleneck in computational efficiency. The multi-step denoising process requires substantial computing resources and inference time, making it impractical for deployment on edge devices or interactive applications where latency is critical. Furthermore, even with the powerful priors encoded in pretrained diffusion models, existing approaches still struggle to fully address the challenges of effect-aware and mask-robust removal [6, 12, 35, 37, 42, 46]. This indicates that simply using pretrained diffusion models is insufficient, and targeted training strategies are necessary to unlock their full potential for this specific task.

To solve this problem, in this paper, we propose OSOR, a **One-Step** diffusion model for **Object Removal** that simultaneously achieves three goals: efficiency through single-step inference, effect-aware removal of shadows and reflections, and robustness to imperfect user masks. Our approach introduces three key technical contributions to address the aforementioned challenges.

Occupancy-guided discriminator for single-step diffusion. While adversarial learning based step distillation [33, 40] has been successfully utilized in image generation, we find that applying these methods directly to object removal leads to unsatisfactory results, particularly producing blurry boundaries around the removed object. This difficulty arises because single-step diffusion lacks the iterative refinement opportunity that allows multi-step methods to gradually correct errors in the generated content [2, 35, 46]. To address this, we design an occupancy-guided discriminator that provides precise boundary supervision by computing fractional occupancy values at multiple scales from the input mask for each patch location. Additionally, we propose formulating object removal as a latent restoration task derived from image restoration principles, where weakly noising the input image reduces training difficulty [17, 23].

A lightweight alpha head for imperfect removal mask correction. OSOR introduces an alpha head, which is implemented as a lightweight projection appended to the diffusion backbone. Leveraging the rich semantic knowledge already present in pretrained diffusion models [9, 16, 44], this component can accurately recover the appropriate mask for removal with minimal additional parameters and computational overhead. Besides, we further propose a two-stage training curriculum: the model first learns the removal task with perfect masks, then adapts to handle imperfect masks to improve robustness.

Semantic-anchored verification pipeline enabling effect-aware removal data at scale. Based on existing image editing datasets, this pipeline combines semantic information with pixel-space gradient analysis to verify successful removal and detect the presence of object effects. This allows us to automatically generate substantial amounts of labeled data for training effect-aware removal models. Using this pipeline, we construct CORNE, a high-quality dataset containing 280K verified object removal pairs with effect-aware annotations [42, 46]. Furthermore, to address the lack of comprehensive benchmarks for evaluating object removal in specific domains, we construct AnimeEraseBench and TextEraseBench, which evaluate removal capability on anime images and text overlays, respectively. In summary, our contributions are threefold.

- We propose OSOR, a single-step object removal model, which introduces an occupancy-guided discriminator, a lightweight alpha head, and a semantic-anchored verification pipeline (SAVP) to achieve efficient, effect-aware, and mask-robust object removal.
- Based on SAVP, we curate CORNE, a high-quality and effect-aware object removal training dataset with 280K removal pairs. Besides, we release AnimeEraseBench and TextEraseBench, two removal benchmarks for evaluation on removal in anime scenarios and text objects, respectively.
- Extensive experimental results on 6 benchmarks with 7 comparison methods demonstrate the superior removal quality and efficiency (Fig. 1). For instance, 2.24 dB higher PSNR and $27\times$ faster than the second-best method in AnimeEraseBench.

2 Related Work

2.1 Image Inpainting and Object Removal

Object removal is related to mask-conditioned inpainting, but it imposes a stricter requirement to preserve surrounding context while removing the target object and its visual effects. GAN-based inpainting enables fast inference but often struggles with boundary continuity and high-frequency realism under irregular or large masks [10, 25, 36, 41]. Transformer-based inpainting improves global structure modeling for large missing regions [5, 15]. Diffusion-based inpainting further boosts realism by leveraging strong generative priors [1, 2, 21, 26, 30]. However, general inpainting backbones are optimized for generic completion rather than removal behavior, so outputs remain sensitive to mask quality when object effects extend beyond the provided spatial conditioning.

2.2 Object Removal Models and Datasets

Recent works adapt diffusion models to object removal through stronger conditioning, task-oriented guidance, or effect-aware modeling [6, 12, 19, 35, 37, 42, 46]. Progress is also driven by improved supervision sources, including video-derived paired frames and captured counterfactual pairs that better reflect real

effects [31,37,38]. In parallel, instruction-based corpora provide abundant editing triplets but can be noisy, motivating verification and filtering to obtain reliable supervision [3,13]. A persistent difficulty in interactive removal is that user masks often under-cover soft effects. Several methods strengthen object effect coupling in conditioning or supervision, for example via object effect attention [46]. Another direction is to predict a soft effective editing region to represent uncertain boundaries and residual effects beyond the provided mask. This view is related to alpha compositing and matting, which model soft transitions with an opacity map instead of a hard mask [14,27,39]. Recent diffusion-based matting further supports predicting alpha-like maps from diffusion priors as a representation of soft boundaries and layer uncertainty [9,16,44].

2.3 Efficient Diffusion and One-Step Generation

Reducing diffusion inference cost has been studied through distillation and consistency training that reduces denoising steps [22,32,34]. A growing line of work targets one-step or near one-step generation, distilling pretrained diffusion models into single-pass generators via distribution matching and adversarial objectives [33,40]. Most efficient diffusion methods are developed for global generation and do not explicitly enforce the strict context preservation required by mask-conditioned editing. Single-pass editing is also more sensitive to boundary ambiguity near the mask, which motivates task-specific supervision when applying one-step inference to object removal.

3 Methodology

OSOR performs one-step object-and-effect removal by restoring a clean background from an intermediate noised latent. Training relies on effect-aware supervision and paired backgrounds, which we obtain with SAVP, a semantic-anchored verification pipeline over noisy instruction-based triplets, yielding the CORNE dataset with effect-aware masks. We train OSOR in two phases, as summarized in Fig. 3. Phase I focuses on boundary-consistent one-step restoration under well-localized masks using an occupancy-guided multi-scale discriminator with patch-level targets. Phase II adds a lightweight alpha head and incomplete-mask conditioning to improve robustness to conservative or misaligned user masks. Implementation details and verification thresholds are provided in the supplementary material.

3.1 SAVP and the CORNE Dataset

Effect-aware object removal requires paired backgrounds and masks that cover both the object and its visual effects such as cast shadows and reflections. We introduce SAVP, a semantic-anchored verification pipeline that extracts reliable removal supervision from noisy instruction-based triplets. We apply SAVP to the single-edit subset of NHR-Edit [13] with *add* or *remove* instructions. For




Fig. 2: Overview of SAVP and CORNE. Starting from single-edit instruction triplets, SAVP verifies semantically aligned and localized differences, then fuses the validated difference region with promptable segmentation to form an effect-aware mask. It further derives object-core masks for Phase II incomplete-mask conditioning.

add, we set $(I_{\text{shot}}, I_{\text{gt}}) = (I_{\text{edit}}, I_{\text{orig}})$, and for *remove* we swap the order. SAVP verifies that the resulting pair exhibits localized and semantically consistent differences (Fig. 2), and outputs paired backgrounds with effect-aware masks to form CORNE.

Semantic-anchored verification. Given a single-edit instruction triplet, SAVP forms an ordered image pair $(I_{\text{shot}}, I_{\text{gt}})$ and verifies that the visual difference is localized and semantically aligned with the instruction (Fig. 2). We compute a multi-feature difference heatmap from log-luminance, chromaticity, and gradient magnitude, then binarize and clean it to obtain a difference mask m_{diff} . Connected components in m_{diff} yield candidate boxes b_{diff} . We run GroundingDINO [18] with the instruction text on I_{shot} to obtain semantic boxes b_{sem} . We first apply a global rejection based on the fragmentation of dominant components and the noise ratio of small components to discard pairs with dispersed artifacts. We then traverse b_{diff} in descending area and match each candidate to its best-overlapping semantic box in b_{sem} . A candidate is accepted as b_{val} only if its best IoU exceeds a threshold and its area remains within a scale ratio bound, producing a refined mask $m_{\text{diff}}^{\text{val}}$. If a candidate violates the scale bound we discard the entire pair as a collapse case, otherwise we drop the candidate and continue. We keep a triplet only if at least one validated region exists. Implementation details and thresholds are deferred to the supplementary.

Effect-aware mask synthesis. The validated difference mask $m_{\text{diff}}^{\text{val}}$ localizes the edit but can be fragmented and may miss parts of the object. We therefore obtain an object-core mask m_{obj} with SAM2 [29] on I_{shot} using b_{val} as box prompts, and fuse it with the validated difference region (Fig. 2),

$$m_{\text{fuse}} = m_{\text{obj}} \cup m_{\text{diff}}^{\text{val}}. \quad (1)$$

We optionally apply a lightweight dilation to m_{fuse} to obtain the final effect-aware target mask m_{gt} used in training.

Effect decomposition for incomplete-mask conditioning. Phase II requires tight object-core conditioning masks to simulate conservative user inputs.




Fig. 3: Two-phase training curriculum. Phase I adapts a diffusion inpainting backbone with hard latent blending and occupancy-guided patch supervision for boundary-consistent one-step removal. Phase II predicts a soft alpha map under incomplete-mask conditioning and performs adaptive blending to remove residual shadows and reflections beyond the provided mask.

We define the effect residual on the pre-expansion fused mask,

$$m_{\text{eff}} = m_{\text{fuse}} \setminus m_{\text{obj}}. \quad (2)$$

We select effect-heavy cases by an effects ratio $\|m_{\text{eff}}\|_1 / \|m_{\text{fuse}}\|_1$. For effect-heavy cases, Phase II samples the conditioning mask m_{in} from a set of conservative masks that includes the tight object-core mask m_{obj} , while supervising with the effect-aware target m_{gt} .

3.2 One-step Latent Restoration

OSOR builds on diffusion-family inpainting backbones (SDXL-Inpainting [26] and FLUX Fill [2]) and operates in the latent space of a pretrained VAE with encoder $E(\cdot)$ and decoder $D(\cdot)$. Given an input *shot* image x and a user-provided mask m , we encode $\bar{z} = E(x)$ and apply forward noising at an intermediate noise level t [23],

$$z_t = \alpha_t \bar{z} + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \quad (3)$$

where α_t and σ_t are schedule coefficients. The backbone is conditioned on the full input latent, the mask, and a fixed text embedding. We use the constant prompt “Remove the instance of object”, denote its embedding by e , and form $c = \langle \bar{z}, m, e \rangle$. This exposes full scene context through $\bar{z}$ while using m only to localize the intended edit.

Given (z_t, c, t) , the backbone predicts its native denoising output

$$u_\theta = f_\theta(z_t, c, t), \quad (4)$$




Fig. 4: Mask-derived patch targets for a four-scale discriminator. Left shows the input mask and its overlay on the shot image for visualization. Right compares three target constructions at each scale. HM uses nearest-neighbor downsampling. SM applies Gaussian smoothing after downsampling. OG uses area pooling to produce fractional occupancies. Differences grow on coarser grids.

and we obtain a one-step estimate of the clean latent via the corresponding one-step mapping,

$$\hat{z}_0 = \frac{z_t - \sigma_t u_\theta}{\alpha_t}. \quad (5)$$

3.3 Phase I: Boundary-consistent One-step Removal

Phase I adapts a pretrained backbone for reliable one-step removal when the affected region is well localized. We use CORNE supervision and take the effect-aware mask m_{gt} as the target region. In Phase I we set the conditioning mask as $m = m_{\text{gt}}$. Following Sec. 3.2, we obtain z_t and the one-step prediction $\hat{z}_0$. To preserve non-edited content exactly, we perform hard blending in latent space (Fig. 3a),

$$z_{\text{out}} = m_z \odot \hat{z}_0 + (1 - m_z) \odot \bar{z}, \quad (6)$$

where $\bar{z} = E(x)$ and m_z is the backbone-specific mask representation used for latent blending. We decode $\hat{x} = D(z_{\text{out}})$. Hard blending keeps the unmasked region identical to the input and restricts gradients from training objectives to the edited region, which stabilizes one-step adaptation.

Occupancy-guided multi-scale discriminator. Single-step restoration can preserve global structure but often shows seams near mask boundaries, where each discriminator patch mixes preserved context and synthesized content (Fig. 4). We use a multi-scale patch discriminator in the PatchGAN family [11, 36, 43]. It consists of a frozen feature trunk ϕ and lightweight trainable patch heads $\{h_\xi^k\}_k$. In our implementation, ϕ is a pretrained OpenCLIP ConvNeXt [4, 20, 28] that outputs multi-resolution features $f_k = \phi_k(x)$, and each head predicts a score map $D_\xi^k(x) = \sigma(h_\xi^k(f_k))$. We derive mask-based supervision targets at four scales, where target discretization becomes more pronounced at coarser heads (Fig. 4).

Mapping a binary mask to a coarse logit grid is ambiguous at boundary patches. Nearest-neighbor downsampling produces hard labels and makes partially covered patches over-confident [36]. Gaussian smoothing yields soft targets




Fig. 5: Qualitative comparison of OSOR and existing methods on RORD-Val, RemovalBench and TextEraseBench.

but depends on the bandwidth σ , which is not tied to patch occupancy and becomes more consequential at larger downsampling factors [43]. We instead compute an occupancy map $\tilde{w}_k \in [0, 1]$ by area pooling the mask to each discriminator scale. Each value equals the masked-area fraction within the patch of a logit location, giving exact fractional targets for boundary patches.

Occupancy-guided objectives. We write the adversarial and reconstruction objectives using a spatial guidance map $w \in [0, 1]^{H \times W}$. In Phase I, we set $w = m_{\text{gt}}$, and $\tilde{w}_k$ is obtained by area pooling w to the k -th discriminator output resolution. Let x^{bg} denote the paired ground-truth background.

The discriminator objective is

$$\begin{aligned} \mathcal{L}_D(w) = & \sum_k \mathbb{E}[-\log D_\xi^k(x^{\text{bg}})] \\ & + \sum_k \mathbb{E}[-(1 - \tilde{w}_k) \odot \log D_\xi^k(\hat{x}) - \tilde{w}_k \odot \log(1 - D_\xi^k(\hat{x}))] + \lambda_{\text{r1}} \mathcal{R}_{\text{r1}}. \end{aligned} \quad (7)$$

We use an R1 regularizer on real samples [24] and compute it on the head inputs since ϕ is frozen (see supplementary).

For the generator, we adopt the non-saturating form and normalize by the occupied area so the loss scale is insensitive to mask size,

$$\mathcal{L}_{\text{adv}}(w) = \sum_k \mathbb{E} \left[\frac{\sum_p \tilde{w}_k(p) (-\log D_\xi^k(\hat{x})_p)}{\sum_p \tilde{w}_k(p) + \varepsilon} \right]. \quad (8)$$

We also use a mask-normalized reconstruction term,

$$\mathcal{L}_{\text{rec}}(w) = \mathbb{E} \left[\frac{\|w \odot (\hat{x} - x^{\text{bg}})\|_1}{\|w\|_1 + \varepsilon} \right], \quad (9)$$




Fig. 6: Qualitative comparison of OSOR and existing methods on CORNE-Val and AnimeEraseBench.

together with $\mathcal{L}_{\text{per}} = \text{LPIPS}(\hat{x}, x^{\text{bg}})$ [45]. The generator objective is

$$\mathcal{L}_G(w) = \lambda_{\text{rec}} \mathcal{L}_{\text{rec}}(w) + \lambda_{\text{per}} \mathcal{L}_{\text{per}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}(w). \quad (10)$$

For parameter-efficient adaptation of large pretrained backbones, we update only lightweight adapters and the terminal output projection while keeping the remaining pretrained weights frozen. Phase I solves the adversarial game with $w = m_{\text{gt}}$ using

$$\min_{\theta} \max_{\xi} \mathcal{L}_G(m_{\text{gt}}) - \mathcal{L}_D(m_{\text{gt}}), \quad (11)$$

which is optimized by alternating updates of θ and ξ .

3.4 Phase II: Alpha-aware Robust Removal with Adaptive Blending

Phase I assumes the conditioning mask covers both the object and its effects. In practice, user masks are often conservative or misaligned and frequently miss soft shadows and reflections. Phase II therefore trains OSOR with an *incomplete* conditioning mask m_{in} and predicts a soft alpha map for adaptive blending (Fig. 3b).

Alpha prediction and adaptive blending. We augment the generator with a lightweight alpha head so that a single forward pass produces both the denoising prediction and an alpha logit map [9, 16, 44]. Concretely, we extend the terminal output projection of the backbone to emit an additional set of logits ℓ_{θ} alongside its native denoising output. For SDXL-Inpainting, we expand the final convolutional output layer of the U-Net; for FLUX Fill, we expand the final output projection of the transformer. This design reuses all backbone computation and adds a small overhead. Given z_t , c , and t , the generator predicts

$$(u_{\theta}, \ell_{\theta}) = f_{\theta}(z_t, c, t), \quad \hat{\alpha} = \sigma(\ell_{\theta}), \quad (12)$$

where $\hat{\alpha} \in [0, 1]$ is predicted at the latent resolution and estimates the effective editing region. For clarity, Fig. 3b visualizes $\hat{\alpha}$ as the alpha output of the head.

We recover $\hat{z}_0$ as in Sec. 3.2 and replace hard blending with alpha compositing in latent space [14, 27, 39],

$$z_{\text{out}} = \hat{\alpha} \odot \hat{z}_0 + (1 - \hat{\alpha}) \odot \bar{z}, \quad (13)$$

with $\bar{z} = E(x)$. Backbone-specific output parameterizations and the exact projection modifications are provided in the supplementary.

Incomplete-mask conditioning.

In Phase II, we replace the conditioning mask in c by an incomplete mask m_{in} , using $c = \langle \bar{z}, m_{\text{in}}, e \rangle$. We sample m_{in} from a family of conservative masks derived from m_{gt} . This family includes tight object-core masks m_{obj} obtained via promptable segmentation, as well as simple perturbations of m_{gt} such as dilation, erosion, translation, and random hole dropping, as illustrated in Fig. 7. These conditioning masks intentionally under-




Fig. 7: Examples of incomplete conditioning masks m_{in} generated from object-core segmentation and simple geometric perturbations.

cover the true affected region, so the model must infer missing effects from the image.

Alpha-guided training. A key design choice is that adversarial and reconstruction losses are always evaluated on the *effect-aware target region* m_{gt} , rather than on the predicted $\hat{\alpha}$, to avoid degenerate solutions where the model reduces the loss by shrinking $\hat{\alpha}$. Accordingly, we reuse the unified objectives in Eqs. (7)–(10) with $w = m_{\text{gt}}$ in Phase II as well. The predicted alpha is used only in the latent compositing of Eq. (13) and is explicitly supervised to match the effect-aware extent. Let m_{gt}^z denote m_{gt} downsampled to the latent resolution (as illustrated in Fig. 3b). We add

$$\mathcal{L}_{\alpha} = \lambda_{\text{bce}} \text{BCE}(\ell_{\theta}, m_{\text{gt}}^z) + \lambda_{\text{dice}} \text{Dice}(\hat{\alpha}, m_{\text{gt}}^z). \quad (14)$$

This trains $\hat{\alpha}$ to recover the full effect-aware extent under incomplete conditioning, enabling removal of shadows and reflections beyond the user input. The overall Phase II objective is

$$\min_{\theta} \max_{\xi} \mathcal{L}_G(m_{\text{gt}}) + \mathcal{L}_{\alpha} - \mathcal{L}_D(m_{\text{gt}}). \quad (15)$$

We initialize Phase II from the Phase I weights and maintain the same parameter-efficient adaptation strategy. LoRA updates the main blocks, while the terminal projection that emits $(u_{\theta}, \ell_{\theta})$ is fine-tuned to calibrate one-step output statistics.

4 Experiments

We defer full implementation details, benchmark specifications, and metric definitions to the supplementary material. Unless stated otherwise, all methods

Table 1: Quantitative comparison on paired-background benchmarks CORNE-Val, RORD-Val, and AnimeEraseBench under object-only masks and effect-aware masks.

Dataset	Method	Latency (s)	Object Mask						Object-Effect Mask					
			FID↓	CMMD↓	LPIPS↓	PSNR↑	SSIM↑	CFD↓	FID↓	CMMD↓	LPIPS↓	PSNR↑	SSIM↑	CFD↓
CORNE-Val	OmniPaint	26.12	<u>18.6077</u>	<u>0.0093</u>	<u>0.0563</u>	28.436	0.9166	0.1950	<u>19.111</u>	<u>0.0094</u>	<u>0.0603</u>	27.950	0.9119	0.2806
	ObjectClear	6.44	22.0457	0.0111	0.0956	<u>29.207</u>	<u>0.9191</u>	0.2040	22.3819	0.0113	0.0976	<u>29.012</u>	<u>0.9166</u>	0.3352
	AttentiveEraser	8.69	25.4818	0.0193	0.0790	25.650	0.8636	0.2244	26.3795	0.0176	0.0825	25.602	0.8584	0.3867
	FLUX-Fill	25.27	59.5678	0.0285	0.0824	23.124	0.9130	0.2766	53.2040	0.0260	0.0784	23.629	0.9119	0.7265
	OmniEraser	20.53	60.3871	0.1079	0.1801	21.540	0.7623	0.2091	59.3292	0.1049	0.1800	21.720	0.7604	0.3250
	SDXL-INP	6.01	75.1054	0.1506	0.2338	16.203	0.7149	0.2883	63.1940	0.1357	0.2276	16.804	0.7177	0.4665
	CLIPAway	5.96	68.1043	0.1740	0.1864	21.446	0.7641	0.2746	51.2921	0.1587	0.1781	22.850	0.7676	0.4661
	OSOR(SDXL)	0.42	29.3846	0.0249	0.1272	27.262	0.8371	<u>0.1608</u>	29.2953	0.0255	0.1268	27.422	0.8376	0.2300
	OSOR(FLUX)	0.80	12.3241	0.0073	0.0462	32.031	0.9373	0.1538	12.5203	0.0069	0.0460	32.189	0.9377	<u>0.2339</u>
RORD-Val	OmniPaint	16.51	27.5658	0.0668	0.1110	23.293	0.7899	0.3707	28.1457	0.0675	0.1147	23.078	0.7848	0.3754
	ObjectClear	8.69	25.4667	0.1458	0.1206	25.804	0.8500	0.3248	26.3523	0.1456	0.1234	25.633	0.8462	0.3323
	AttentiveEraser	8.68	36.5740	0.1951	0.1685	22.798	0.7023	0.3608	34.9941	0.1822	0.1688	23.080	0.6995	0.3558
	FLUX-Fill	24.80	118.0815	0.5883	0.1821	20.171	0.8156	0.5415	111.9136	0.5451	0.1818	20.587	0.8127	0.5607
	OmniEraser	25.99	43.2935	0.2409	0.2174	21.900	0.6550	0.3788	43.4599	0.2395	0.2183	21.881	0.6537	0.3788
	SDXL-INP	6.38	95.2105	0.7893	0.3247	17.155	0.6466	0.5551	88.0085	0.7111	0.3234	17.398	0.6469	0.5788
	CLIPAway	4.95	89.5176	0.9300	0.2962	19.447	0.5851	0.5805	81.5943	0.9031	0.2952	19.665	0.5850	0.5662
	OSOR(SDXL)	0.42	35.5989	0.1854	0.1795	23.767	0.6827	<u>0.2644</u>	35.6523	0.1858	0.1797	23.826	0.6823	<u>0.2624</u>
	OSOR(FLUX)	0.62	<u>27.3621</u>	<u>0.0952</u>	<u>0.1140</u>	<u>25.599</u>	<u>0.8175</u>	0.2465	<u>28.1271</u>	<u>0.1016</u>	<u>0.1154</u>	<u>25.613</u>	<u>0.8163</u>	0.2324
AnimeEraseBench	OmniPaint	24.93	<u>27.4034</u>	<u>0.0231</u>	<u>0.1129</u>	<u>25.542</u>	0.8573	0.4369	<u>26.8119</u>	<u>0.0223</u>	0.1172	25.351	0.8515	0.4922
	ObjectClear	10.18	36.6865	0.0241	0.1788	25.429	0.8409	0.3872	36.4259	0.0228	0.1800	<u>25.530</u>	0.8381	0.4406
	AttentiveEraser	8.69	37.5448	0.0341	0.1333	22.587	0.7314	0.3802	37.7600	0.0337	0.1375	22.649	0.7259	0.4625
	FLUX-Fill	25.19	76.5397	0.0558	0.1269	22.668	<u>0.8597</u>	0.4419	59.5985	0.0458	<u>0.1115</u>	23.960	<u>0.8625</u>	0.7075
	OmniEraser	19.78	57.5116	0.1304	0.2339	19.190	0.6284	0.4249	56.6766	0.1220	0.2350	19.531	0.6232	0.4839
	SDXL-INP	6.21	89.4030	0.1280	0.2754	14.827	0.6516	0.4952	70.5512	0.1081	0.2719	14.991	0.6516	0.5614
	CLIPAway	5.96	88.9794	0.2142	0.3191	19.240	0.5748	0.4732	79.8083	0.2064	0.3143	19.581	0.5747	0.5293
	OSOR(SDXL)	0.48	55.7614	0.0666	0.2367	23.090	0.6857	<u>0.2820</u>	55.4380	0.0658	0.2368	23.130	0.6858	<u>0.3713</u>
	OSOR(FLUX)	0.89	26.8352	0.0165	0.1070	27.780	0.8859	0.2703	26.2539	0.0168	0.1077	27.871	0.8855	0.3505

follow the same mask protocol and paired-background evaluation described there. We compare with SDXL-Inpainting [26] and FLUX Fill [2] and with OmniEraser [37], CLIPAway [6], AttentiveEraser [35], ObjectClear [46], and OmniPaint [42] using official code and recommended settings, and we measure latency on an NVIDIA A100.

4.1 Comparison with Previous Methods

Quantitative Results. Table 1 reports paired-background results under object-only masks and effect-aware masks. OSOR runs in under one second per image. On CORNE-Val, OSOR-FLUX achieves the best scores across all reported metrics under both mask settings, while OSOR-SDXL gives the lowest CFD under the effect-aware mask. On RORD-Val, OSOR-FLUX attains the lowest CFD under both masks and remains competitive on perceptual metrics. On AnimeEraseBench, OSOR-FLUX again ranks first across fidelity and perceptual measures under both masks. The results change little when switching between the two mask settings, which matches the goal of Phase II training. Table 2 reports benchmarks evaluated under object-only masks. OSOR-FLUX performs best on TextEraseBench and OmniPaint-Bench for FID, CMMD, LPIPS, PSNR, and SSIM. On RemovalBench, OSOR-SDXL achieves the lowest CFD, and OSOR-FLUX remains competitive across the remaining metrics while retaining sub-second latency.

Table 2: Quantitative comparison under object-only masks on TextEraseBench, OmniPaint-Bench, and RemovalBench.

Dataset	Method	Latency (s)	FID↓	CMMD↓	LPIPS↓	PSNR↑	SSIM↑	CFD↓
TextEraseBench	OmniPaint	24.89	29.3857	0.0348	<u>0.1039</u>	28.404	0.8811	0.2296
	ObjectClear	10.16	<u>29.1359</u>	<u>0.0293</u>	0.1361	<u>29.689</u>	<u>0.8945</u>	0.1957
	AttentiveEraser	8.69	34.6198	0.0397	0.1124	26.540	0.7974	<u>0.1994</u>
	FLUX-Fill	25.27	85.1376	0.3464	0.1250	22.157	0.8781	0.4745
	OmniEraser	19.80	53.7494	0.1129	0.2071	21.300	0.7134	0.2217
	SDXL-INP	6.22	68.4836	0.1823	0.2383	19.300	0.7284	0.3422
	CLIPAway	6.22	64.7926	0.3275	0.2440	23.106	0.7018	0.3011
	OSOR(SDXL)	0.43	38.1242	0.0418	0.1800	27.748	0.7735	0.2128
	OSOR(FLUX)	0.83	19.6169	0.0225	0.0803	31.997	0.9161	0.2066
OmniPaint-Bench	OmniPaint	34.13	<u>50.0921</u>	<u>0.0808</u>	<u>0.1522</u>	23.455	<u>0.8126</u>	0.2700
	ObjectClear	8.90	58.3634	0.1353	0.2357	<u>23.543</u>	0.7991	0.2535
	AttentiveEraser	8.68	62.1880	0.1330	0.1732	22.406	0.7762	0.3290
	FLUX-Fill	25.27	124.7098	0.3085	0.2056	19.460	0.7902	0.4750
	OmniEraser	14.39	75.0069	0.1743	0.2502	22.063	0.7062	0.3194
	SDXL-INP	6.16	115.0825	0.3088	0.2828	18.081	0.7059	0.5792
	CLIPAway	5.96	117.2142	0.5298	0.3005	19.434	0.6972	0.5457
	OSOR(SDXL)	0.42	67.4797	0.1967	0.2648	23.168	0.7419	0.2104
	OSOR(FLUX)	1.06	49.1927	0.0703	0.1419	24.936	0.8330	<u>0.2303</u>
RemovalBench	OmniPaint	30.59	41.1790	0.0273	0.2135	24.599	0.7488	0.2180
	ObjectClear	10.16	49.2744	0.0335	0.2871	25.884	0.7726	<u>0.1810</u>
	AttentiveEraser	8.69	80.3447	0.1788	0.1942	23.787	0.7002	0.2094
	FLUX-Fill	25.27	178.6896	0.4454	0.2695	20.277	0.7099	0.3984
	OmniEraser	14.39	65.4691	0.3066	0.2929	23.030	0.6712	0.1835
	SDXL-INP	6.03	165.4925	0.4756	0.3388	16.983	0.6251	0.3227
	CLIPAway	5.96	169.0450	0.4072	0.3787	20.461	0.6537	0.3569
	OSOR(SDXL)	0.43	63.6405	0.0709	0.3313	24.676	0.6943	0.1585
	OSOR(FLUX)	0.93	<u>43.8899</u>	<u>0.0300</u>	<u>0.1972</u>	<u>25.803</u>	<u>0.7612</u>	<u>0.1810</u>

Table 3: Noise-level ablation on RORD-Val with object-effect masks.

Timestep	FID↓	CMMD↓	LPIPS↓	PSNR↑	SSIM↑	CFD↓
200	60.1520	0.2463	0.1926	22.232	0.6718	0.3610
400	50.0639	0.2143	0.1855	22.696	0.6741	0.3013
600	55.4021	0.2263	0.1884	22.563	0.6720	0.3418
800	51.8961	0.2505	0.1924	22.644	0.6710	0.3299

Table 4: Patch-target ablation on RORD-Val with object-effect masks.

Target	PSNR↑	SSIM↑	FID↓	LPIPS↓	CFD↓
Hard Mask	22.463	0.6739	56.1508	0.2007	0.3399
Gaussian Soft	21.992	0.6708	64.1290	0.2103	0.3430
Occupancy	22.696	0.6741	50.0639	0.1855	0.3013

Qualitative Results. Fig. 5 and Fig. 6 highlight common failure modes in existing methods. CLIPAway often hallucinates new content inside the masked region. AttentiveEraser removes the object but leaves cast shadows and reflections beyond an object-only mask. OmniEraser, ObjectClear, and OmniPaint are generally strong on object removal and often suppress associated effects, yet occasional cases still exhibit mild residues, boundary inconsistencies, or unintended content. OSOR more consistently removes both the object and its associated effects while preserving cleaner background structure and boundaries across the shown cases.



Fig. 8: Alpha compositing under imperfect masks.

Table 5: Ablation of alpha compositing on RORD-Val under object-only and effect-aware conditioning masks.

Conditioning mask	Method	PSNR $\uparrow$	SSIM $\uparrow$	FID $\downarrow$	CMMD $\downarrow$	LPIPS $\downarrow$	CFD $\downarrow$
Object-only m_{obj}	Phase I	21.965	0.6726	59.8017	0.2764	0.2062	0.3133
	Phase II	23.767	0.6827	35.5989	0.1854	0.1795	0.2644
Effect-aware m_{gt}	Phase I	22.696	0.6741	50.0639	0.2143	0.1855	0.3013
	Phase II	23.826	0.6823	35.6523	0.1858	0.1797	0.2624

4.2 Ablation Study

Noise level for one-step denoising. We study the noise level t for one-step restoration on SDXL-Inpainting and evaluate on RORD-Val with object-effect masks. We sweep $t \in \{200, 400, 600, 800\}$ while keeping all other settings fixed, including the occupancy-guided multi-scale discriminator. Table 3 shows that $t = 400$ yields the best overall trade-off across fidelity and perceptual metrics. We use $t = 400$ in all subsequent experiments unless stated otherwise.

Occupancy-guided multi-scale discriminator. We ablate the construction of mask-derived patch targets on SDXL-Inpainting and evaluate on RORD-Val with object-effect masks. We fix $t = 400$ and compare hard targets from nearest-neighbor downsampling, Gaussian-smoothed targets, and our occupancy targets from area pooling. Table 4 shows that occupancy targets improve both perceptual and full-reference metrics, supporting fractional supervision at boundary patches.

Alpha compositing with conservative masks. We evaluate whether the alpha head is needed for robust removal when the conditioning mask under-covers object effects. We compare Phase I with hard latent blending in Eq. (6) and Phase II with alpha compositing in Eq. (13) on RORD-Val under two conditioning masks. The first uses the effect-aware mask m_{gt} . The second uses an object-only mask m_{obj} that excludes effect regions and serves as a conservative input. Table 5 shows that Phase II improves both fidelity and perceptual quality under both conditioning settings, with a clearer advantage when the conditioning mask is object-only. Figure 8 visualizes the mechanism. The predicted alpha extends into effect regions that are excluded from m_{obj} , enabling the model to modify a broader effective region than hard blending.

5 Conclusion

We presented OSOR, a one-step diffusion inpainting framework for effect-aware object removal. OSOR formulates removal as latent restoration from an intermediate noised latent and predicts the clean background in a single denoising pass. To improve boundary consistency in single-step training, we introduce an

occupancy-guided multi-scale discriminator that uses fractional mask occupancies as patch-level targets. To handle conservative or misaligned user masks, we add a lightweight alpha head and train with incomplete-mask conditioning so the model can remove effects beyond the provided boundary. We further propose SAVP to extract effect-aware supervision from noisy instruction-based triplets and curate CORNE with 280K verified removal pairs, together with AnimeEraseBench and TextEraseBench for evaluation. Experiments show that OSOR reaches strong perceptual quality while reducing inference time by $4\times$ to $30\times$, and it processes 1024×1024 images within one second on a single A100 GPU.

Acknowledgements

This work was supported by Transsion Holdings.

References

1. Avrahami, O., Lischinski, D., Fried, O.: Blended diffusion for text-driven editing of natural images. In: CVPR. pp. 18187–18197 (2022)
2. Black Forest Labs: FLUX.1 Fill [dev]. <https://huggingface.co/black-forest-labs/FLUX.1-Fill-dev> (2024)
3. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR. pp. 18392–18402 (2023)
4. Cherti, M., Beaumont, R., Wightman, R., Wortsman, M., Ilharco, G., Gordon, C., Schuhmann, C., Schmidt, L., Jitsev, J.: Reproducible scaling laws for contrastive language-image learning. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17–24, 2023. pp. 2818–2829 (2023). <https://doi.org/10.1109/CVPR52729.2023.00276>
5. Dong, Q., Cao, C., Fu, Y.: Incremental transformer structure enhanced image inpainting with masking positional encoding. In: CVPR. pp. 11348–11358 (2022)
6. Ekin, Y., Yildirim, A.B., Caglar, E.E., Erdem, A., Erdem, E., Dundar, A.: CLIP-Away: Harmonizing focused embeddings for removing objects via diffusion models. In: Advances in Neural Information Processing Systems. vol. 37, pp. 17572–17601 (2024). <https://doi.org/10.52202/079017-0559>
7. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: NeurIPS. pp. 2672–2680 (2014)
8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS. pp. 6840–6851 (2020)
9. Hu, X., Peng, X., Luo, D., Ji, X., Peng, J., Jiang, Z., Zhang, J., Jin, T., Wang, C., Ji, R.: Diffumattng: Synthesizing arbitrary objects with matting-level annotation. In: Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXVIII. vol. 15126, pp. 396–413 (2024). https://doi.org/10.1007/978-3-031-73113-6_23
10. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Globally and locally consistent image completion. ACM TOG **36**(4), 107:1–107:14 (2017)

11. Isola, P., Zhu, J., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 5967–5976 (2017). <https://doi.org/10.1109/CVPR.2017.632>
12. Jiang, L., Wang, Z., Bao, J., Zhou, W., Chen, D., Shi, L., Chen, D., Li, H.: Smarteraser: Remove anything from images using masked-region guidance. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 24452–24462 (2025). <https://doi.org/10.1109/CVPR52734.2025.02277>
13. Kuprashevich, M., Alekseenko, G., Tolstykh, I., Fedorov, G., Suleimanov, B., Dokholyan, V., Gordeev, A.: NoHumansRequired: Autonomous high-quality image editing triplet mining. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 6059–6068 (March 2026)
14. Levin, A., Lischinski, D., Weiss, Y.: A closed-form solution to natural image matting. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(2), 228–242 (2008). <https://doi.org/10.1109/TPAMI.2007.1177>
15. Li, W., Lin, Z., Zhou, K., Qi, L., Wang, Y., Jia, J.: Mat: Mask-aware transformer for large hole image inpainting. In: CVPR. pp. 10748–10758 (2022)
16. Li, X., Yang, Z., Quan, R., Yang, Y.: DRIP: Unleashing diffusion priors for joint foreground and alpha prediction in image matting. In: Advances in Neural Information Processing Systems. vol. 37, pp. 79868–79888 (2024). <https://doi.org/10.52202/079017-2537>
17. Lin, X., Yu, F., Hu, J., You, Z., Shi, W., Ren, J.S., Gu, J., Dong, C.: Harnessing diffusion-yielded score priors for image restoration. *ACM Trans. Graph.* **44**(6), 208:1–208:21 (2025). <https://doi.org/10.1145/3763346>
18. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Jiang, Q., Li, C., Yang, J., Su, H., Zhu, J., Zhang, L.: Grounding DINO: marrying DINO with grounded pre-training for open-set object detection. In: Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XLVII. vol. 15105, pp. 38–55 (2024). https://doi.org/10.1007/978-3-031-72970-6_3
19. Liu, Y., Zhou, H., Cui, B., Shang, W., Lin, R.: Erase diffusion: Empowering object removal through calibrating diffusion pathways. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 2418–2427 (2025). <https://doi.org/10.1109/CVPR52734.2025.00231>
20. Liu, Z., Mao, H., Wu, C., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 11966–11976 (2022). <https://doi.org/10.1109/CVPR52688.2022.01167>
21. Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: CVPR. pp. 11451–11461 (2022)
22. Luo, S., Tan, Y., Huang, L., Li, J., Zhao, H.: Latent consistency models: Synthesizing high-resolution images with few-step inference. *CoRR* **abs/2310.04378** (2023). <https://doi.org/10.48550/ARXIV.2310.04378>
23. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J., Ermon, S.: Sdedit: Guided image synthesis and editing with stochastic differential equations. In: The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022 (2022)

24. Mescheder, L.M., Geiger, A., Nowozin, S.: Which training methods for gans do actually converge? In: Dy, J.G., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018. vol. 80, pp. 3478–3487 (2018), <http://proceedings.mlr.press/v80/mescheder18a.html>
25. Pathak, D., Krähenbühl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR. pp. 2536–2544 (2016)
26. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: SDXL: improving latent diffusion models for high-resolution image synthesis. In: The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024 (2024)
27. Porter, T.K., Duff, T.: Compositing digital images. In: Christiansen, H. (ed.) Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1984, Minneapolis, Minnesota, USA, July 23-27, 1984. pp. 253–259 (1984). <https://doi.org/10.1145/800031.808606>
28. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event. vol. 139, pp. 8748–8763 (2021), <http://proceedings.mlr.press/v139/radford21a.html>
29. Ravi, N., Gabeur, V., Hu, Y., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., Mintun, E., Pan, J., Alwala, K.V., Carion, N., Wu, C., Girshick, R.B., Dollár, P., Feichtenhofer, C.: SAM 2: Segment anything in images and videos. In: The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025 (2025), <https://openreview.net/forum?id=Ha6RTeWMd0>
30. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10674–10685 (2022)
31. Sagong, M., Yeo, Y., Jung, S., Ko, S.: RORD: A real-world object removal dataset. In: 33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022. p. 542 (2022)
32. Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. In: The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022 (2022), <https://openreview.net/forum?id=TIIdIXIpzhoI>
33. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXXXVI. vol. 15144, pp. 87–103 (2024). https://doi.org/10.1007/978-3-031-73016-0_6
34. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. In: International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA. vol. 202, pp. 32211–32252 (2023)
35. Sun, W., Dong, X., Cui, B., Tang, J.: Attentive eraser: Unleashing diffusion model’s object removal potential via self-attention redirection guidance. In: AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA. pp. 20734–20742 (2025). <https://doi.org/10.1609/AAAI.V39I19.34285>

36. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. In: IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2022, Waikoloa, HI, USA, January 3-8, 2022. pp. 3172–3182 (2022). <https://doi.org/10.1109/WACV51458.2022.00323>
37. Wei, R., Yin, Z., Zhang, S., Zhou, L., Wang, X., Ban, C., Cao, T., Sun, H., He, Z., Liang, K., Ma, Z.: OmniEraser: Remove objects and their effects in images with paired video-frame data. arXiv preprint arXiv:2501.07397 (2025). <https://doi.org/10.48550/arXiv.2501.07397>
38. Winter, D., Cohen, M., Fruchter, S., Pritch, Y., Rav-Acha, A., Hoshen, Y.: Objectdrop: Bootstrapping counterfactuals for photorealistic object removal and insertion. In: Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part LXXVII. vol. 15135, pp. 112–129 (2024)
39. Xu, N., Price, B.L., Cohen, S., Huang, T.S.: Deep image matting. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 311–320 (2017). <https://doi.org/10.1109/CVPR.2017.41>
40. Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T., Park, T.: One-step diffusion with distribution matching distillation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024. pp. 6613–6623 (2024). <https://doi.org/10.1109/CVPR52733.2024.00632>
41. Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Free-form image inpainting with gated convolution. In: 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019. pp. 4470–4479 (2019). <https://doi.org/10.1109/ICCV.2019.00457>
42. Yu, Y., Zeng, Z., Zheng, H., Luo, J.: OmniPaint: Mastering object-oriented editing via disentangled insertion-removal inpainting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17324–17334 (October 2025)
43. Zeng, Y., Fu, J., Chao, H., Guo, B.: Aggregated contextual transformations for high-resolution image inpainting. *IEEE Trans. Vis. Comput. Graph.* **29**(7), 3266–3280 (2023). <https://doi.org/10.1109/TVCG.2022.3156949>
44. Zhang, L., Agrawala, M.: Transparent image layer diffusion using latent transparency. *ACM Trans. Graph.* **43**(4), 100:1–100:15 (2024). <https://doi.org/10.1145/3658150>
45. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018. pp. 586–595 (2018). <https://doi.org/10.1109/CVPR.2018.00068>
46. Zhao, J., Zhou, S., Wang, Z., Yang, P., Loy, C.C.: Objectclear: Complete object removal via object-effect attention. *CoRR* **abs/2505.22636** (2025). <https://doi.org/10.48550/ARXIV.2505.22636>