UMAMI: Unifying Masked Autoregressive Models and Deterministic Rendering for View Synthesis

 $\begin{array}{cccccccc} \textbf{Thanh-Tung Le}^{*1} & \textbf{Tuan Pham}^{*1} & \textbf{Tung Nguyen}^2 \\ \textbf{Deying Kong}^3 & \textbf{Xiaohui Xie}^{1\dagger} & \textbf{Stephan Mandt}^{1\dagger} \\ & ^1 UCI & ^2 UCLA & ^3 Google \\ \end{array}$



Figure 1: UMAMI synthesizes photorealistic novel views from sparse inputs. Shown are single-view generation, three-view extrapolation, and six-view reconstruction. The hybrid model fuses deterministic rendering with diffusion-based completion for unseen regions, yielding fast and consistent results without explicit 3D priors.

Abstract

Novel view synthesis (NVS) seeks to render photorealistic, 3D-consistent images of a scene from unseen camera poses given only a sparse set of posed views. Existing deterministic networks render observed regions quickly but blur unobserved areas, whereas stochastic diffusion-based methods hallucinate plausible content yet incur heavy training- and inference-time costs. In this paper, we propose a hybrid framework that unifies the strengths of both paradigms. A bidirectional transformer encodes multi-view image tokens and Plücker-ray embeddings, producing a shared latent representation. Two lightweight heads then act on this representation: (i) a feed-forward regression head that renders pixels where geometry is well constrained, and (ii) a masked autoregressive diffusion head that completes occluded or unseen regions. The entire model is trained end-to-end with joint photometric and diffusion losses, without handcrafted 3D inductive biases, enabling scalability across diverse scenes. Experiments demonstrate that our method attains state-of-the-art image quality while reducing rendering time by an order of magnitude compared with fully generative baselines.

^{*} These authors contributed equally to this work. † denotes corresponding authors.

1 Introduction

Novel view synthesis (NVS) has long been a central problem in computer vision and graphics, aiming to generate realistic, 3D-consistent images of a scene from new camera viewpoints, using a given set of input views with known poses. Traditional methods often require dense input views, treating NVS as a sequential 3D reconstruction and rendering task [49, 34]. Recently, modern deep network priors [98, 11, 21] have been proposed to address the sparse views reconstruction problems, and achieve realistic rendering results.

Two dominant strategies have emerged for sparse-view NVS using deep networks: deterministic and generative-based methods. Deterministic methods often build generalizable networks that predict novel views by incorporating explicit 3D inductive biases [98, 11, 10, 83] or by leveraging priors from large-scale reconstruction models [27, 38, 31] with minimum inductive bias. While these approaches can be effective and fast in rendering observed regions, they often struggle with uncertainty in unobserved areas, leading to blurry predictions. Conversely, generative NVS approaches [87, 102, 21] can generate plausible content for unseen regions. These methods typically employ pretrained diffusion models conditioned on input views and camera poses. However, despite their strong generative capabilities, they often require extensive training data and computational resources, and their iterative sampling process leads to slow rendering speeds.

In this work, we address the question: "Can we combine the rendering efficiency of deterministic models with the generation capabilities of generative models?" We aim to unify these disparate approaches into a single, efficient framework. We observe that conventional diffusion models [102, 21], which iteratively generate full images using large UNet or Transformer backbones, can be inefficient if significant portions of the target view are already observable and could be rendered directly by a feed-forward network.

To this end, we introduce UMAMI, Unifying Masked Autoregressive Models and DeterminIstic Rendering for View Synthesis, a novel hybrid framework for NVS from sparse inputs. Our approach integrates a masked autoregressive model, trained with a diffusion loss [40], alongside a deterministic rendering head. Specifically, drawing inspiration from recent feed-forward NVS models [31], we employ a transformer with bidirectional attention. The model encodes input multi-view images tokens and masked target image tokens, conditioned on Plücker ray embeddings for both input and target views to a representation. The learned representation fulfills a dual role: (1) it conditions a lightweight MLP diffusion backbone that reconstructs unobserved regions through a diffusion loss [40], and (2) it serves as input to another MLP that directly renders pixel intensities for observed regions, trained with a photometric loss. Our method is designed to be purely data-driven, minimizing reliance on predefined inductive biases in its representation and rendering. This "inductive biasfree" design promotes scalability and generalizability, advantages empirically supported by prior work [21, 102, 31]. Ultimately, UMAMI aims to achieve accurate, training-efficient, and scalable novel view synthesis with photorealistic quality, enjoying both rapid rendering and robust generative completion.

We comprehensively evaluate our model through extensive experiments on RealEstate10K [103] and DL3DV [41], demonstrating competitive performance across both interpolation and extrapolation settings, and under varying input-view configurations.

Our contributions as as follows:

- A hybrid framework for NVS: We propose UMAMI, a novel hybrid architecture that combines
 deterministic and diffusion-based generation to effectively synthesize both visible and
 occluded regions from sparse views.
- We demonstrate that UMAMI achieves state-of-the-art performance across multiple benchmarks and input settings, while offering favorable trade-offs between speed and quality.

2 Related works

Novel view synthesis (NVS) is a rapidly advancing field. This section summarizes key prior works most relevant to our approach, with a more exhaustive review provided in the Appendix.

2.1 Novel view synthesis (NVS)

Novel view synthesis (NVS) has traditionally relied on image-based rendering that blends reference views with proxy geometry [14, 25, 67], light-field techniques that sample the plenoptic function from dense inputs [13], and learning-based variants that predict blending weights or depth maps with CNNs [12]. While multiview-stereo reconstructions enlarge the valid viewing volume [28, 7, 53], the breakthrough NeRF model introduced a differentiable volumetric representation whose photometric training signal became the new benchmark for NVS [49]. Subsequent work has pushed NeRF toward higher fidelity [1, 78, 2], faster inference [57, 24, 58], and fewer input views [50, 84], or has hybridized it with explicit structures such as dense or sparse voxels [72, 42, 20], low-rank decompositions and hashing [4, 8, 9], or point/gaussian primitives [92, 100, 19, 34]. Despite significant progress in rendering quality, these per-scene optimization methods often suffer from slow training times and limited generalization to novel scenes.

2.2 Deterministic NVS

To address the limitations of per-scene optimization, deterministic NVS methods train a single network across multiple scenes for fast, feed-forward inference. Some approaches, such as PixelSplat [6], MVSplat [11], and NoPoSplat [96], learn to predict 3D Gaussian parameters directly. While efficient, their reliance on specific 3D representations (e.g., NeRF [49], 3D Gaussians [34]) can hinder scalability. Alternatively, data-driven methods like LVSM [31] and SRT [63] leverage Transformeronly backbones to map input images and target poses to novel views, demonstrating the potential to synthesize views without explicit 3D representations given sufficient data and careful network design. Although scalable and fast, the deterministic nature of these methods typically restricts view generation to regions observed in the input context. Our method, in contrast, aims to synthesize novel views even when parts of the scene are occluded or outside the context views.

2.3 Generative NVS

In addition to deterministic approaches, generative approaches have adapted powerful image and video diffusion models (DMs) [3, 69] for NVS [102, 80, 87, 36], leveraging their strong generative priors. Early diffusion-based NVS models [66, 44, 48] often utilized image DMs conditioned on input images. Contemporary methods increasingly adopt video DMs [21, 102], conditioned on camera poses, to achieve finer-grained control and generate high-quality views of unseen regions. However, training these large-scale generative models demands substantial data and computational resources, potentially impacting rendering performance.

3 Background

3.1 Novel View Synthesis

Deterministic approaches focus on learning a mapping $f_{\theta}(\mathbf{I}^{\text{ctx}}, \boldsymbol{\pi}^{\text{ctx}}, \boldsymbol{\pi}^{\text{tgt}})$ that directly generates the target image \mathbf{I}^{tgt} . Here, \mathbf{I}^{ctx} and \mathbf{I}^{tgt} represent context and target images, while $\boldsymbol{\pi}^{\text{ctx}}$ and $\boldsymbol{\pi}^{\text{tgt}}$ denote their respective camera poses. This mapping f_{θ} may be realized through pure neural networks [31, 71] or by integrating 3D inductive biases [98, 6, 11, 96]. Although generally efficient, a fundamental limitation of deterministic methods is the inability to generate unseen region due to the deterministic nature.

Generative approaches learn to sample \mathbf{I}^{tgt} from a learned conditional distribution $p_{\theta}(\mathbf{I}^{\text{tgt}}|\mathbf{I}^{\text{ctx}},\boldsymbol{\pi}^{\text{ctx}},\boldsymbol{\pi}^{\text{tgt}})$. This distribution is often modeled using powerful generative frameworks such as diffusion models [102, 21]. The advantage of such generative techniques lies in their ability to convincingly hallucinate regions absent in the input views. Nevertheless, this capability comes at a significant computational cost for both training and inference, thereby posing challenges to their widespread practical use in NVS applications.

3.2 Masked Autoregressive Image Generatation

Unlike diffusion models, autoregressive (AR) models [76, 51, 22] approach the generation of an ordered token sequence $\{x^1, x^2, \dots, x^n\}$ (with $1 \le i \le n$ defining the order) by formulating

the problem as "next token prediction." This is mathematically expressed by factorizing the joint probability:

$$p(x^{1},...,x^{n}) = \prod_{i=1}^{n} p(x^{i}|x^{1},...,x^{i-1}),$$
(1)

where the conditional probability $p(x^i|x^1,\ldots,x^{i-1})$ is modeled by a neural network.

Departing from traditional AR methods [22, 51], the Masked Autoregressive (MAR) model [40] presents an different approach that unifies random-order AR principles with masked generative modeling through the use of a Diffusion Loss. In MAR, an autoregressive network produces a feature vector $z = f(\cdot) \in \mathbb{R}^D$. This vector, alongside a small MLP $\epsilon_{\theta}(\cdot)$, is used to model the conditional distribution p(x|z) for a token $x \in \mathbb{R}^d$. The model is trained using the denoising criterion:

$$\mathcal{L}(z,x) = \mathbb{E}_{\epsilon,t} \left[||\epsilon - \epsilon_{\theta}(x_t|t,z)||^2 \right], \tag{2}$$

where $\epsilon \in \mathbb{R}^d$ is Gaussian noise and $t \in \mathbb{R}$ is the timestep.

Compared to traditional diffusion models [59, 52], MAR sample an image by iteratively unmasking tokens using the MLP diffusion conditioned on learned latent from transformer. MAR demonstrates computational efficiency while showcasing competitive performance with its counterparts.

Building upon MAR's efficient generative capabilities, our work introduces a novel hybrid method for the NVS task. Specifically, we leverage MAR's generative framework within a hybrid network that incorporates deterministic rendering. Furthermore, we propose a unique sampler specifically designed to efficiently generate novel views, thereby avoiding the iterative full-image generation typical of large backbone architectures [102, 21]. This approach enables our generative solution to achieve rendering speeds an order of magnitude faster than previous generative NVS methods. We believe this to be the first proposal of a hybrid method that successfully unifies a generative model with a deterministic head to tackle the NVS challenge.

4 Methods

In this section, we first outline our problem formulation (Section 4.1) and then details our hybrid model (Section 4.2). Subsequently, we present the training loss (Section 4.3) and conclude by proposing a novel hybrid sampler (Section 4.4).

4.1 Problem Formulation

Given sparse input images with known camera poses $\{(\mathbf{I}^{\text{ctx}}, \boldsymbol{\pi}^{\text{ctx}})\}$, our goal is to model the conditional distribution $p(\mathbf{I}^{\text{tgt}}|\mathbf{I}^{\text{ctx}}, \boldsymbol{\pi}^{\text{ctx}}, \boldsymbol{\pi}^{\text{tgt}})$ to synthesize realistic novel view \mathbf{I}^{tgt} given its camera poses $\boldsymbol{\pi}^{\text{tgt}}$.

Data Representation To jointly encode image content and camera pose information, we follow the established convention of concatenating each image with its corresponding Plücker ray embeddings [54] along the channel dimension. This concatenated representation is subsequently processed through an MLP-based tokenizer to produce discrete latent tokens. For simplicity, we forego a VAE-based approach and directly tokenize each image into 8×8 patches. We denote the resulting token sequence from context images and their poses $\{(\mathbf{I}^{\text{ctx}}, \boldsymbol{\pi}^{\text{ctx}})\}$ as $\mathbf{c} = (c^1, c^2, ..., c^N)$; and similarly define the target token sequence from $\{(\mathbf{I}^{\text{tgt}}, \boldsymbol{\pi}^{\text{tgt}})\}$ as $\mathbf{x} = (x^1, x^2, ..., x^M)$. Ignoring the facts that the target camera pose are also embedded within \mathbf{x} and treat them purely as image tokens, we can write the target conditional distribution as $p(\mathbf{x}|\mathbf{c})$.

4.2 Hybrid Masked Autoregressive Models for Novel View Synthesis

As discussed in Section 3, deterministic-based NVS approaches [31, 63] model $p(\mathbf{x}|\mathbf{c})$ to be a deterministic function F of inputs: $p(\mathbf{x}|\mathbf{c}) = \delta(\mathbf{x} - F(\mathbf{c}))$, where δ is the Dirac delta function. While they have shown strong performance in generating high-fidelity outputs for regions covered by input views, they struggle to handle unseen regions due to their inability to model inherent ambiguity. In contrast, generative models based on diffusion [21, 102, 80] can generate plausible completions for unobserved regions, but often incur significantly higher computational costs due to iterative sampling over the full image. This trade-off motivates our hybrid design in UMAMI, which is based on the factorization:

$$p(\mathbf{x}|\mathbf{c}) = \delta(\mathbf{x}_D - F(\mathbf{c})) \cdot p(\mathbf{x}_S|\mathbf{x}_D, \mathbf{c})$$
(3)

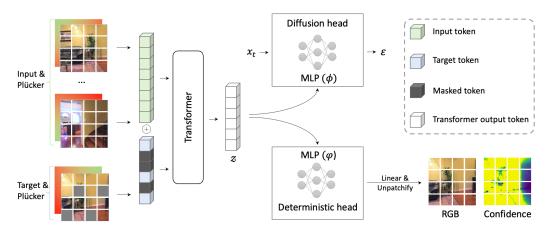


Figure 2: UMAMI synthesizes target images from their camera poses and context views (each paired with its Plücker pose). During training, we randomly mask the target image, replace masked areas with learnable tokens, and concatenate these with the target's Plücker embedding. Input views are also tokenized. A Transformer processes both tokenized inputs and the masked target representation to produce a latent z. This code inputs to two MLP heads: a deterministic head (φ) outputs RGB and confidence, while a diffusion head (ϕ) models the distribution of target tokens conditioned on z. The model is trained end-to-end using a weighted loss combination (Section 4.3). At inference, the target image is initialized with learned masked tokens for our proposed hybrid sampling (Section 4.4).

where \mathbf{x}_D and \mathbf{x}_S are disjoint subsets of \mathbf{x} such that $\mathbf{x} = \mathbf{x}_D \cup \mathbf{x}_S$. Intuitively, \mathbf{x}_D corresponds to the tokens that are fully determined by the input context \mathbf{c} (e.g., seen or deterministically visible regions) and can be computed directly as a function $F(\mathbf{c})$. In contrast, \mathbf{x}_S represents tokens in uncertain or unseen regions, which require sampling from a complex conditional distribution $p(\mathbf{x}_S|\mathbf{x}_D,\mathbf{c})$.

Model Architecture The architecture of UMAMI is illustrated in the Figure 2. UMAMI is a masked autoregressive model designed to support both efficient deterministic prediction and flexible stochastic generation by progressively unmasking target tokens. At the core of our model is a transformer backbone [77] that extract the target latent representation **z** from the partially masked **x** and context **c**. Following previous works [21, 31], we adopt a decoder-only, bi-directional transformer backbone.

To generate the target tokens, UMAMI uses two specialized output heads. The **deterministic head**, parameterized by φ in Figure 2, calculates $F(\mathbf{x})$ using the extracted latents \mathbf{z} from the transformer backbone and reconstructs tokens in \mathbf{x}_D in a single forward pass, leveraging regions of high confidence inferred from the context. In contrast, the **diffusion head**, parameterized by ϕ in Figure 2, models the conditional distribution over \mathbf{x}_S and performs iterative denoising to progressively generate plausible content in uncertain or unseen regions. Following MAR [40], both heads are small MLP networks with SiLU activation [18] that operate on each token latent individually, and the diffusion head takes an additional time embedding as input. This dual-head design enables UMAMI to adaptively combine the speed and accuracy of deterministic prediction with the generative capacity of diffusion models, effectively addressing both observed and novel view synthesis scenarios.

In practice, the separation between deterministic and uncertain regions is not known a priori. To address this, we introduce a pixel-wise confidence prediction that estimates an uncertainty score for each pixel. The confidence score of a patch is defined as the minimum confidence among its pixels. Given a threshold τ , patches with confidence above that threshold are assigned to \mathbf{x}_D , while the remaining are treated as \mathbf{x}_S and handled via the stochastic generation process.

4.3 Training Losses

We train UMAMI using a masked autoregressive generative framework [40, 5, 39]. At each training step, a binary mask **m** is sampled uniformly to mask a subset of the target image patches. Crucially, only the target image is masked (e.g. each selected patch is replaced with a learnable token) while the corresponding target camera pose embeddings are preserved. The model is then optimized to reconstruct the masked patches conditioned on the context and unmasked target information, using a combination of deterministic and diffusion losses:

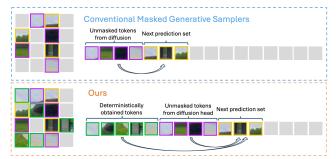


Figure 3: **Hybrid Masked Autoregressive Sampler.** Top: Conventional Masked Generative samplers [40, 5, 39] predict multiple tokens simultaneously using random ordering. Bottom (Ours): A deterministic first pass for high confidence tokens, followed by simultaneous random-order sampling for the remaining tokens, significantly boosts rendering times for the NVS task.

Deterministic reconstruction loss The deterministic head produces token predictions $\hat{\mathbf{x}}$ for masked patches, which yields reconstruction of the target images $\hat{\mathbf{I}}^{tgt}$. To supervise this process, we employ standard photometric losses for novel view synthesis, defined as:

$$\mathcal{L}_{render} = MSE(\hat{\mathbf{I}}^{tgt}, \mathbf{I}^{tgt}) + \lambda_{p} Perceptual(\hat{\mathbf{I}}^{tgt}, \mathbf{I}^{tgt}), \tag{4}$$

where λ_p is the weight for balancing the perceptual loss [33]. Importantly, this loss is computed over the full image rather than individual patches to encourage spatial consistency.

Confidence loss As discussed in Section 4.2, we augment the deterministic head to output a pixel-level confidence map \mathbf{s}_p , where each value indicates the model's confidence in its prediction. The confidence-aware loss balances the regression error with a regularization term that penalizes overconfidence:

$$\mathcal{L}_{\text{conf}} = \mathbf{m} \odot (\mathbf{s}_p \odot || \hat{\mathbf{I}}^{\text{tgt}} - \mathbf{I}^{\text{tgt}} ||_2^2 - \lambda_s \cdot \log \mathbf{s}_p), \tag{5}$$

where λ_s is a hyper-parameter controlling the regularization term [81]. The loss is averaged over all masked parts. We compute a patch-wise confidence map s by taking the minimum value of \mathbf{s}_p within each patch.

Diffusion loss To model the conditional distribution over uncertain tokens, we incorporate a diffusion model following the formulation of DDPM [26]. Specifically, we use a linear noise schedule to corrupt the ground truth tokens and train the model to reverse this process. Given a noisy token x_t at timestep t and its corresponding latent t extracted from the transformer backbone, the diffusion head predicts the added noise $\hat{\epsilon}$. The denoising objective is defined over all masked tokens:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{\epsilon, t} \left[||\epsilon - \hat{\epsilon}(x_t|t, z)||_2^2 \right], \tag{6}$$

where $\epsilon \sim \mathcal{N}(0, I)$ is Gaussian noise, and t is sampled uniformly over the diffusion steps.

To better allocate learning effort, we emphasize uncertain regions during training by predicting a token-wise weighting scheme derived from the patch-level confidence map s. Specifically, we define the weight for each token as $\max(s, \lambda_d)/\lambda_d$, where λ_d is a hyperparameter. This weighting encourages the model to focus more heavily on regions with lower confidence, enhancing generative quality in areas with higher ambiguity.

Total loss Our model is trained end-to-end using a weighted sum of the aforementioned losses.

4.4 Hybrid Masked Autoregressive Sampling

The overall sampling process is illustrated in Figure 3. Given a masked target image and its corresponding camera pose, UMAMI performs hybrid inference by first identifying and reconstructing the set of deterministic tokens \mathbf{x}_D , and then generating the remaining uncertain tokens \mathbf{x}_S through a diffusion-based process.

In the first stage, the model performs a single forward pass through the deterministic head to predict \mathbf{x}_D , guided by the confidence map predicted from the transformer backbone. Tokens with confidence scores greater then a predefined threshold τ are reconstructed deterministically.

In the second stage, the remaining masked tokens x_S , are iteratively sampled using the diffusion head. We employ a cosine unmasking schedule following the approach of [40], which gradually reveals

Table 1: Quantitative results on RealEstate10K across different validation splits. Best results are highlighted in red, second-best in orange.

M (1 1	D 00	Re10K-2view-extra		Re10K-2view-interp			Re10K-3view			
Method	Params (M)	PSNR ↑	LPIPS \downarrow	SSIM ↑	PSNR ↑	LPIPS \downarrow	SSIM ↑	PSNR ↑	LPIPS \downarrow	SSIM \uparrow
Deterministic										
MVSplat [11]	12.0	23.30	0.160	0.830	26.39	0.128	0.869	25.64	0.142	0.857
DepthSplat [91]	360	24.57	0.158	0.848	27.44	0.119	0.887	22.54	0.177	0.824
LVSM [31]	171	28.51	0.117	0.882	29.67	0.098	0.906	30.04	0.090	0.936
Diffusion-based										
ViewCrafter [99]	N/A	-	-	-	21.42	0.203	0.710	22.81	0.164	0.830
SEVA [102]	1300	24.00	0.100	0.797	25.66	0.061	0.847	27.57	0.073	0.892
UMAMI	271	28.95	0.107	0.897	28.85	0.101	0.899	31.06	0.084	0.946

more tokens at later iterations using T_S unmasking steps. As $|\mathbf{x}_S| \leq |\mathbf{x}|$, especially in scenarios where target views significantly overlap with context views, we introduce a dynamic strategy to adjust the number of unmasking steps accordingly. Specifically, given a maximum step budget T_{\max} for unmasking the entire token set \mathbf{x} , the number of steps allocated for \mathbf{x}_S is computed using a simple linear scaling rule: $T_S = \lceil |\mathbf{x}_S|/|\mathbf{x}| \cdot T_{\max} \rceil$. The hyperparameter T_{\max} is fixed across experiments, and T_S is automatically determined by the number of tokens to be unmasked. Despite its simplicity, we find this strategy to be effective in practice and well-suited for varying levels of token uncertainty.

5 Experiments

Datasets We evaluate UMAMI on two scene-level novel view synthesis benchmarks: RealEstate 10K (CC-BY-4.0) [103] and DL3DV (CC-BY-4.0) [41]. RealEstate 10K consists of 80K indoor and outdoor video clips sourced from YouTube, while DL3DV features over 10K videos captured across a wide range of real-world locations. We train separate models for each dataset at a resolution of 256×256 . For the RealEstate 10K dataset, we adopt the evaluation split from PixelSplat [6], which primarily features target views located between the 2 input views, making it suitable for assessing interpolation performance. We refer to this split as Re10K-2View-Interp. To evaluate extrapolation ability, we construct a complementary split by swapping the roles of the context and target views, which we denote as Re10K-2View-Extra. Additionally, we incorporate the 3-view validation split introduced in Reconfusion [89], labeled as Re10K-3View, respectively. For the DL3DV dataset, we follow the validation setup from Zhou et al. [102], using the 1-view, 3-view, and 6-view input configurations, which we name DL3DV-1View, DL3DV-3View, and DL3DV-6View, respectively.

Experiment Details Each model is trained for 100K iterations with a batch size of 32, using the AdamW optimizer [47] with a learning rate of 2×10^{-4} and a cosine decay schedule. Training takes approximately two days on $8\times$ NVIDIA A100 GPUs. During training, we randomly sample 1 or 2 context views and select between 1 and 3 target views per training example. In the main experiments, we report results for predicting a single target view, while results for generating multiple target views are included in the Appendix. To accelerate convergence, we initialize our model using the pretrained transformer backbone from LVSM [31]. We use a fixed threshold value of $\tau=0.95$ and a maximum sampling steps $T_{\rm max}=32$ across our experiments, as we found those values balance well between generation quality and speed. For diffusion sampling, we use 50 DDPM steps with a CFG value of 2.0 and a sampling temperature of 0.9. An details on other hyperparameters of our model are in the Appendix.

Baselines To the best of our knowledge, we are the first method that perform a hybrid render on deterministic and generative method, thus we have no direct competitors. Therefore, we compare our method to different deterministic and generative baselines. For deterministic methods, we compare UMAMI to MVSplat [11], LVSM [31]. For generative approaches, we compare UMAMI with ViewCrafter [99] and SEVA [102].

5.1 Experiment results

Quantitative results Tables 1 and 2 present a comprehensive comparison of our method UMAMI against both deterministic and diffusion-based baselines on the RealEstate10K and DL3DV datasets, respectively. On RealEstate10K, UMAMI consistently achieves top-tier performance across all splits. On the Re10K-2view-interp split, it matches LVSM closely, trailing by only 0.82 PSNR (28.85 vs.

Table 2: Quantitative results on DL3DV across 1-view, 3-view, and 6-view settings. Best results are highlighted in red, second-best in orange.

Method	DL3DV-1view			DL3DV-3view			DL3DV-6view		
Method	PSNR ↑	LPIPS \downarrow	SSIM ↑	PSNR ↑	LPIPS \downarrow	SSIM ↑	PSNR ↑	LPIPS \downarrow	SSIM \uparrow
Deterministic									
DepthSplat [91]	9.63	0.580	0.349	12.52	0.405	0.452	15.72	0.481	0.513
Diffusion-based	Diffusion-based								
ViewCrafter [99]	8.97	0.616	0.323	11.50	0.576	0.400	13.78	0.558	0.469
SEVA [102]	13.01	0.484	0.360	15.95	0.316	0.480	17.98	0.232	0.546
UMAMI	12.81	0.574	0.269	16.37	0.386	0.444	17.33	0.326	0.476

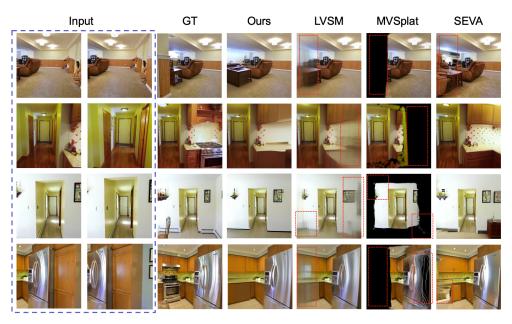


Figure 4: **Qualitative results on Re10K** Evaluation of UMAMI on the challenging *Re10K-2View-Extra* extrapolation set, comparing it with LVSM [31], MVSplat [11], and SEVA [102]. UMAMI not only renders sharp details in observed regions but also generates plausible content for unseen areas. More results can be viewed in the Appendix.

29.67). On the Re10K-2view-extra split, it obtains the highest PSNR (28.95) and SSIM (0.897), outperforming the second-best LVSM by 0.44 PSNR and 0.015 SSIM, while maintaining a second-best LPIPS of 0.107; showing its capabilities of doing extrapolation. On the Re10K-3view setting, UMAMI surpasses all baselines with the best PSNR (31.06) and SSIM (0.946), while having the second best LPIPS of 0.084.

On the DL3DV benchmark, UMAMI delivers competitive performance across all input configurations. It achieves the second-best PSNR in both the 1-view (12.81) and 6-view (17.33) settings, and ranks second in LPIPS across all three input setups. Notably, in the 3-view setting, UMAMI achieves the highest PSNR among all methods. While its SSIM lags behind DepthSplat, UMAMI outperforms ViewCrafter in both the 3-view and 6-view scenarios. These results underscore the robustness and adaptability of our hybrid framework, demonstrating its effectiveness in handling diverse view configurations and maintaining a strong balance between reconstruction fidelity and perceptual quality.

Varying overlap ratios Following NoPoSplat [96], we evaluate UMAMI on Re10k test set with varying camera overlaps based on ratio of image overlap: small (0.05%-0.3%), medium (0.3%-0.55%), and large (0.55%-0.8%), determined using dense feature matching method, RoMA [17]. The results, shown in the Table 3, demonstrate that UMAMI outperforms NoPoSplat on all metrics and validation sets. As with NoPoSplat, we also observed that performance improves as the overlap ratio increases, which suggests less scene occlusion. Notably, our method remains robust even with a small camera overlap, outperforming NoPoSplat across all metrics. This confirms UMAMI's effectiveness across various datasets and overlap ratios.

Table 3: Quantitative comparison across overlaping ratios (Small, Medium, Large). Best results are highlighted in red, second-best in orange.

Method	PSNR ↑	Small SSIM ↑	LPIPS ↓	PSNR ↑	Medium SSIM ↑	LPIPS ↓	PSNR ↑	Large SSIM ↑	LPIPS ↓
NoPoSplat	22.514	0.784	0.210	24.899	0.839	0.160	27.411	0.883	0.119
Ours	23.558	0.806	0.176	26.713	0.862	0.130	29.782	0.907	0.094

Qualitative results presented in Figure 4, highlight the performance of our method against several methods. Firstly, MVSplat [11], as a deterministic method employing 3D Gaussians, is unable to generate content beyond the provided context images, resulting in black rendered areas in unobserved regions. Similarly, while LVSM [31] avoids such black areas by forgoing 3D inductive biases, its non-generative nature results in blurry predictions for unseen pixels. Our method overcomes these limitations of deterministic approaches, demonstrating the ability to both accurately render observed regions and plausibly generate content in unobserved areas. Finally, in comparison to SEVA [102], a considerably larger model with 1.3B parameters (versus our 271M parameters), our approach achieves comparable performance on visible regions and produces results with fewer artifacts.

5.2 Ablation study and Analysis

Table 4: Ablation study on threshold τ . Higher τ Table 5: Ablation study on the number of context improves image quality but increases transformer calls and runtime. Table 5: Ablation study on the number of context views. Increasing the number of context views (N_c) improves image synthesis quality by pro-

au	Time (s)	# Trans. Calls	LPIPS \downarrow
0	0.02	1.00	0.398
0.5	2.71	12.31	0.394
0.8	4.30	18.99	0.389
0.9	4.62	20.30	0.387
0.95	4.77	21.14	0.386
1	7.63	33.00	0.377

Table 5: Ablation study on the number of context views. Increasing the number of context views (N_c) improves image synthesis quality by providing more deterministic tokens and reducing the number of average transformer calls, due to higher confidence in a larger portion of the scene.

N_c	# Deter Tokens	# Trans. Calls	LPIPS ↓
1	119.24	29.55	0.574
3	394.68	21.14	0.386
6	527.25	16.98	0.326

Effect on threshold τ during sampling We control the balance between UMAMI's deterministic and diffusion heads during sampling using a threshold τ in Table 4. Setting $\tau=0$ engages only the deterministic head, enabling UMAMI to predict all tokens in 0.02s with one transformer call. Conversely, setting $\tau=1$ relies exclusively on the diffusion head for sampling target views, which yields the optimal LPIPS score in our experiments. We observe that incrementally increasing τ from 0 to 1 enhances LPIPS performance, though at the cost of increased runtime due to more frequent transformer and diffusion sampling operations. Thus, by adjusting τ , our dual-head model offers a flexible mechanism to trade off inference speed against generative quality.

Effect on the number of context views We conduct an ablation study on the effect of varying the number of context views, as shown in 5. Results indicate that increasing the number of input views leads to a significant boost in mean LPIPS. Furthermore, with more context available, the model exhibits higher confidence, resulting in a greater proportion of tokens being handled deterministically. This, in turn, reduces the number of transformer calls required during sampling, leading to improved computational efficiency.

Run time analysis Unlike conventional methods with fixed rendering times, UMAMI offers operational flexibility by adaptively engaging its deterministic and generative heads. For instance, UMAMI renders an image in approximately 5s when $\tau=0.95$ (details in Table 4). This is considerably faster than generative counterparts like SEVA [102], which takes about 1 minute to sample an image. While purely deterministic methods [11] achieve sub-second rendering, they sacrifice the ability to generate content for unobserved target regions. UMAMI thus provides a compelling trade-off: it achieves strong generative capabilities for a modest increase in runtime compared to deterministic approaches, while remaining significantly more efficient than other generative models.

Backbone initialization As mentioned in Section 5, we initialize our model using pretrained weight from LVSM [31]. To further demonstrate the strength of our method, we train the model from scratch

Table 6: Ablation results on DL3DV across 1-view, 3-view, and 6-view settings comparing pretrained and random initialization. Best results are highlighted in red, second-best in orange.

Method	DL3DV-1view			DL3DV-3view			DL3DV-6view		
Method	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS \downarrow	SSIM ↑
Pretrained	12.81	0.574	0.269	16.37	0.386	0.444	17.33	0.326	0.476
Random Init.	11.80	0.543	0.256	14.46	0.374	0.370	15.43	0.318	0.404

on DL3DV with randomly initialized weights, without relying on LVSM pretrained on Re10K. This variant shares the same settings as the pretrained version, except for a larger batch size (512 vs. 32) to stabilize training. The results, presented in the Table 6, reveal that even without pretraining, our model performs comparably—showing slightly lower PSNR and SSIM but improved LPIPS. Remarkably, even with random initialization, our method consistently outperforms ViewCrafter and DepthSplat, underscoring that our performance stems from the strength of our hybrid deterministic-generative design, rather than dependence on LVSM initialization and Re10K pretraining.

6 Discussion and Conclusion

Limitations and Future Work While our method achieves competitive performance, it also has several limitations. First, because we operate directly in pixel space, each image is represented by a large number of tokens (e.g., $32 \times 32 = 1024$), which increases memory and computational requirements. A promising direction for future work is to adapt our framework to operate in the latent space of a pretrained VAE [35], which would reduce the token count while preserving semantic content. Second, unlike recent diffusion-based NVS approaches [21, 102, 99], our model does not make use of any pretrained text-to-image priors. Integrating such powerful generative priors [40, 16] could enhance the model's ability to hallucinate plausible unseen regions and improve visual fidelity in sparse-view settings. We also leave for future exploration techniques to further accelerate sampling and incorporate temporal consistency for video-based novel view synthesis. On the social impact side, this work could enable deepfake information, so users will be required to follow usage guidelines.

Conclusion We have presented UMAMI, a hybrid framework for novel view synthesis that unifies deterministic and generative modeling to handle both seen and unseen regions effectively. By leveraging a confidence-aware mechanism, our model adaptively allocates computation between a fast deterministic head and a diffusion-based head, achieving a strong balance between efficiency and image quality. Extensive experiments on RealEstate10K and DL3DV demonstrate that UMAMI is competitive with both deterministic and diffusion-only baselines across various input configurations. Our results suggest a promising direction for designing more efficient approaches to novel view synthesis.

Acknowledgements

Stephan Mandt acknowledges funding from the National Science Foundation (NSF) through an NSF CAREER Award IIS-2047418, IIS-2007719, the NSF LEAP Center, and the Hasso Plattner Research Center at UCI. Xiaohui Xie acknowledges funding from NIH 1P01CA288662-01A1 and Kay Family Foundation. Parts of this research were supported by the Intelligence Advanced Research Projects Activity (IARPA) via the Department of Interior/ Interior Business Center (DOI/IBC) contract number 140D0423C0075. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

- [1] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5855–5864, 2021.
- [2] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pages 19697–19705, 2023.
- [3] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- [4] E. R. Chan, C. Z. Lin, M. A. Chan, K. Nagano, B. Pan, S. De Mello, O. Gallo, L. J. Guibas, J. Tremblay, S. Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 16123–16133, 2022.
- [5] H. Chang, H. Zhang, L. Jiang, C. Liu, and W. T. Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern* recognition, pages 11315–11325, 2022.
- [6] D. Charatan, S. L. Li, A. Tagliasacchi, and V. Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, pages 19457–19467, 2024.
- [7] G. Chaurasia, S. Duchene, O. Sorkine-Hornung, and G. Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM transactions on graphics (TOG)*, 32(3):1–12, 2013.
- [8] A. Chen, Z. Xu, A. Geiger, J. Yu, and H. Su. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, pages 333–350. Springer, 2022.
- [9] A. Chen, Z. Xu, X. Wei, S. Tang, H. Su, and A. Geiger. Factor fields: A unified framework for neural fields and beyond. *arXiv preprint arXiv:2302.01226*, 2023.
- [10] A. Chen, Z. Xu, F. Zhao, X. Zhang, F. Xiang, J. Yu, and H. Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14124–14133, 2021.
- [11] Y. Chen, H. Xu, C. Zheng, B. Zhuang, M. Pollefeys, A. Geiger, T.-J. Cham, and J. Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. In *European Conference on Computer Vision*, pages 370–386. Springer, 2024.
- [12] I. Choi, O. Gallo, A. Troccoli, M. H. Kim, and J. Kautz. Extreme view synthesis. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 7781– 7790, 2019.
- [13] A. Davis, M. Levoy, and F. Durand. Unstructured light fields. In *Computer Graphics Forum*, volume 31, pages 305–314. Wiley Online Library, 2012.
- [14] P. E. Debevec, C. J. Taylor, and J. Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 465–474. 2023.
- [15] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 13142–13153, 2023.
- [16] H. Deng, T. Pan, H. Diao, Z. Luo, Y. Cui, H. Lu, S. Shan, Y. Qi, and X. Wang. Autoregressive video generation without vector quantization. *arXiv preprint arXiv:2412.14169*, 2024.

- [17] J. Edstedt, Q. Sun, G. Bökman, M. Wadenbäck, and M. Felsberg. Roma: Robust dense feature matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19790–19800, 2024.
- [18] S. Elfwing, E. Uchibe, and K. Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. arxiv e-prints, art. arXiv preprint arXiv:1702.03118, 2017.
- [19] W. Feng, J. Li, H. Cai, X. Luo, and J. Zhang. Neural points: Point cloud representation with neural fields for arbitrary upsampling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18633–18642, 2022.
- [20] S. Fridovich-Keil, A. Yu, M. Tancik, Q. Chen, B. Recht, and A. Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5501–5510, 2022.
- [21] R. Gao, A. Holynski, P. Henzler, A. Brussee, R. Martin-Brualla, P. Srinivasan, J. T. Barron, and B. Poole. Cat3d: Create anything in 3d with multi-view diffusion models. *Advances in Neural Information Processing Systems*, 2024.
- [22] K. Gregor, I. Danihelka, A. Mnih, C. Blundell, and D. Wierstra. Deep autoregressive networks. In *International Conference on Machine Learning*, pages 1242–1250, 2014.
- [23] H. He, Y. Xu, Y. Guo, G. Wetzstein, B. Dai, H. Li, and C. Yang. Cameractrl: Enabling camera control for text-to-video generation, 2024.
- [24] P. Hedman, P. P. Srinivasan, B. Mildenhall, J. T. Barron, and P. Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5875–5884, 2021.
- [25] B. Heigl, R. Koch, M. Pollefeys, J. Denzler, and L. Van Gool. Plenoptic modeling and rendering from image sequences taken by a hand-held camera. In *Mustererkennung 1999:* 21st DAGM Symposium, pages 94–101, 1999.
- [26] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020.
- [27] Y. Hong, K. Zhang, J. Gu, S. Bi, Y. Zhou, D. Liu, F. Liu, K. Sunkavalli, T. Bui, and H. Tan. LRM: Large reconstruction model for single image to 3d, 2024.
- [28] M. Jancosek and T. Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3121–3128, 2011.
- [29] H. Jiang, Z. Jiang, K. Grauman, and Y. Zhu. Few-view object reconstruction with unknown categories and camera poses. In *Proceedings of the International Conference on 3D Vision (3DV)*, pages 31–41, 2024.
- [30] H. Jiang, Z. Jiang, Y. Zhao, and Q. Huang. Leap: Liberate sparse-view 3d modeling from camera poses, 2023.
- [31] H. Jin, H. Jiang, H. Tan, K. Zhang, S. Bi, T. Zhang, F. Luan, N. Snavely, and Z. Xu. LVSM: A large view synthesis model with minimal 3D inductive bias, 2024.
- [32] M. M. Johari, Y. Lepoittevin, and F. Fleuret. Geonerf: Generalizing nerf with geometry priors, 2022.
- [33] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and superresolution. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 694–711, 2016.
- [34] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023.
- [35] D. P. Kingma, M. Welling, et al. Auto-encoding variational bayes, 2013.

- [36] X. Kong, S. Liu, X. Lyu, M. Taher, X. Qi, and A. J. Davison. Eschernet: A generative model for scalable view synthesis, 2024.
- [37] J. Kulhánek, E. Derner, T. Sattler, and R. Babuška. Viewformer: Nerf-free neural rendering from few images using transformers. In *European Conference on Computer Vision (ECCV)*, 2022.
- [38] J. Li, H. Tan, K. Zhang, Z. Xu, F. Luan, Y. Xu, Y. Hong, K. Sunkavalli, G. Shakhnarovich, and S. Bi. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model, 2023.
- [39] T. Li, H. Chang, S. Mishra, H. Zhang, D. Katabi, and D. Krishnan. Mage: Masked generative encoder to unify representation learning and image synthesis. In *Proceedings of the IEEE/CVF* Conference on Computer Vision and Pattern Recognition, pages 2142–2152, 2023.
- [40] T. Li, Y. Tian, H. Li, M. Deng, and K. He. Autoregressive image generation without vector quantization. *Advances in Neural Information Processing Systems*, 37:56424–56445, 2024.
- [41] L. Ling, Y. Sheng, Z. Tu, W. Zhao, C. Xin, K. Wan, L. Yu, Q. Guo, Z. Yu, Y. Lu, et al. Dl3dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22160–22169, 2024.
- [42] L. Liu, J. Gu, K. Z. Lin, T. Chua, and C. Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pages 15651–15663, 2020.
- [43] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick. Zero-1-to-3: Zero-shot one image to 3d object, 2023.
- [44] Y. Liu, C. Lin, Z. Zeng, X. Long, L. Liu, T. Komura, and W. Wang. Syncdreamer: Generating multiview-consistent images from a single-view image, 2023.
- [45] Y. Liu, S. Peng, L. Liu, Q. Wang, P. Wang, C. Theobalt, X. Zhou, and W. Wang. Neural rays for occlusion-aware image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7824–7833, 2022.
- [46] X. Long, Y. Guo, C. Lin, Y. Liu, Z. Dou, L. Liu, Y. Ma, S. Zhang, M. Habermann, C. Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion, 2023.
- [47] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [48] A. Mercier, R. Nakhli, M. Reddy, R. Yasarla, H. Cai, F. Porikli, and G. Berger. HexaGen3D: StableDiffusion is just one step away from fast and diverse text-to-3D generation. arXiv preprint arXiv:2401.07727, 2024.
- [49] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421, 2020.
- [50] M. Niemeyer, J. T. Barron, B. Mildenhall, M. S. Sajjadi, A. Geiger, and N. Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5480–5490, 2022.
- [51] A. V. D. Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. In International Conference on Machine Learning, pages 1747–1756, 2016.
- [52] W. Peebles and S. Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023.
- [53] E. Penner and L. Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics*, 36(6):1–11, 2017.

- [54] J. Plücker. On a new geometry of space. *Philosophical Transactions of the Royal Society of London*, pages 725–791, 1865.
- [55] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3D using 2D diffusion. arXiv, 2022.
- [56] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents, 2022.
- [57] C. Reiser, S. Peng, Y. Liao, and A. Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14335–14345, 2021.
- [58] C. Reiser, R. Szeliski, D. Verbin, P. Srinivasan, B. Mildenhall, A. Geiger, J. Barron, and P. Hedman. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *ACM Transactions on Graphics*, 42(4):1–12, 2023.
- [59] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. 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.
- [60] R. Rombach, P. Esser, and B. Ommer. Geometry-free view synthesis: Transformers and no 3d priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14356–14366, 2021.
- [61] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi. Palette: Image-to-image diffusion models. In *Proceedings of SIGGRAPH*, 2022.
- [62] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2022.
- [63] M. S. Sajjadi, H. Meyer, E. Pot, U. Bergmann, K. Greff, N. Radwan, S. Vora, M. Lučić, D. Duckworth, A. Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6229–6238, 2022.
- [64] K. Sargent, Z. Li, T. Shah, C. Herrmann, H.-X. Yu, Y. Zhang, E. R. Chan, D. Lagun, L. Fei-Fei, D. Sun, et al. ZeroNVS: Zero-shot 360-degree view synthesis from a single real image. arXiv preprint arXiv:2310.17994, 2023.
- [65] R. Shi, H. Chen, Z. Zhang, M. Liu, C. Xu, X. Wei, L. Chen, C. Zeng, and H. Su. Zero123++: A single image to consistent multi-view diffusion base model, 2023.
- [66] Y. Shi, P. Wang, J. Ye, L. Mai, K. Li, and X. Yang. MVDream: Multi-view diffusion for 3d generation, 2023.
- [67] S. Sinha, D. Steedly, and R. Szeliski. Piecewise planar stereo for image-based rendering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1881–1888, 2009.
- [68] V. Sitzmann, S. Rezchikov, B. Freeman, J. Tenenbaum, and F. Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 19313–19325, 2021.
- [69] Stability AI. Stable diffusion 3 (technical preview). Press release, 2024.
- [70] M. Suhail, C. Esteves, L. Sigal, and A. Makadia. Generalizable patch-based neural rendering. In *European Conference on Computer Vision (ECCV)*, pages 156–174, 2022.
- [71] M. Suhail, C. Esteves, L. Sigal, and A. Makadia. Light field neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8269–8279, 2022.

- [72] C. Sun, M. Sun, and H. Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5459–5469, 2022.
- [73] S. Szymanowicz, E. Insafutdinov, C. Zheng, D. Campbell, J. Henriques, C. Rupprecht, and A. Vedaldi. Flash3d: Feed-forward generalisable 3d scene reconstruction from a single image, 2024.
- [74] J. Tang, Z. Chen, X. Chen, T. Wang, G. Zeng, and Z. Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. In *European Conference on Computer Vision*, pages 1–18. Springer, 2024.
- [75] J. Tung, G. Chou, R. Cai, G. Yang, K. Zhang, G. Wetzstein, B. Hariharan, and N. Snavely. Megascenes: Scene-level view synthesis at scale. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024.
- [76] A. Van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, et al. Conditional image generation with pixelcnn decoders. *Advances in neural information processing systems*, 29, 2016.
- [77] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [78] D. Verbin, P. Hedman, B. Mildenhall, T. Zickler, J. T. Barron, and P. P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [79] V. Voleti, C. Yao, M. Boss, A. Letts, D. Pankratz, D. Tochilkin, C. Laforte, R. Rombach, and V. Jampani. SV3D: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 439–457, 2025.
- [80] V. Voleti, C.-H. Yao, M. Boss, A. Letts, D. Pankratz, D. Tochilkin, C. Laforte, R. Rombach, and V. Jampani. SV3D: Novel multi-view synthesis and 3D generation from a single image using latent video diffusion. In *European Conference on Computer Vision*, 2024.
- [81] S. Wan, T.-Y. Wu, W. H. Wong, and C.-Y. Lee. Confnet: predict with confidence. In 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 2921–2925. IEEE, 2018.
- [82] P. Wang, H. Tan, S. Bi, Y. Xu, F. Luan, K. Sunkavalli, W. Wang, Z. Xu, and K. Zhang. PF-LRM: Pose-free large reconstruction model for joint pose and shape prediction, 2023.
- [83] Q. Wang, Z. Wang, K. Genova, P. P. Srinivasan, H. Zhou, J. T. Barron, R. Martin-Brualla, N. Snavely, and T. Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699, 2021.
- [84] Z. Wang, S. Wu, W. Xie, M. Chen, and V. A. Prisacariu. Nerf--: Neural radiance fields without known camera parameters, 2021.
- [85] Z. Wang, Z. Yuan, X. Wang, Y. Li, T. Chen, M. Xia, P. Luo, and Y. Shan. Motionctrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024.
- [86] D. Watson, W. Chan, R. Martin-Brualla, J. Ho, A. Tagliasacchi, and M. Norouzi. Novel view synthesis with diffusion models, 2022.
- [87] D. Watson, S. Saxena, L. Li, A. Tagliasacchi, and D. J. Fleet. Controlling space and time with diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2024.

- [88] X. Wei, K. Zhang, S. Bi, H. Tan, F. Luan, V. Deschaintre, K. Sunkavalli, H. Su, and Z. Xu. Meshlrm: Large reconstruction model for high-quality mesh, 2024.
- [89] R. Wu, B. Mildenhall, P. Henzler, K. Park, R. Gao, D. Watson, P. P. Srinivasan, D. Verbin, J. T. Barron, B. Poole, et al. Reconfusion: 3D reconstruction with diffusion priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21551–21561, 2024.
- [90] D. Xie, S. Bi, Z. Shu, K. Zhang, Z. Xu, Y. Zhou, S. Pirk, A. Kaufman, X. Sun, and H. Tan. Lrm-zero: Training large reconstruction models with synthesized data, 2024.
- [91] H. Xu, S. Peng, F. Wang, H. Blum, D. Barath, A. Geiger, and M. Pollefeys. Depthsplat: Connecting gaussian splatting and depth. *arXiv* preprint arXiv:2410.13862, 2024.
- [92] Q. Xu, Z. Xu, J. Philip, S. Bi, Z. Shu, K. Sunkavalli, and U. Neumann. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5438–5448, 2022.
- [93] Y. Xu, H. Tan, F. Luan, S. Bi, P. Wang, J. Li, Z. Shi, K. Sunkavalli, G. Wetzstein, Z. Xu, and K. Zhang. Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model, 2023.
- [94] Y. Yan, Z. Xu, H. Lin, H. Jin, H. Guo, Y. Wang, K. Zhan, X. Lang, H. Bao, X. Zhou, and S. Peng. Streetcrafter: Street view synthesis with controllable video diffusion models, 2024.
- [95] J. Yang, Z. Cheng, Y. Duan, P. Ji, and H. Li. Consistnet: Enforcing 3d consistency for multi-view images diffusion, 2023.
- [96] B. Ye, S. Liu, H. Xu, X. Li, M. Pollefeys, M.-H. Yang, and S. Peng. No pose, no problem: Surprisingly simple 3d gaussian splats from sparse unposed images. *arXiv preprint arXiv:2410.24207*, 2024.
- [97] J. Ye, P. Wang, K. Li, Y. Shi, and H. Wang. Consistent-1-to-3: Consistent image to 3d view synthesis via geometry-aware diffusion models, 2023.
- [98] A. Yu, V. Ye, M. Tancik, and A. Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021.
- [99] W. Yu, J. Xing, L. Yuan, W. Hu, X. Li, Z. Huang, X. Gao, T. Wong, Y. Shan, and Y. Tian. Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis, 2024.
- [100] Q. Zhang, S. Baek, S. Rusinkiewicz, and F. Heide. Differentiable point-based radiance fields for efficient view synthesis. In *SIGGRAPH Asia Conference Papers*, pages 1–12, 2022.
- [101] C. Zheng and A. Vedaldi. Free3d: Consistent novel view synthesis without 3d representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9720–9731, 2024.
- [102] J. J. Zhou, H. Gao, V. Voleti, A. Vasishta, C.-H. Yao, M. Boss, P. Torr, C. Rupprecht, and V. Jampani. Stable virtual camera: Generative view synthesis with diffusion models. *arXiv* preprint arXiv:2503.14489, 2025.
- [103] T. Zhou, R. Tucker, J. Flynn, G. Fyffe, and N. Snavely. Stereo magnification: Learning view synthesis using multiplane images. In SIGGRAPH, 2018.

Supplement to "UMAMI: Unifying Masked Autoregressive Models and Deterministic Rendering for View Synthesis"

A Failure cases

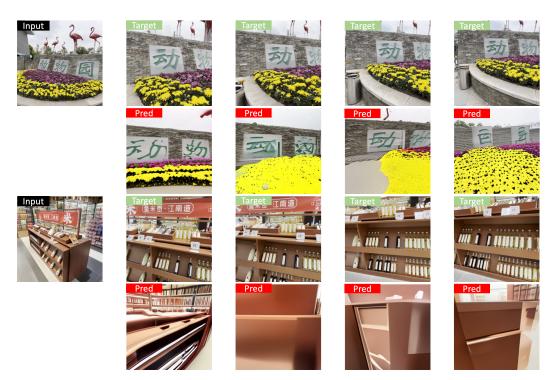


Figure 5: **Failure cases.** Our method may produce noticeable artifacts when target camera poses are too distant from the input view. Increasing the scale of training data and model parameters could improve the robustness of UMAMI.

B Related works

B.1 Feed-forward deterministic NVS methods

Early generalizable methods for Novel View Synthesis (NVS) demonstrated the potential of neural networks, trained across various scenes, to enable fast inference of novel views or underlying 3D representations in a feed-forward manner. Prominent examples include PixelNeRF [98], MVSNeRF [10], and IBRNet [83], which typically predict volumetric 3D representations by incorporating 3D-specific priors like epipolar geometry or plane sweep cost volumes. Subsequent research has extended these capabilities, improving performance particularly under challenging conditions such as sparse input views [45, 32, 29, 30], and adapting these techniques for emerging representations like 3D Gaussian Splatting (3DGS) [6, 73, 11, 74].

Recently, 3D Large Reconstruction Models (LRMs) have emerged [27, 38, 82, 93, 88, 90], leveraging the power of scalable transformer architectures [77] trained on extensive datasets to learn generic 3D priors. While these methods successfully avoid explicit architectural reliance on epipolar projection or cost volumes, they still typically depend on pre-defined 3D representations such as tri-plane NeRFs, meshes, or 3DGS, along with their corresponding rendering equations. This reliance can limit their flexibility and overall potential.

An alternative line of work attempts to directly learn a geometry-free rendering function [70, 63, 68, 60, 37]. However, these approaches often face limitations in model capacity and scalability, which can hinder their ability to capture high-frequency details. Notably, Scene Representation Transformers (SRT) [63] aimed to avoid explicit, handcrafted 3D representations by learning a latent

scene representation via a transformer, an objective shared by our encoder-decoder architecture. Despite this similarity, certain design choices in SRT, such as its CNN-based token extractor and the use of cross-attention in the decoder, have been shown to lead to less effective performance. To address the issue, LVSM [31] proposes a method that is fully transformer-based, leveraging bidirectional self-attention for enhanced representational power. Furthermore, they introduce a novel and more scalable decoder-only architecture that directly learns the NVS function with minimal 3D inductive bias and without relying on an intermediate latent representation.

Our proposed method adopts the versatile and scalable decoder-only transformer backbone from LVSM, which has demonstrated its efficacy in NVS tasks by leveraging a data-driven approach with minimal handcrafted 3D inductive bias. However, a crucial distinction lies in the nature of our approach: unlike the deterministic LVSM, our method is generative. We aim to address the inherent limitations of deterministic methods by harnessing the generative capabilities of masked autoregressive diffusion models in an efficient manner.

B.2 Generative-based NVS methods

The pursuit of generative-based (NVS) has recently seen significant advancements through the integration of diffusion models, drawing inspiration from successes in broader NVS [68, 63] and generative image-to-image tasks [61, 56, 62].

An early exploration in this domain was 3DiM [86], which trained image-to-image diffusion models for object-level multi-view rendering without explicit 3D representations. However, by training from scratch on limited 3D data, 3DiM's applicability was restricted to category-specific scenarios and lacked zero-shot generalization capabilities. Building on this, Zero-1-to-3 [43] adopted a similar geometry-free pipeline but significantly improved generalization and output quality by fine-tuning a pretrained 2D diffusion model on a larger 3D object dataset [15]. Despite these improvements, a key challenge for Zero-1-to-3 and other early image-based diffusion models for NVS (e.g., for distant viewpoints [65]) was multi-view inconsistency, as they typically generated each target view independently and probabilistically, leading to jitter or inconsistencies when rendering a camera trajectory.

To address this multi-view inconsistency, subsequent research diverged into several directions. One line of work focused on integrating explicit 3D inductive biases—such as 3D representations or epipolar attention—into the diffusion denoising process. Examples include SyncDreamer [44], ConsistNet [95], Consistent-1-to-N [97], and MegaScenes [75], though these often came at the cost of increased computation. Another set of approaches, including Instant3D [38], MVDream [66], and Wonder3D [46], aimed to predict a single grid of multiple, specific views simultaneously. While this improved consistency across those fixed views, it sacrificed the ability for fine-grained camera control. Works like MVDream [66], SyncDreamer [44], and more recently HexGen3D [48], generate multiple fixed views from a conditional image but do not support arbitrary viewpoint selection. To achieve consistent 3D object geometry from these image-based models, further steps like NeRF distillation, using techniques such as Score Distillation Sampling (SDS) [55, 64] or direct optimization on sampled images [89, 21], are often necessary. However, distillation techniques such as SDS can introduce substantial computational overhead due to test-time optimization.

More recently, a promising trend has emerged with models that jointly predict multiple target views while maintaining accurate camera control and ensuring view consistency, often through mechanisms like cross-view attention. This category includes methods such as Free3D [101], EscherNet [36], CAT3D [21], and SV3D [79]. Several video model-based approaches [85, 23, 99, 94, 102] also fall into this paradigm, increasing NVS performance. Despite these advancements, achieving high-quality generation with these recent models often necessitates substantial computational resources and extensive training data. Furthermore, their reliance on full-image iterative sampling typically results in slow inference times, limiting practical applicability. Our proposed method, UMAMI, addresses this critical issue by enabling photorealistic novel view rendering while maintaining efficient inference times.

Table 7: Hyperparameters for training UMAMI. We use the same set of hyperparameters for both RealEstate10K and DL3DV experiments.

Component	Parameter	Value
Image Tokenizer	Image size Patch size Channels	256 8 9 (3 RGB + 6 for Plücker)
Transformer	Layers Hidden dim Head dim QK Norm	24 768 64 True
Training	Batch size / GPU Num GPUS Learning rate Optimizer (β_1, β_2) Grad clip norm Mixed precision Weight decay Train steps Warmup steps	4 8 0.0002 (0.9, 0.95) 3.0 True 0.02 100k 1000
Data Setup	Input / Target views Center Crop	1 to 2 / 1 to 3 True
Loss Weights	L2 loss LPIPS loss Perceptual loss Diffusion loss Confidence loss	1.0 0.0 0.5 10

Implementation details

C.1 Hyperparamters

We report the hyperparameters used in Table 7.

C.2 Algorithm

We describe the sampling process of UMAMI in Algorithm 1.

Algorithm 1 Hybrid Inference in UMAMI

Require: Trained model, context views $\{(I_{\text{ctx}}, \pi_{\text{ctx}})\}$, target pose π_{tgt} , threshold τ , max unmasking

- 1: Tokenize context views into c, initialize target tokens x with masked tokens
- 2: Encode (\mathbf{c}, \mathbf{x}) with Transformer to obtain latent \mathbf{z}
- 3: Predict confidence map s_p and patch-level scores s
- 4: Partition target tokens:
 - Deterministic tokens: $\mathbf{x}_D \leftarrow \{x_i \mid s_i \geq \tau\}$
 - Stochastic tokens: $\mathbf{x}_S \leftarrow \{x_i \mid s_i < \tau\}$
- 5: Predict \mathbf{x}_D in one pass using deterministic head: $\hat{\mathbf{x}}_D = \phi(\mathbf{z}_D)$ 6: Compute sampling steps: $T_S = \lceil |x_S|/|x| \cdot T_{\max} \rceil$
- 7: for $t = T_S$ to 1 do
- Sample random unmasked set $\mathbf{x}_t \subset \mathbf{x}_S$ following a cosine scheduler. 8:
- Update \mathbf{x}_t by DDPM sampling using φ head. 9:
- 10: **end for**
- 11: Merge $\hat{\mathbf{x}}_D$ and $\hat{\mathbf{x}}_S$ into full target image \hat{I}_{tgt}

D Additional quantitative results

D.1 Multiple images generation

Table 8: Multi-view	generation results	on RealEstate10K.

Dataset	# gen views	PSNR ↑	LPIPS ↓	SSIM ↑
Re10K-2views-Extra	1 3	28.95 28.65	0.107 0.109	0.897 0.892
Re10K-2views-Interp	1 3	28.85 28.52	0.101 0.105	0.899 0.894

As shown in Table 7, our model is trained to predict up to three target views simultaneously. This joint prediction encourages consistency across generated images. In Table 8, we report results for generating one and three views. The generation quality is comparable across both settings. Notably, we use a fixed number of unmasking steps ($T_{\rm max}=32$) for all cases, which means generating multiple views in parallel can improve inference efficiency without sacrificing quality.

E Additional qualitative results

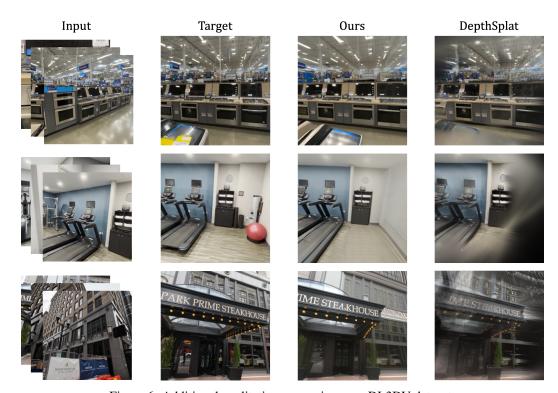


Figure 6: Additional qualitative comparisons on DL3DV dataset.

Additional qualitative evaluations are presented on the DL3DV dataset [41], where our method is compared against DepthSplat [91] under a three-view input configuration (Figure 6). As depicted UMAMI demonstrates notably sharp rendering, particularly in unobserved regions. This is achieved by leveraging its generative capabilities to synthesize plausible details unobserved region of input images.

Furthermore, to investigate the impact of the diffusion threshold hyperparameter, τ , on UMAMI's performance, its value was systematically varied, with findings illustrated in Figure 7. An initial setting of $\tau=0$, corresponding to a fully deterministic operation of UMAMI, achieved rapid inference.

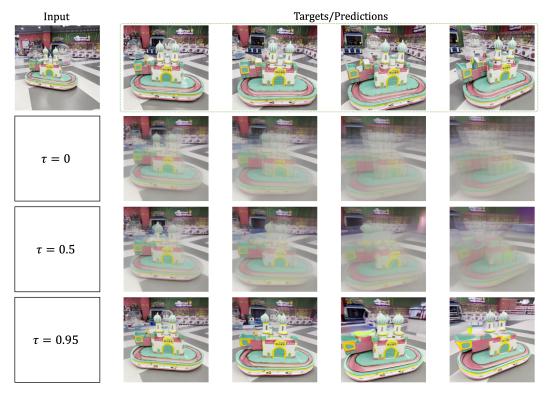


Figure 7: Impact of the diffusion threshold hyperparameter τ on rendering outcomes. The top row shows the single input view alongside four corresponding target views. The subsequent rows (2-4) illustrate the results as τ is incrementally increased. While a lower τ promotes deterministic behavior and faster inference, higher values of τ lead to notably sharper image rendering quality.

However, this configuration resulted in image blurring, an artifact attributable to unobserved regions in the input view. Progressively increasing τ to 0.5 and subsequently to 0.95 yielded a significant enhancement in rendering quality. This improvement, however, was accompanied by an increase in running time. Finally, to demonstrate the complete sampling dynamics of our method, the unmasking processes for $\tau=0.95$ and for full unmasking diffusion process ($\tau=1$) are presented in the supplementary video.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: In the abstract, we mainly claimed that our method can achieve state-of-the-art accuracy and can achieve faster speed compared to generative methods. There are supported with experimental evidence in our main paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have provided a dedicated subsection that discusses the limitations of our approach.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This work does not include theoretical analysis.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The "Experiments" section details the experimental setup, including GPU configurations, learning rates, optimizers, and the datasets used. Additionally, the code will be made publicly available upon acceptance.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The code is currently proprietary. We intend to open-source it upon acceptance. Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The settings are specified in the "Experiment" section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: The experiment works with large datasets, and the resources needed are such that running it multiple times for statistical significance would present a significant cost challenge.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We have included the information in the "Experiments" section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We have followed the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We have pointed out this work could be used to generate deepfake information. Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: As we only train our model using standardized NVS datasets, this does not apply to our work.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We explicitly stated the licenses for both datasets used and included references to their associated research papers.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new datasets or code are being submitted in conjunction with this manuscript.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.

• For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: This paper does not involve LLMs as any important, original, or non-standard components

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.