

You Only Need One Stage: Novel-View Synthesis From A Single Blind Face Image

Taoyue Wang¹, Xiang Zhang¹, Xiaotian Li¹, Huiyuan Yang², Lijun Yin¹

¹School of Computing, State University of New York at Binghamton

²Department of Computer Science, Missouri University of Science and Technology

twang61@binghamton.edu, zxiang4@binghamton.edu, xli210@binghamton.edu, hyang@mst.edu, lijun@cs.binghamton.edu

Abstract

We propose a novel one-stage method, **NVB-Face**, for generating consistent Novel-View images directly from a single **Blind Face** image. Existing approaches to novel-view synthesis for objects or faces typically require a high-resolution RGB image as input. When dealing with degraded images, the conventional pipeline follows a two-stage process: first restoring the image to high resolution, then synthesizing novel views from the restored result. However, this approach is highly dependent on the quality of the restored image, often leading to inaccuracies and inconsistencies in the final output. To address this limitation, we extract single-view features directly from the blind face image and introduce a feature manipulator that transforms these features into 3D-aware, multi-view latent representations. Leveraging the powerful generative capacity of a diffusion model, our framework synthesizes high-quality, consistent novel-view face images. Experimental results show that our method significantly outperforms traditional two-stage approaches in both consistency and fidelity.

1 Introduction

Reconstructing 3D structures or generating novel-view images from a single face image has long been a fundamental research topic, with applications spanning digital human modeling and 3D animation. Humans can effortlessly imagine different viewpoints of a given face due to their extensive prior knowledge of facial structures and multi-view consistency. However, in computer vision, faithfully synthesizing novel views while preserving the exact identity, expression, background, and other attributes remains an extremely challenging and ill-posed problem.

Currently, mainstream approaches for novel-view face generation fall into three major categories: regression-based models such as 3D Morphable Models (3DMM) (Paysan et al. 2009; Wu et al. 2019; Deng et al. 2019b; Dib et al. 2021; Zhuang et al. 2022), Neural Radiance Fields (NeRF)-based methods (Chan et al. 2021; Gu et al. 2021; Zhang et al. 2022; Chan et al. 2022; Yin et al. 2023; Yuan et al. 2023; An et al. 2023; Trevithick et al. 2023; Bhattarai, Nießner, and Sevastopolsky 2024), and generative methods leveraging Stable Diffusion and ControlNet (Rombach et al. 2022; Zhang, Rao,




Figure 1: We compare our method with typical two-stage pipelines, such as CodeFormer (Zhou et al. 2022) + PanoHead-PTI (An et al. 2023), which first restore the degraded image and then synthesize novel views. It is evident that when the restoration stage fails to recover accurate details, these errors are further amplified during the novel view synthesis, leading to results that deviate significantly from the original identity and appearance. In contrast, our method generates novel views in a single stage directly from the low-quality input. This end-to-end design suppresses error accumulation, resulting in more reliable and faithful novel-view images.

and Agrawala 2023; Gu et al. 2024, 2025). Despite their differences, these methods share a critical limitation—they **all require a high-quality, single-view face image as input**. Unfortunately, obtaining high-quality face images is not always feasible, as most images inevitably suffer from some degree of degradation. When dealing with low-resolution, blurry, noisy, or compression-degraded face images, the conventional approach involves a two-stage pipeline: first restoring the input image and then performing novel-view synthesis. However, this two-stage inference process introduces a fundamental issue: the performance of novel-view generation is entirely dependent on the quality of the restored image. Since the two stages operate independently without explicit connections, any errors in the upstream image restoration step get amplified in the downstream novel-view synthesis, significantly degrading the final output quality (as shown in Figure 1). Moreover, such a pipeline is inherently inefficient. Novel-view synthesis can only be performed after obtaining and evaluating the restoration results, often requiring additional selection or filtering to ensure quality. This dependency

introduces a major bottleneck for large-scale deployment, making it less practical for real-world applications.

To address this issue, we propose an end-to-end framework based on Stable Diffusion (Rombach et al. 2022) that enables the direct generation of high-quality novel-view images from a single blind face image in a single inference stage. Our overall pipeline begins with extracting features from a low-quality input face image. These features are then transformed into novel-view feature representations corresponding to the target camera viewpoint. Finally, we leverage the strong generative capability of Stable Diffusion to reconstruct high-resolution novel-view images from the transformed low-quality features. Specifically, to enable viewpoint transformation in latent space, we introduce a Transformer-based 3D Feature Construction Module, which builds a latent 3D face representation grid from the single-view input. By leveraging camera parameters, we are able to directly project this 3D representation into novel-view features within the latent space. It is important to highlight that, unlike previous diffusion-based novel-view synthesis methods (Papantoniou et al. 2024; Gu et al. 2024, 2025) that rely on ControlNet (Zhang, Rao, and Agrawala 2023), where camera parameters are first used to generate image templates as conditional inputs, our method directly performs viewpoint-conditioned mapping on the generated 3D feature grid using the camera parameters. This not only provides a more intuitive and flexible modeling strategy, but also significantly improves multi-view consistency due to the use of an explicit 3D feature representation. Furthermore, our pipeline employs a single-stage, end-to-end inference process, which avoids the error accumulation commonly found in multi-stage approaches, resulting in higher computational efficiency and stronger generalization.

Our contributions can be summarized as follows.

1. We propose a novel tuning-free framework that directly generates high-quality novel-view face images from a single blind face input at specified viewpoints. To the best of our knowledge, our approach is the first work to explore this end-to-end manner.
2. We introduce a novel 3D latent space representation of facial features, which enables consistent and accurate multi-view feature projection. This structured representation helps maintain cross-view consistency during novel-view synthesis.
3. Extensive qualitative and quantitative experiments demonstrate that our one-stage novel-view synthesis method achieves state-of-the-art performance, validating that our framework produces superior perceptual quality compared to traditional two-stage pipelines.

2 Related Work

Our proposed method integrates two challenging tasks: **blind face restoration** and **novel view synthesis**. In this section, we provide a brief overview of related works in both areas and discuss why directly combining these two tasks into a single inference stage for generating high-quality novel views from a degraded image is inherently difficult.

Blind Image Restoration. Restoring a degraded image into a high-resolution, detail-rich version is an inherently

ill-posed problem. Over the past decade, researchers have leveraged the powerful generative capabilities of deep models, particularly GANs (Yang et al. 2021; Dib et al. 2021; Wang et al. 2021a, 2022; Zhou et al. 2022; Wang et al. 2023a) and diffusion models (Wang et al. 2023b; Yue, Wang, and Loy 2023; Wang et al. 2024b; Lin et al. 2024; Yue and Loy 2024; Sun et al. 2024; Yue, Wang, and Loy 2023; Wu et al. 2025; Wang et al. 2024a), to reconstruct missing details in degraded images. GAN-based methods typically employ a U-Net-style generator to map degraded images to high-quality outputs, while a discriminator is used to enforce realism. However, GANs suffer from two major limitations. First, they are notoriously difficult to train, as maintaining the balance between the generator and discriminator remains a long-standing challenge. Second, their performance heavily depends on large-scale training data, often requiring millions of samples to generalize well. To address these issues, recent works have explored diffusion-based approaches for image restoration. Early diffusion-based methods (Wang et al. 2023b; Lin et al. 2024; Yue and Loy 2024; Sun et al. 2024) for blind image restoration typically condition on the degraded image, using modules such as ControlNet (Zhang, Rao, and Agrawala 2023) or a similar image encoder (Wang et al. 2024b) to extract features and inject them into a frozen pretrained Stable Diffusion model. The strong generative capability of Stable Diffusion then reconstructs the high-quality image. Given the compelling advantages of diffusion models, we also adopt a diffusion-based backbone. However, instead of following existing diffusion-based image restoration paradigms which require to first restore a single-view high-resolution image before performing a separate multi-view synthesis step, we propose a novel pipeline specifically designed to integrate seamlessly with the task of novel-view face synthesis.

Novel-View Face Synthesis. Traditional regression-based models, such as 3D Morphable Models (3DMM) (Paysan et al. 2009; Wu et al. 2019; Deng et al. 2019b; Dib et al. 2021; Zhuang et al. 2022), estimate parametric geometric priors of human faces and synthesize novel views by rendering 3D facial reconstructions. However, due to the limited expressiveness of the parameter space, 3DMM-based methods struggle to capture rich identity variations, expressions, and other facial attributes. As a result, researchers proposed NeRF-based methods (Chan et al. 2021; Gu et al. 2021; Zhang et al. 2022; Trevithick et al. 2023) or their more efficient evolution, triplane-based representations (Chan et al. 2022; Bhattarai, Nießner, and Sevastopolsky 2024; An et al. 2023; Yin et al. 2023; Yuan et al. 2023). Among these, EG3D (Chan et al. 2022) was the first method to leverage 3D-aware GANs to generate a latent triplane representation and render high-quality, multi-view consistent face images. Following this work, a series of studies (Yin et al. 2023; Yuan et al. 2023; Trevithick et al. 2023; Bhattarai, Nießner, and Sevastopolsky 2024) explored GAN inversion techniques to fit EG3D’s framework. However, all of these approaches rely on first extracting the camera viewpoint information from the input image to ensure consistency with EG3D’s triplane representation. If the input image is degraded, extracting accurate camera parameters becomes highly challenging, making it

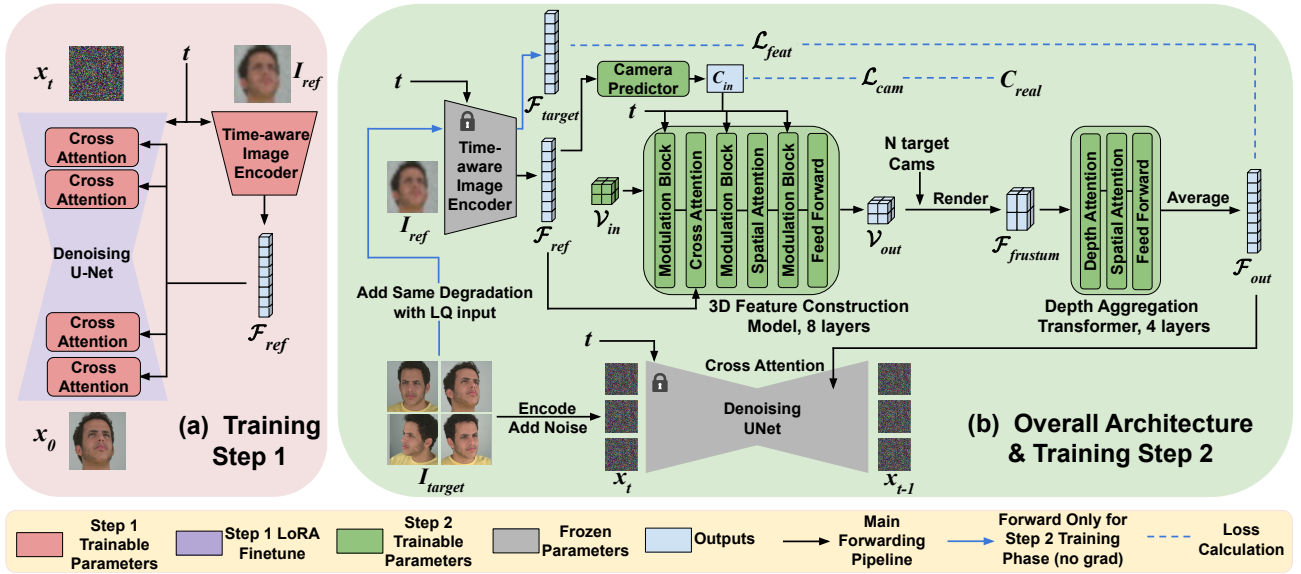


Figure 2: An overview of the proposed NVB-Face architecture. (a) Our first training step focuses solely image restoration. (b) In the second training step, we update only the parameters of the newly introduced modules (highlighted in dark green), keeping the rest of the network frozen. After training, this two-step process forms our complete inference pipeline.

necessary to first restore the image before novel-view synthesis — a process that inevitably introduces error accumulation and dependency on restoration quality. More recently, (Gu et al. 2024) introduced the first diffusion-based approach for novel-view face generation. However, our experiments reveal a major limitation of this approach that while it works well for high-quality input images, it fails significantly for degraded inputs, leading to severe information loss.

To overcome these challenges, we propose a single-stage framework that seamlessly integrates blind face restoration and novel-view synthesis into a unified process. Unlike previous methods, our approach is able to directly take any face image of arbitrary quality and generate high-resolution and consistent novel-view outputs at specified camera viewpoints.

3 Method

Given an input image \mathcal{I}_{ref} with an arbitrary level of degradation, our goal is to reconstruct the original details of \mathcal{I}_{ref} while generating a high-resolution image that preserves the same identity, expression, and attributes, but from the specified new camera viewpoint \mathcal{C}^i . To achieve this, we propose NVB-Face, a framework that is trained in two steps. First, an image encoder extracts latent features \mathcal{F}_{ref} from \mathcal{I}_{ref} , which are then fed into a Stable Diffusion (SD) model (Section 3.1). Leveraging its strong generative capabilities, the SD model restores fine details and reconstructs a high-quality version of the degraded image while maintaining identity consistency. Next, we introduce a 3D-aware viewpoint transformation model that takes the single-view feature \mathcal{F}_{ref} and transforms it into new features \mathcal{F}_{out}^i according to the i th camera viewpoints \mathcal{C}^i . The transformed features are then passed into our pretrained SD model to synthesize high-resolution novel-view images (Section 3.2). Finally, we design a tailored loss

function that effectively bridges blind face restoration and novel view synthesis into a unified training objective. We further provide theoretical and empirical justifications for how this loss formulation ensures seamless optimization of both tasks within a single-stage inference pipeline (Section 3.3).

3.1 Image Restoration

As depicted in Figure 2a, given an low-quality (LQ) image, the restoration component of our NVB-Face first employs an image encoder Enc to extract the feature \mathcal{F}_{ref} of reference image \mathcal{I}_{ref} , which are then fed into the SD model through cross-attention to generate a high-quality image. It is important to note that we did not apply average pooling to the output features of the last layer. Instead, we retained the original feature $\mathcal{F}_{ref} \in \mathbb{R}^{H \times W \times C}$ dimensions before feeding them into the Cross-Attention module. Since our task is image restoration, keeping the full latent spatial resolution of the features ensures that more fine-grained details from the input image are maintained, in order to faithfully reconstruct high-resolution outputs that accurately preserve the original contents.

Following the approach of (Liu et al. 2023; Wang et al. 2024b), we integrate temporal information into the image encoder to enhance generation quality and training stability. We incorporate the same time step embedding used in SD model into each image encoder block, making the image encoder time-aware, to ensure that the extracted features are synchronized with each time step of the downstream SD model.

$$\mathcal{F}_{ref} = Enc(\mathcal{I}_{ref}, t) \quad (1)$$

Since we replace the original CLIP (Radford et al. 2021) text encoder in the SD model with our custom image encoder, the existing cross-attention module, which was designed for

text conditioning, is no longer effective. To address this, we jointly fine-tune both image encoder and the cross-attention parameters. Additionally, to make the model better suited for our specific task while preserving the inherent generative capabilities of the SD model, we fine-tune all components of the SD model using LoRA (Hu et al. 2022; Luo et al. 2023), except for the cross-attention module.

The training objective remains consistent with the original SD model, where the model is trained to predict noise at each time step.

$$\mathcal{L}_{SD} = \mathbb{E}_{x_0, \mathcal{F}_{ref}, t, \epsilon \sim \mathcal{N}(0,1)} \left[\|\epsilon - \epsilon_\theta(x_t, t, \mathcal{F}_{ref})\|^2 \right], \quad (2)$$

where x_0 is the latent image encoded by VAE encoder, and is diffused t time steps into a Gaussian-distributed $x_t \sim \mathcal{N}(0, 1)$. And the ϵ_θ is a U-Net structured noise predictor.

In addition to using conventional high-quality face dataset, we also leverage multi-view face datasets during the first training step. This ensures that, regardless of the subject ID or viewpoint of the features generated in the second step, the SD model can consistently reconstruct high-resolution images with accurate identity and detail preservation. Details about the datasets used in our experiments will be discussed in Section 4.1.

3.2 Novel View Synthesis

3D Feature Construction Model. To generate novel-view images of \mathcal{I}_{ref} , previous approaches (Gu et al. 2024; Papantoniou et al. 2024) have attempted to synthesize a template image from the target viewpoint and feed it into ControlNet (Zhang, Rao, and Agrawala 2023) to condition the model with camera parameters. However, there are two major issues. First, in our case, the SD model has already been fine-tuned during the first training step specifically for the image restoration task, which may partially compromise its generative capability for other tasks such as novel view generation. Second, ControlNet cannot inherently guarantee multi-view consistency, necessitating the integration of a trainable View Consistency Module within the SD model (Gu et al. 2024), thereby increasing the overall model complexity.

Inspired by (Peebles and Xie 2023) and (Hong et al. 2023), We propose a transformer-based **3D Feature Construction Model**, denoted as *Trans*, which transforms the single-view feature \mathcal{F}_{ref} extracted from the image encoder into a 3D feature volume \mathcal{V}_{out} that fuses multi-view information, in order to enforces multi-view consistency in 3D space.

$$\begin{aligned} \mathcal{V}_{out} &= Trans(\mathcal{V}_{in}, \mathcal{F}_{ref}, \mathcal{C}_{in}, t), \\ \mathcal{C}_{in} &= CameraPredictor(\mathcal{F}_{ref}) \end{aligned} \quad (3)$$

where \mathcal{C}_{in} represents the camera parameters of the input image, while t denotes the time step embedding, which is synchronized with the image encoder and the SD model. This design is inspired by (Hong et al. 2023), which advocates constructing a 3D representation by leveraging both image features and corresponding camera parameters. The primary objective is to disentangle facial context (e.g., identity, expression) from pose, thereby enabling a more stable and consistent 3D feature representation. However, during inference, it is often impractical to obtain accurate viewpoint

information from degraded input images. To address this issue, we introduce a **Camera Predictor** module that estimates the camera parameters directly from the input image features. The predicted viewpoint is then used to supervise the generation of the 3D feature representation, ensuring it remains view-consistent even in the absence of explicit camera inputs at inference time.

To generate the 3D feature volume \mathcal{V}_{out} , we define a spatial positional embedding \mathcal{V}_{in} to serve as the query, which is then fed into Cross-Attention to integrate information from \mathcal{F}_{ref} , as illustrated in Figure 2b. This mechanism injects information beyond the current viewpoint, including identity, expression, attributes, and other contextual details. To condition the 3D representation on the input camera viewpoint, we introduce a Time-Aware Camera Modulation Block based on an Adaptive Layer Normalization, denoted as *ModLN* (the Modulation Block in Figure 2b)

$$\begin{aligned} \gamma, \beta &= MLP(\mathcal{C}_{in}), \\ ModLN(f') &= LayerNorm(f) \cdot (1 + \gamma) + \beta + t, \end{aligned} \quad (4)$$

which is applied between every attention module and feed-forward layer within the Transformer.

2D Features Sampling and Aggregation. Following (Liu et al. 2023) and (Yang et al. 2024), the 3D grid features are then warped into frustum volume features corresponding to the specified camera viewpoint \mathcal{C}^i by ray sampling and feature interpolation along each viewing direction

$$\mathcal{F}_{frustum}^i = Warp(\mathcal{V}_{out}, \mathcal{C}^i), \mathcal{F}_{frustum}^i \in \mathbb{R}^{D \times H \times W \times C}. \quad (5)$$

Since our method operates entirely in the latent space, we do not adopt volumetric rendering as used in NeRF (Mildenhall et al. 2021; Zhang et al. 2022), which fuses depth information based on explicit physical modeling. Instead, we design a **Depth Aggregation Transformer** *Aggr*, whose modules contain cross-depth attentions and spatial attentions (Fig 2b rightmost), to enhance the expressive ability of the resulting 2D features from each novel view. Finally, we apply average pooling along the depth dimension of the frustum volume to produce novel-view features \mathcal{F}_{out}^i that match the spatial dimensions of the input feature \mathcal{F}_{ref} .

$$\mathcal{F}_{out}^i = AvgPool(Aggr(\mathcal{F}_{frustum}^i)), \mathcal{F}_{out}^i \in \mathbb{R}^{H \times W \times C} \quad (6)$$

It is important to note that during the second training step, only the parameters of the 3D Feature Construction Model, Depth Aggregation Transformer and the Camera Predictor are updated. The parameters of the image encoder and the SD model remain frozen, as they were already fine-tuned for the image restoration task during the first training step. Modifying these parameters would negatively impact restoration quality. In summary, our approach decouples the target task into two independent objectives during training but fuses into a unified pipeline during inference. This modular design enhances the model’s scalability and robustness, while also simplifying the training complexity.

3.3 Loss Functions

In the second training step, we use the same diffusion loss as in Eq 2 but additionally introduce a feature loss \mathcal{L}_{feat} to align the generated novel-view features with their ground truth counterparts. The ground truth features are obtained by first degrading the novel-view ground truth images \mathcal{I}_{tgt} using the same degradation level applied to \mathcal{I}_{ref} . These degraded images are then passed through the pretrained image encoder from the first step to extract the ground truth novel-view features \mathcal{F}_{tgt} . Since the SD model has already learned how to restore high-resolution images from degraded features, applying the same noise pattern ensures that the generated features and ground truth features share the same distribution. This allows the model to focus solely on novel-view feature generation rather than on restoration, effectively decoupling the two tasks. The feature loss is formulated as below

$$\mathcal{F}_{tgt}^i = Enc(Degrade(\mathcal{I}_{tgt}^i), t),$$

$$\mathcal{L}_{feat} = \frac{1}{N} \sum_{i=1}^N \left((\mathcal{F}_{out}^i - \mathcal{F}_{tgt}^i)^2 + \lambda \left(1 - \frac{\mathcal{F}_{out}^i \cdot \mathcal{F}_{tgt}^i}{\|\mathcal{F}_{out}^i\| \|\mathcal{F}_{tgt}^i\|} \right) \right) \quad (7)$$

where N is the number of camera views and \mathcal{L}_{feat} comprises two components, MSE Loss and Square Cosine Loss, to jointly constrain the generated features at pixel level and feature space level. Additionally, to ensure that the camera parameters predicted by the Camera Predictor \mathcal{C}_{in} align with the ground truth of the input image, we compute an MSE loss between the predicted and ground-truth camera parameters, which referred to as the camera loss,

$$\mathcal{L}_{cam} = \|\mathcal{C}_{in} - \mathcal{C}_{real}\|_2^2 \quad (8)$$

Consequently, the total loss function in the second training step is defined as

$$\mathcal{L}_{total} = \mathcal{L}_{SD} + \lambda_1 \mathcal{L}_{feat} + \lambda_2 \mathcal{L}_{cam}. \quad (9)$$

4 Experiments

4.1 Experimental Settings

Datasets. Similar to the dataset setting in (Gu et al. 2024), we use a mixed training dataset for novel view synthesis in the second training step, which includes both photo-realistic multi-view dataset NeRSemble (Kirschstein et al. 2023) and synthetic data generated using PanoHead (An et al. 2023). NeRSemble dataset contains dynamic expression videos of 220 subjects captured from 16 synchronized fixed viewpoints. We sample 3000 sets of multi-view frames from this dataset. To expand the range of viewpoints, increase subject diversity, and enhance background variance, we further augment the training data with 3000 sets of images generated using PanoHead. This augmentation addresses the low diversity inherent in NeRSemble, as it is a lab-controlled dataset with limited variability. For our first training step, we enhance the model’s adaptability to multi-view faces by incorporating the above-mentioned datasets. Additionally, we mix in the high-resolution face dataset FFHQ (Karras, Laine, and Aila 2019) to improve the model’s generalization ability. During training, all input images are cropped and aligned following the preprocessing strategy used in EG3D. To generate the




Figure 3: Qualitative comparisons on NeRSemble (Kirschstein et al. 2023) dataset. As shown in our results, this end-to-end strategy achieves superior perceptual quality and preserves identity and expression information more effectively than two-stage methods, minimizing the loss of critical facial attributes.

LQ input images \mathcal{I}_{ref} , we adopt the Real-ESRGAN (Wang et al. 2021b) pipeline for degradation simulation. For quantitative evaluation, we sample 300 unseen sets of multi-view images from the NeRSemble dataset. To evaluate our model’s generalization ability on in-the-wild data, we perform qualitative evaluation on LFW-Test (Huang et al. 2008) dataset and quantitative evaluation for single view reconstruction on CelebA-Test (Karras et al. 2017) dataset.

Training details. In the first training step, the image encoder, the cross-attention layers of the SD model, and the LoRA parameters (with a rank of 16) are jointly optimized for 200K iterations, with a learning rate of $5e-5$ and a batch size of 8. In the second training step, we set the hyper-parameters in Eq 7 as $\lambda = 2.0$, $\lambda_1 = 10$ and $\lambda_2 = 0.05$ in Eq 9. For each input image, we randomly sample 7 different target viewpoints from the remaining views (the total views $N = 8$) to supervise the generation of the latent 3D representation. This step is trained with a learning rate of $1e-4$, a batch size of 2 for 300K iterations. During inference, we adopt the DDIM (Song, Meng, and Ermon 2020) denoising sampler with 50 steps. All training and inference experiments are conducted on a cluster of 8 NVIDIA RTX 6000 Ada GPUs.

4.2 Qualitative Comparisons

We conduct comparative experiments with several state-of-the-art novel-view synthesis methods, including PanoHead-PTI (An et al. 2023), GOAE (Yuan et al. 2023), TriPlaneNet (Bhattacharai, Nießner, and Sevastopolsky 2024) and DiffPortrait3D (Gu et al. 2024), and present qualitative comparisons on the NeRSemble (Kirschstein et al. 2023) and LFW-Test (Huang et al. 2008) datasets in Figure 3 and Figure 4, respectively. Among the compared methods, DiffPortrait3D is based on SD model, while the others are GAN-based. Since these existing methods are not designed to handle degraded facial images, we first apply CodeFormer (Zhou et al. 2022) for blind face restoration before passing the results to each method for novel view generation. Notably, GAN-based approaches require accurate camera parameters as input to perform novel view synthesis. However, due to the degraded quality of the CodeFormer-restored images, it



Figure 4: Qualitative comparisons on LFW-Test (Huang et al. 2008) dataset. Our method produces consistently stable results across varying levels of input degradation. Compared to other approaches, our generated images preserve the most information from the original input and exhibit higher visual realism, even under severe degradation.

is often difficult to extract reliable camera parameters from them. To ensure successful generation for these methods, we can only instead use the original high-quality images to extract the input camera parameters. Actually this is impossible in real-world blind restoration, hence highlights a critical limitation of two-stage methods: poor restoration quality makes accurate camera parameter estimation for novel view synthesis intractable. Our single-stage approach avoids this entirely, as it does not rely on external parameter estimation.

As shown in Figure 3, our method achieves results that are closest to the ground truth in terms of both identity and expression preservation. In contrast, restoration-based pipelines inevitably introduce errors during the blind restoration process, such as identity shift, expression shift, and visual artifacts. When these imperfect restorations are used as inputs for novel view generation, the errors are often amplified, leading to substantial information loss and noticeable degradation in image quality. We further evaluate our method on the LFW-Test dataset to demonstrate its robustness under varying levels of degradation in real-world conditions. We adjust the hyper-parameters of Real-ESRGAN (Wang et al. 2021b) to generate degradation level 1 and 2, from weak to strong. As shown in Figure 4, when the input is a clean image, our method produces results comparable to prior approaches. However, as the degradation level increases, the performance of other methods drops significantly, resulting in poor-quality novel views. While our method may also exhibit some identity and expression shifts, these changes are generally plausible predictions based on the degraded input, and the generated results remain substantially more stable

compared to other approaches.

4.3 Quantitative Comparisons

Similar to (Gu et al. 2024), we evaluate our method on both multi-view generation and single-view reconstruction tasks. For quantitative evaluation, we adopt a comprehensive set of metrics, including LPIPS (Zhang et al. 2018), DISTS (Ding et al. 2020), SSIM (Wang et al. 2004), ID similarity, FID (Heusel et al. 2017), and POSE error. The ID similarity is computed by extracting face embeddings from the generated images using the method in (Deng et al. 2019a), followed by measuring the cosine similarity with the ground truth embeddings. The POSE error is computed by estimating the facial pose using the approach from (Deng et al. 2019b) and calculating the Mean Squared Error (MSE) with respect to the estimated pose from ground truth images. For the multi-view generation task, we use our unseen test split of the NeRSemble dataset, where the model generates all the other novel views from a single degraded input image. For the single-view reconstruction task, we use degraded images from the CelebA-Test dataset and aim to reconstruct the original view from the same pose. Competing methods follow a two-stage pipeline, where the degraded image is first restored using CodeFormer (Zhou et al. 2022) before novel-view synthesis.

As shown in Table 1a, our proposed method outperforms all baseline approaches across all evaluation metrics. These results validate that our model can not only faithfully preserve the original content, but also robustly generate novel views, even from heavily degraded inputs.

	PanoHead-PTI	GOAE	TriplaneNet	DiffPortrait3D	Ours
SSIM \uparrow	0.73/0.62	0.73/0.58	0.75/0.63	0.68/0.62	0.78/0.63
LPIPS \downarrow	0.48/0.50	0.45/0.53	0.51/0.51	0.52/0.49	0.17/0.49
DISTS \downarrow	0.26/0.28	0.25/0.27	0.32/0.29	0.26/0.27	0.10/0.22
FID \downarrow	83.19/65.88	90.00/71.12	111.84/63.63	80.16/67.06	5.67/22.71
ID \uparrow	0.32/0.22	0.27/0.21	0.30/0.34	0.29/0.21	0.77/0.46
POSE \downarrow	0.0446/-	0.0425/-	0.0456/-	0.0341/-	0.0084/-

(a)

	w/o \mathcal{L}_{feat}	w/ \mathcal{L}_{feat}
SSIM \uparrow	0.75	0.78
LPIPS \downarrow	0.24	0.17
DISTS \downarrow	0.16	0.10
FID \downarrow	11.72	5.67
ID \uparrow	0.57	0.77

(b)

Table 1: (a) Quantitative comparisons with other methods on NeRSemble dataset (novel-view synthesis) / CelebA-Test (single-view reconstruction). (b) Ablation study on NeRSemble dataset to assess the effectiveness of the feature loss in our framework. Specifically, we compare the quantitative performance of our model with and without feature loss.



Figure 5: Qualitative ablation study on LFW-Test dataset to compare our method with and without feature loss.

4.4 Ablation Studies

Without Feature Loss. We conduct ablation studies on both NeRSemble and LFW-Test datasets to evaluate the impact of the feature loss on both quantitative and qualitative performance. In the ablated setting, we remove the feature loss and optimize only the diffusion loss, keeping all other training procedures identical. As shown in Table 1b, removing the feature loss results in a significant drop in performance on NeRSemble test set. To further investigate which aspect of performance is most affected, we evaluate the model on the LFW-Test dataset, which contains challenging in-the-wild samples. As shown in Figure ??, removing the feature loss leads to severe multi-view inconsistency in the generated outputs. We attribute this issue to the lack of explicit constraints in the latent space during training. Without feature loss, the model is only optimized to produce visually plausible final images, while ignoring the alignment and consistency of novel-view features in latent space. Although the powerful generative capacity of the diffusion model can superficially compensate for misaligned features during training, this shortcut fails to generalize. When tested on unseen data, the unaligned latent features lead to noticeable multi-view inconsistency, severely degrading the reliability of the generated results.

Feature Correction in Training Step 2 We further investigate whether suboptimal image restoration in training step 1 would negatively affect the final novel-view synthesis results. To this end, we use the image restoration model of training step 1 to recover high-resolution images from degraded inputs, and compare these with the final novel-view outputs



Figure 6: Feature correction of our training step 2.

produced by our whole NVB-Face model. As shown in Figure 6, even when the restored images from training step 1 are not visually satisfactory—indicating that the image encoder failed to extract all critical information—our training step 2 is still able to correct the imperfect features and produce more realistic novel-view images. This is made possible by our viewpoint transformation model, which is trained independently in training step 2. As a result, it can generate robust novel-view features even when the input features are partially flawed. As discussed earlier, traditional two-stage methods tend to amplify errors during novel-view synthesis when the restoration quality is suboptimal. In contrast, our single-stage approach is able to mitigate such error propagation, resulting in more stable and reliable outputs.

5 Conclusion

In this paper, we propose NVB-Face, the first framework capable of directly generating novel views from a single blind face image. NVB-Face decouples the overall task into two carefully designed and independently optimized steps, while introducing tailored loss functions that seamlessly integrate the two components. In contrast to traditional two-stage pipelines that perform face restoration followed by novel-view synthesis, our end-to-end approach significantly reduces error accumulation and better preserves identity and other essential image attributes. Furthermore, by constructing novel-view features in a 3D latent space, our method effectively enforces multi-view consistency across generated images. Extensive experiments demonstrate that our NVB-Face reliably synthesizes high-quality and consistent novel views for arbitrary in-the-wild blind face images.

6 Acknowledgments

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via [2022-21102100003]. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- An, S.; Xu, H.; Shi, Y.; Song, G.; Ogras, U. Y.; and Luo, L. 2023. Panohead: Geometry-aware 3d full-head synthesis in 360deg. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 20950–20959.
- Bhatarai, A. R.; Nießner, M.; and Sevastopolsky, A. 2024. Triplanenet: An encoder for eg3d inversion. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3055–3065.
- Chan, E. R.; Lin, C. Z.; Chan, M. A.; Nagano, K.; Pan, B.; De Mello, S.; Gallo, O.; Guibas, L. J.; Tremblay, J.; Khamis, S.; et al. 2022. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16123–16133.
- Chan, E. R.; Monteiro, M.; Kellnhofer, P.; Wu, J.; and Wetzstein, G. 2021. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5799–5809.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019a. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4690–4699.
- Deng, Y.; Yang, J.; Xu, S.; Chen, D.; Jia, Y.; and Tong, X. 2019b. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 0–0.
- Dib, A.; Thebault, C.; Ahn, J.; Gosselin, P.-H.; Theobalt, C.; and Chevallier, L. 2021. Towards high fidelity monocular face reconstruction with rich reflectance using self-supervised learning and ray tracing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12819–12829.
- Ding, K.; Ma, K.; Wang, S.; and Simoncelli, E. P. 2020. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5): 2567–2581.
- Gu, J.; Liu, L.; Wang, P.; and Theobalt, C. 2021. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *arXiv preprint arXiv:2110.08985*.
- Gu, Y.; Tran, P.; Zheng, Y.; Xu, H.; Li, H.; Karmanov, A.; and Li, H. 2025. DiffPortrait360: Consistent Portrait Diffusion for 360 View Synthesis. *arXiv preprint arXiv:2503.15667*.
- Gu, Y.; Xu, H.; Xie, Y.; Song, G.; Shi, Y.; Chang, D.; Yang, J.; and Luo, L. 2024. DiffPortrait3D: Controllable Diffusion for Zero-Shot Portrait View Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10456–10465.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Hong, Y.; Zhang, K.; Gu, J.; Bi, S.; Zhou, Y.; Liu, D.; Liu, F.; Sunkavalli, K.; Bui, T.; and Tan, H. 2023. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Huang, G. B.; Mattar, M.; Berg, T.; and Learned-Miller, E. 2008. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Kirschstein, T.; Qian, S.; Giebenhain, S.; Walter, T.; and Nießner, M. 2023. Nersemble: Multi-view radiance field reconstruction of human heads. *ACM Transactions on Graphics (TOG)*, 42(4): 1–14.
- Lin, X.; He, J.; Chen, Z.; Lyu, Z.; Dai, B.; Yu, F.; Qiao, Y.; Ouyang, W.; and Dong, C. 2024. Diffbir: Toward blind image restoration with generative diffusion prior. In *European Conference on Computer Vision*, 430–448. Springer.
- Liu, Y.; Lin, C.; Zeng, Z.; Long, X.; Liu, L.; Komura, T.; and Wang, W. 2023. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*.
- Luo, S.; Tan, Y.; Patil, S.; Gu, D.; von Platen, P.; Passos, A.; Huang, L.; Li, J.; and Zhao, H. 2023. Lcm-lora: A universal stable-diffusion acceleration module. *arXiv preprint arXiv:2311.05556*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Papantoniou, F. P.; Lattas, A.; Moschoglou, S.; Deng, J.; Kainz, B.; and Zafeiriou, S. 2024. Arc2face: A foundation model for id-consistent human faces. In *European Conference on Computer Vision*, 241–261. Springer.
- Paysan, P.; Knothe, R.; Amberg, B.; Romdhani, S.; and Vetter, T. 2009. A 3D face model for pose and illumination invariant face recognition. In *2009 sixth IEEE international conference on advanced video and signal based surveillance*, 296–301. Ieee.

- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Sun, L.; Wu, R.; Zhang, Z.; Yong, H.; and Zhang, L. 2024. Improving the stability of diffusion models for content consistent super-resolution. *CoRR*.
- Trevithick, A.; Chan, M.; Stengel, M.; Chan, E.; Liu, C.; Yu, Z.; Khamis, S.; Chandraker, M.; Ramamoorthi, R.; and Nagano, K. 2023. Real-time radiance fields for single-image portrait view synthesis. *ACM Transactions on Graphics (TOG)*, 42(4): 1–15.
- Wang, J.; Gong, J.; Zhang, L.; Chen, Z.; Liu, X.; Gu, H.; Liu, Y.; Zhang, Y.; and Yang, X. 2024a. OSDFace: One-Step Diffusion Model for Face Restoration. *arXiv preprint arXiv:2411.17163*.
- Wang, J.; Yue, Z.; Zhou, S.; Chan, K. C.; and Loy, C. C. 2024b. Exploiting diffusion prior for real-world image super-resolution. *International Journal of Computer Vision*, 132(12): 5929–5949.
- Wang, X.; Li, Y.; Zhang, H.; and Shan, Y. 2021a. Towards real-world blind face restoration with generative facial prior. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9168–9178.
- Wang, X.; Xie, L.; Dong, C.; and Shan, Y. 2021b. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1905–1914.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Wang, Z.; Zhang, J.; Chen, R.; Wang, W.; and Luo, P. 2022. Restoreformer: High-quality blind face restoration from undegraded key-value pairs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17512–17521.
- Wang, Z.; Zhang, J.; Chen, T.; Wang, W.; and Luo, P. 2023a. RestoreFormer++: Towards real-world blind face restoration from undegraded key-value pairs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wang, Z.; Zhang, Z.; Zhang, X.; Zheng, H.; Zhou, M.; Zhang, Y.; and Wang, Y. 2023b. Dr2: Diffusion-based robust degradation remover for blind face restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1704–1713.
- Wu, F.; Bao, L.; Chen, Y.; Ling, Y.; Song, Y.; Li, S.; Ngan, K. N.; and Liu, W. 2019. Mvf-net: Multi-view 3d face morphable model regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 959–968.
- Wu, R.; Sun, L.; Ma, Z.; and Zhang, L. 2025. One-step effective diffusion network for real-world image super-resolution. *Advances in Neural Information Processing Systems*, 37: 92529–92553.
- Yang, J.; Cheng, Z.; Duan, Y.; Ji, P.; and Li, H. 2024. Consistnet: Enforcing 3d consistency for multi-view images diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7079–7088.
- Yang, T.; Ren, P.; Xie, X.; and Zhang, L. 2021. Gan prior embedded network for blind face restoration in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 672–681.
- Yin, F.; Zhang, Y.; Wang, X.; Wang, T.; Li, X.; Gong, Y.; Fan, Y.; Cun, X.; Shan, Y.; Oztireli, C.; et al. 2023. 3d gan inversion with facial symmetry prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 342–351.
- Yuan, Z.; Zhu, Y.; Li, Y.; Liu, H.; and Yuan, C. 2023. Make encoder great again in 3d gan inversion through geometry and occlusion-aware encoding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2437–2447.
- Yue, Z.; and Loy, C. C. 2024. Difface: Blind face restoration with diffused error contraction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Yue, Z.; Wang, J.; and Loy, C. C. 2023. Resshift: Efficient diffusion model for image super-resolution by residual shifting. *Advances in Neural Information Processing Systems*, 36: 13294–13307.
- Zhang, J.; Li, X.; Wan, Z.; Wang, C.; and Liao, J. 2022. Fdnerf: Few-shot dynamic neural radiance fields for face reconstruction and expression editing. In *SIGGRAPH Asia 2022 Conference Papers*, 1–9.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhou, S.; Chan, K.; Li, C.; and Loy, C. C. 2022. Towards robust blind face restoration with codebook lookup transformer. *Advances in Neural Information Processing Systems*, 35: 30599–30611.
- Zhuang, Y.; Zhu, H.; Sun, X.; and Cao, X. 2022. Mofan-erf: Morphable facial neural radiance field. In *European conference on computer vision*, 268–285. Springer.