



VM-NeRF: Tackling Sparsity in NeRF with View Morphing

Matteo Bortolon^{1,2,3}(✉), Alessio Del Bue², and Fabio Poiesi^{1,2}

¹ TeV, Fondazione Bruno Kessler, Povo, Italy
{mbortolon,poiesi}@fbk.eu

² PAVIS, Istituto Italiano di Tecnologia, Genoa, Italy
{alessio.delbue}@iit.it

³ DISI, University of Trento, Trento, Italy

Abstract. NeRF aims to learn a continuous neural scene representation by using a finite set of input images taken from various viewpoints. A well-known limitation of NeRF methods is their reliance on data: the fewer the viewpoints, the higher the likelihood of overfitting. This paper addresses this issue by introducing a novel method to generate geometrically consistent image transitions between viewpoints using View Morphing. Our VM-NeRF approach requires no prior knowledge about the scene structure, as View Morphing is based on the fundamental principles of projective geometry. VM-NeRF tightly integrates this geometric view generation process during the training procedure of standard NeRF approaches. Notably, our method significantly improves novel view synthesis, particularly when only a few views are available. Experimental evaluation reveals consistent improvement over current methods that handle sparse viewpoints in NeRF models. We report an increase in PSNR of up to 1.8dB and 1.0dB when training uses eight and four views, respectively. Source code: <https://github.com/mbortolon97/VM-NeRF>.

1 Introduction

Novel View Synthesis (NVS) is the problem of synthesising unseen camera views from a set of known views¹ [8, 29]. NVS is a key technology that can enable compelling augmented or virtual reality experiences [10], new entertainment technology [6], and robotics applications [11]. NVS has undergone a significant improvement after the introduction of Neural Radiance Fields (NeRF) [2, 17] – a trainable implicit neural representation of a 3D scene that can photorealistically render unseen (novel) views. NeRF is a data-driven model that can synthesise high-quality novel views but in general requiring several multi-view images, e.g. about hundreds of images taken from different and uniformly distributed camera viewpoints around an object of interest [17]. If these viewpoints are few and/or not uniformly distributed, the resulting NeRF model may fail to produce

¹ Throughout the paper, we will use the term *viewpoint* to refer to the camera pose, *view* to refer to the scene seen through a certain viewpoint and to *image* to refer to the photometric content captured from a view.

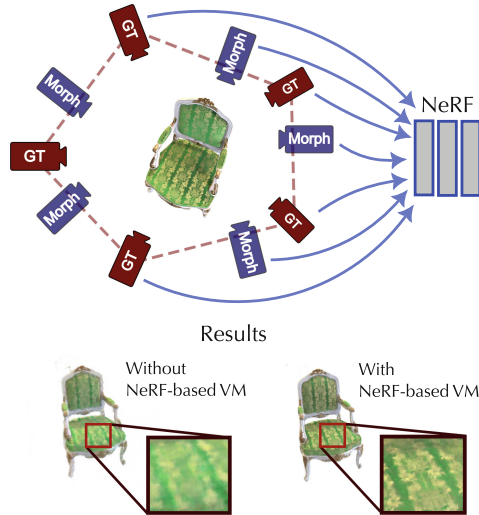


Fig. 1. Given a set of known views (ground truth), View Morphing-NeRF (VM-NeRF) generates image transitions between views (morph) that can be effectively used to train a NeRF model in the case of few-shot view synthesis. Results are of a higher quality when VM-NeRF is used.

satisfactory novel views [12, 16]. This detrimental effect is a known drawback of NeRF-based approaches and it is due to the likelihood of overfitting on known viewpoints while decreasing generalisation on novel views that are furthest from the given viewpoints, namely the few-shot view synthesis problem [12].

In this paper, we propose to tackle the problem of training a NeRF model on scenes captured with a sparse set of viewpoints by using a novel geometry-based strategy based on View Morphing [24] (Fig. 1). This purely geometric method can synthesise or morph a new viewpoint that lies in-between two given camera views while ensuring realistic image transitions. Traditionally, view morphing requires a set of accurate point matches between known image pairs in order to successfully perform the morph. As this matching stage is hard to integrate into a NeRF-based learning pipeline, our intuition is to leverage the per-image depth information implicitly estimated by NeRF to obtain dense coordinate matches among views after an image rectification stage (Fig. 2). To this end, we have to relax and modify several steps of the view morphing strategy to be duly integrated in the NeRF learning paradigm. This technique does not require any prior knowledge about the captured 3D scene, and it can synthesise 3D projective transformations (e.g. 3D rotations, translations, shears) of objects by operating entirely on the input images. We evaluate our approach by using the dataset of the original NeRF’s paper [17] and we show that PSNR improves up to 1.8dB and 1.0dB when eight and four views are used for training, respectively. We compare our approach with DietNeRF [12], AugNeRF [5] and RegNeRF [19], and show that our approach can produce higher-quality renderings.

To summarise, our contributions are:

- We present a novel and effective method for NeRF to address the problem of few-shot view synthesis;
- We introduce a new view morphing technique based on the NeRF depth output, named VM-NeRF;
- VM-NeRF can achieve higher-quality rendered images than alternative methods in the literature.

2 Related Work

NVS scene synthesis can be solved either by using traditional 3D reconstruction techniques [23] or by adopting methods based on neural rendering [26]. Neural Radiance Fields (NeRF) is a recent neural rendering method that can learn a volumetric representation of an unknown 3D scene approximating its radiance and density fields from a set of known (ground truth) views by using a multilayer perceptron (MLP) [17]. NeRF optimises its parameters on one scene based on a set of known views, thus overfitting can occur when these views are few.

Current approaches addressing few-shot novel view synthesis can be divided into two groups. The first group uses the same trained network to generate novel views of different scenes. This category of methods trains on datasets characterised by similar scenes, such as DTU [1]. Multiple-scene training can introduce dataset biases and may produce low-quality results in contexts outside the training domain [18, 27]. SparseNeuS [14] and ShaRF [22] train NVS on multiple scenes by conditioning the MLP with features that encode appearance and geometry of the surface at a 3D location. This can be achieved by using an auxiliary deep network jointly trained with NeRF. The second group uses the original per-scene optimisation procedure of NeRF, so a single network trains and tests only on one scene leading to methods without dataset bias. These methods are more likely to encounter overfit problems on the known views, however they reduce this likelihood by adding either semantic or geometric constraints during training. DietNeRF belongs to this category and exploits the feature representations of known images computed with a CLIP pre-trained image encoder, renders random poses, and processes them by imposing semantic consistency through CLIP features [12]. RegNeRF [19] renders random viewpoints around the known ones, and introduces regularisation constraints between known viewpoints and randomly sampled ones.

Single-scene methods working with few viewpoints may overfit on the known images, producing artefacts when novel views are rendered. In general, we can mitigate overfitting via data augmentation [25], and to the best of our knowledge, the only methods that address data augmentation for NeRF are AugNeRF [5] and GeoAug [4]. AugNeRF aims to improve NeRF generalisation by using adversarial data augmentation to enforce each ray and its augmented version to produce the same result. GeoAug [4] perturbs translation and rotation of the known viewpoints during training. Our proposed approach does not perturb

the known input views and rays, instead we create new views (novel 3D projective transformations) using pairs of known views. This allows us to enforce coherence of newly rendered viewpoints between distant viewpoint pairs. At the moment of the acceptance of this paper, we could not replicate the results of GeoAug because the authors have not released their source code.

3 Preliminaries

3.1 NeRF Overview

NeRF’s objective is to synthesise novel views of a scene by optimising a volumetric function given a finite set of input views [17]. Let f_{θ} be the underlying function we aim to optimise. The input to f_{θ} is a 5D datum that encodes a point on a camera ray, i.e. a 3D spatial location (x, y, z) and a 2D viewing direction (θ, ϕ) . Let $\mathbf{c} \in \mathbb{R}^3$ be the view-dependent emitted radiance (colour) and σ be the volume density that f_{θ} predicts at (x, y, z) . Novel views are synthesised by querying 5D data along the camera rays. Traditional volume rendering techniques can be used to transform \mathbf{c} and σ into an image [13, 15]. Because volume rendering is differentiable, f_{θ} can be implemented as a fully-connected deep network and learned.

Rendering a view from a novel viewpoint consists of estimating the integrals of all 3D rays that originate from the camera optic centre and that pass through each pixel of the camera image plane. Let \mathbf{r} be a 3D ray. To make rendering computationally tractable, each ray is represented as a finite set of 3D spatial locations, indexed with i , which are defined between two clipping distances: a near one (t_n) and a far one (t_f). Let Γ be the number of 3D spatial locations sampled between t_n and t_f . Rendering the colour of a pixel is given by

$$\hat{\mathbf{c}}(\mathbf{r}) = \sum_{i=1}^{\Gamma} s(i) \left(1 - e^{-\hat{\sigma}(\mathbf{r})_i \delta_i}\right) \hat{\mathbf{c}}(\mathbf{r})_i, \quad (1)$$

where $\hat{\mathbf{c}}(\mathbf{r})_i$ is the colour and $\hat{\sigma}(\mathbf{r})_i$ is the density predicted by the network at i . $\delta_i = t_{i+1} - t_i$ is the distance between adjacent sampled 3D spatial locations, and $s(i)$ is the inverse of the volume density that is accumulated up to the i^{th} spatial location, which is in turn computed as

$$s(i) = e^{-\sum_{j=1}^{i-1} \hat{\sigma}(\mathbf{r})_j \delta_j}, \quad (2)$$

where $(1 - e^{-\hat{\sigma}(\mathbf{r})_i \delta_i})$ is a density-based weight component: the higher the density value σ of a point, the larger the contribution on the final rendered colour.

Similarly to Eq. 1, we can render the pixel depth as

$$\mathbf{d}(\mathbf{r}) = \sum_{i=1}^{\Gamma} s(i) \left(1 - e^{-\hat{\sigma}(\mathbf{r})_i \delta_i}\right) \mathbf{z}_i, \quad (3)$$

where \mathbf{z}_i is the distance of the i^{th} spatial location with respect to the camera optic centre.

The input required to learn the NeRF parameters is a set of N images and their corresponding camera information. Let $\mathcal{I} = \{\mathbf{I}_k\}_{k=1}^N$ be the training images, and $\mathcal{P} = \{\mathbf{P}_k\}_{k=1}^N$ and $\mathcal{K} = \{\mathbf{K}_k\}_{k=1}^N$ be their corresponding camera poses and intrinsic parameters, respectively. A pose $\mathbf{P} = [\mathbf{R}, \mathbf{t}]$ is composed of rotation \mathbf{R} and translation \mathbf{t} . We can estimate the depth map of a given view k by rendering the depth of all its pixels, therefore we can define the estimated depth maps as $\mathcal{D} = \{\mathbf{D}_k\}_{k=1}^N$.

Learning f_θ is achieved by comparing each ground-truth pixel $\mathbf{c}(\mathbf{r})$ with its predicted counterpart $\hat{\mathbf{c}}(\mathbf{r})$. The goal is to minimise the following L2-norm objective function

$$\mathcal{L} = \frac{1}{|\mathcal{R}|} \sum_{\mathbf{r} \in \mathcal{R}} (\|\mathbf{c}(\mathbf{r}) - \hat{\mathbf{c}}_c(\mathbf{r})\|_2^2 + \|\mathbf{c}(\mathbf{r}) - \hat{\mathbf{c}}_f(\mathbf{r})\|_2^2), \quad (4)$$

where $\hat{\mathbf{c}}_c(\mathbf{r})$ and $\hat{\mathbf{c}}_f(\mathbf{r})$ are the coarse and fine predicted volume colours for ray \mathbf{r} , respectively. Please refer to [17] for more details.

3.2 View Morphing Overview

View morphing objective is to synthesise natural 2D transitions between an image pair $\{\mathbf{I}_k, \mathbf{I}_{k'}\}$ and the approach can be summarised in three steps: *i*) the two images are *prewarped* through rectification, i.e. their image planes are aligned without changing their cameras' optic centres; *ii*) the *morph* is computed between these prewarped images to generate a morphed image whose viewpoint lies on the line connecting the optic centres; *iii*) the image plane of the morphed image is transformed to a desired viewpoint through *postwarping*.

In practice, assuming the two views are prewarped, the morph uses the knowledge of their camera poses $\mathbf{P}_k, \mathbf{P}_{k'}$, and the pixel correspondences between the images, i.e. $q_k : \mathbf{I}_k \Rightarrow \mathbf{I}_{k'}, q_{k'} : \mathbf{I}_{k'} \Rightarrow \mathbf{I}_k$ where q_k is a function that maps a pixel of \mathbf{I}_k to the corresponding pixel in $\mathbf{I}_{k'}$ [24]. Sparse pixel correspondences can be defined by a user or determined by a keypoint detector, they can then be densified via interpolation to create a dense correspondence map. *This procedure is not viable as is in a learning-based pipeline, hence we have to define a novel view morphing strategy for a NeRF-based network architecture.* A warp function for each image can be computed from the correspondence map through linear interpolation

$$\begin{aligned} \hat{\mathbf{I}}_{k,\alpha} &= (1 - \alpha)\hat{\mathbf{I}}_k + \alpha q_k(\hat{\mathbf{I}}_k) \\ \hat{\mathbf{I}}_{k',\alpha} &= (1 - \alpha)q_{k'}(\hat{\mathbf{I}}_{k'}) + \alpha \hat{\mathbf{I}}_{k'} \\ \mathbf{P}_\alpha &= (1 - \alpha)\mathbf{P}_0 + \alpha \mathbf{P}_1, \end{aligned} \quad (5)$$

where $\hat{\mathbf{I}}_k$ are the coordinates of the image of camera k , $\hat{\mathbf{I}}_{k,\alpha}$ are pixel coordinates of the morphed image, and $\alpha \in [0, 1]$ regulates the position of the morphed view along the line connecting the two views. The morphed image can then be computed by averaging the pixel colours of the warped images. Please refer to [24] for more details.

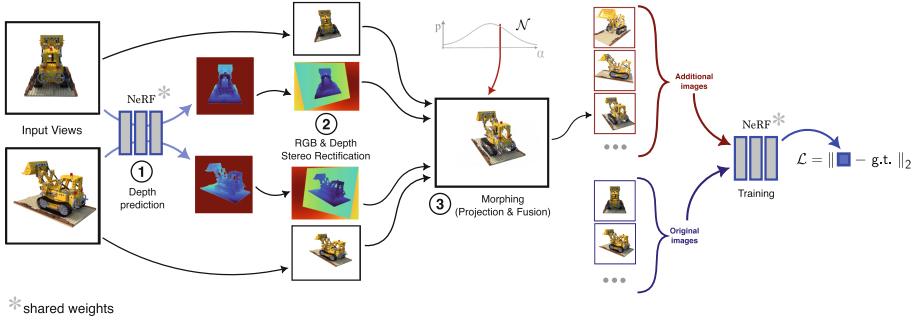


Fig. 2. Block diagram of NeRF-based View Morphing (VM-NeRF). From the left, we (1) predict the depth with NeRF, (2) rectify the input images and predicted depths, and (3) compute the image morphing of a view randomly positioned between the view pair. α determines the new view position and it is sampled from a Gaussian distribution.

4 NeRF-Based View Morphing

The goal of NeRF-based View Morphing (VM-NeRF) is to use the geometrical constraints of the morphing technique to synthesise a set of additional training input views $\mathcal{M} = \{\mathbf{M}_{(k,k'),\alpha}\}$, where $\mathbf{M}_{(k,k'),\alpha}$ is a morphed view generated from the view pair k and k' with a given value of α . Adapting view morphing in a learning-based pipeline is challenging as we need reliable pixel correspondences (q_k and $q_{k'}$) to synthesise morphed views. Our intuition is that it is possible to compute one-to-one correspondences from the disparity information, a function of the depth as in Eq. 3, which we can render with the very same NeRF model. We can then linearly interpolate the photometric content of the view pair to produce the morphed view.

Based on the description in Sect. 3.2, we integrate in NeRF only the steps of prewarping and morphing. We experimentally found that postwarping does not lead to better results. Sect. 4.1 describes how we perform the initial rectification of the two cameras. Section 4.2 describes how the images morphing is computed. Section 4.3 provides detailed information on our practical approach to training NeRF with View Morphing. Figure 2 shows the block diagram of our approach.

4.1 Rectification

Our first step is rectification, which leads to rotating the known camera poses \mathbf{P}_k and $\mathbf{P}_{k'}$ around their optic centres until their image planes become coplanar. We can then compute the common image plane by using a selection of algorithms such as [7, 9]. We represent this plane as the rotation matrix

$$\tilde{\mathbf{R}} = [\mathbf{a}_x, \mathbf{a}_y, \mathbf{a}_z], \quad (6)$$

where $\mathbf{a}_x, \mathbf{a}_y, \mathbf{a}_z$ are the axis components of the coplanar plane resulting from the rectification. Stereo rectification is applied to the original images $\{\mathbf{I}_k, \mathbf{I}_{k'}\}$

and depth maps $\{\mathbf{D}_k, \mathbf{D}_{k'}\}$ predicted in Eq. 3. The new camera pose of view k is equal to $\hat{\mathbf{P}}_k = [\hat{\mathbf{R}}, \mathbf{t}_k]$, where \mathbf{t}_k is the translation of the original camera pose \mathbf{P}_k (same applies to view k').

Rectification algorithms are typically based on the assumptions that viewpoints are aligned horizontally and that the reference viewpoint is the left-hand side of the camera (from an observer positioned behind the cameras) [7, 9]. This is atypical in NeRF, as viewpoints may have arbitrary camera configurations, leading to errors that should be corrected. We mitigate this problem by comparing \mathbf{a}_z with the z component of the original view pose. If this angle is greater than 45° with respect to both \mathbf{P}_k and $\mathbf{P}_{k'}$, we rotate the warping matrices and poses by 90° or 180° . The application of this modification to conventional rectification algorithm allows us to correctly generate the following rectified images $\{\tilde{\mathbf{I}}_k, \tilde{\mathbf{I}}_{k'}\}$ and rectified depth maps $\{\tilde{\mathbf{D}}_k, \tilde{\mathbf{D}}_{k'}\}$.

4.2 Image Morphing

The second step is image morphing, i.e. fusing the rectified images to obtain the new morphed image. This procedure is divided in three steps: *i*) finding the pixel correspondences; *ii*) computing the position of each pixel on the morphed camera; *iii*) fusing pixels that fall in the same position. To determine the image correspondences, we initially compute the disparity maps as functions of the rectified estimated depths

$$\mathbf{E}_k = \frac{f_k}{\tilde{\mathbf{D}}_k} \|\mathbf{o}_k - \mathbf{o}_{k'}\|_2, \quad \mathbf{E}_{k'} = \frac{f_{k'}}{\tilde{\mathbf{D}}_{k'}} \|\mathbf{o}_k - \mathbf{o}_{k'}\|_2, \quad (7)$$

where $\{\mathbf{o}_k, \mathbf{o}_{k'}\}$ are the principal points and $\{f_k, f_{k'}\}$ are the focal lengths of cameras k and k' .

Then, we determine the correspondences of the pixel positions between images defined in Eq. 5 as

$$q_k(\hat{\mathbf{I}}_k) = \hat{\mathbf{I}}_k + \frac{\tilde{\mathbf{b}}_k}{\|\tilde{\mathbf{b}}_k\|_2} \mathbf{1}^\top \odot \mathbf{E}_k, \quad (8)$$

where $\mathbf{1}$ is a vector of ones, \odot indicates the Hadamard product and $\hat{\mathbf{I}}_k$ is the baseline direction with respect to the common plane defined in Eq. 6 that is computed as

$$\tilde{\mathbf{b}}_k = \mathbf{a}_z \times ((\mathbf{o}_k - \mathbf{o}_{k'}) \times \mathbf{a}_z). \quad (9)$$

The same operation is computed for k' . Then, we apply the warp functions of Eq. 5 to compute the position of each pixel on the morphed view, thus obtaining $\hat{\mathbf{I}}_{k,\alpha}$ and $\hat{\mathbf{I}}_{k',\alpha}$.

Lastly, a coalescence operation [3] fuses the pixels of the two views k and k' . The coalescence operation concatenates two sets of coordinates and fuse pixels with the same position, preserving only the pixel values of the points that are nearest to the camera. We use $\{\tilde{\mathbf{D}}_k, \tilde{\mathbf{D}}_{k'}\}$ to determine the distance of the points.

Table 1. Results on the NeRF realistic synthetic 360° dataset.

# views	Method	PSNR ↑	SSIM ↑	LPIPS ↓
100 [17]	NeRF [17]	31.21	0.9513	0.0465
8	NeRF [17]	23.45	0.8673	0.1303
	DietNeRF [12]	22.98	0.8545	0.1258
	AugNeRF [5]	10.04	0.5415	0.3866
	RegNeRF [19]	22.91	0.8756	0.1138
	VM-NeRF	24.39	0.8768	0.1146
8 [12]	NeRF [12]	20.09	0.8220	0.1790
	DietNeRF [12]	23.59	0.8740	0.0970
	VM-NeRF	24.14	0.8729	0.1180
4	NeRF [17]	10.98	0.6550	0.3620
	DietNeRF [12]	12.61	0.6591	0.3302
	AugNeRF [5]	8.14	0.3924	0.4802
	RegNeRF [19]	15.88	0.7932	0.1994
	VM-NeRF	16.90	0.7563	0.2461

Table 2. Ablation study results. Keys: # views: represent the number of views and the relative subset. avg. dist.: average distance between view pairs.

# views	avg. dist.	Method	PSNR ↑	SSIM ↑	LPIPS ↓
8 [12]	5.20	DietNeRF [12]	25.59	0.9120	0.0770
8 [12]	5.20	VM-NeRF	26.90	0.9180	0.0797
8 (s1)	5.18	VM-NeRF	27.87	0.9294	0.0693
8 (s2)	4.38	VM-NeRF	27.13	0.9178	0.0855
8 (s3)	4.73	VM-NeRF	27.48	0.9317	0.0676
8 (s4)	4.69	VM-NeRF	28.39	0.9365	0.0598

4.3 Training with VM-NeRF

VM-NeRF is subject to the same geometric constraints as the original view morphing technique [24]. These constraints impose that singular camera configurations should not exist. These configurations happen whenever the optic centre of a camera is within the field of view of another one [24]. We also discard cameras that are distant from each other more than a threshold γ , as the morphed cameras may be on a transition path that crosses regions where the object of interest is not actually visible (so being rather useless for training a NeRF based model).

Because view morphing allows the synthesis of a new view at any point on the line that connects the known camera pair, we randomly sample new views using a Gaussian distribution centred halfway through the camera pair. Specifically, let us consider a normalised distance between the two cameras. The Gaussian distribution is centred at 0.5 and the standard deviation σ is chosen such that $3\sigma \rightarrow \epsilon$ at the optic centre positions. Therefore, we sample $\alpha \sim \mathcal{N}(0.5, \sigma)$ with $0 \leq \alpha \leq 1$. The depth NeRF can render at the first few iterations is noisy, therefore, we let NeRF warm up on the known views for λ iterations before synthesising and injecting VM-NeRF views in the next training iterations. After the warm-up, for each valid camera pair, we regenerate M new views every η training iterations as the predicted depth improves over time during training.

5 Experiments

5.1 Experimental Setup

We evaluate our method on three training setups using the NeRF realistic synthetic 360° dataset [17], which is composed of eight scenes, i.e. Chair, Drums,

Ficus, Lego, Materials, Ship, Mic, Hot Dog. **First setup:** We select $N = 8$ views out of 100 available for each scene using the Farthest Point Sampling (FPS) [20] (the first view is used for FPS initialisation in each scene). **Second setup:** we use the same $N = 8$ views used in DietNeRF [12]. **Third setup:** we select $N = 4$ views using the previous FPS approach. We test each trained model on all the test views of NeRF realistic synthetic 360°. We quantify the rendering results using the peak signal-to-noise ratio (PSNR) score, the structured similarity index measure (SSIM) [28] and the learned perceptual image patch similarity (LPIPS) [30]. We quantitatively compare our approach against DietNeRF [12] and RegNeRF [19] as the most recent methods for few-shot view synthesis. We also compare against AugNeRF [5] because it is the only data augmentation for NeRF, and data augmentation can be a useful strategy to promote generalisation. We choose to use the Chair scene for our ablation study, which consists of testing VM-NeRF on four different, randomly-chosen, configurations of eight views and on the DietNeRF configuration.

We implement NeRF and our approach in PyTorch Lightning, and run experiments on a single Nvidia A40 with a batch size of 1024 rays. A single scene can be trained in about two days. We use the original implementations of DietNeRF, AugNeRF and RegNeRF to evaluate the different setups. We set the same training parameters as in [17], and set $\gamma = 6$, $\sigma = 0.2$, $M = 1$, $\eta = 5$, $\lambda = 500$.

5.2 Analysis of the Results

Quantitative. Table 1 shows the results averaged over the eight scenes. Our NeRF implementation can achieve nearly the same results reported in [17] on the 100-view setup, i.e. PSNR equal to 31.21 (ours) compared to 31.01 [17].

VM-NeRF can outperform all the other methods in the eight-view setting. Interestingly, the original version of NeRF is the one that performs as second best, followed by DietNeRF and RegNeRF. AugNeRF fails to produce satisfactory results. We can also observe that VM-NeRF achieves slightly better quality than its version with oracle depth maps, i.e. 24.39 vs. 24.22 PSNR. In fact, we observed that VM-NeRF can effectively leverage the depth information that is estimated during training, although it is noisy. We also evaluate VM-NeRF on the same eight views originally tested by DietNeRF [12]. Also here we can achieve higher quality results on average, i.e. 24.14 vs. 23.59. We also improve in the four-view setup where we obtain an improvement of +1.02 PSNR on average. The results also show that the perturbation of the known input views, done by AugNeRF, has adverse effects in all the tested setups.

Qualitative. Figure 3 shows some qualitative results on Chair, Hot Dog and Lego where we can observe that VM-NeRF produce results with better details than DietNeRF. We speculate that this difference with DietNeRF may be due to its CLIP-based approach that is introduced to leverage a semantic consistency loss for regularisation [21]. The CLIP output is a low-dimensional (global) representation vector of the image, which may hinder the learning of high-definition details. Differently, our approach interpolates the original photometric infor-

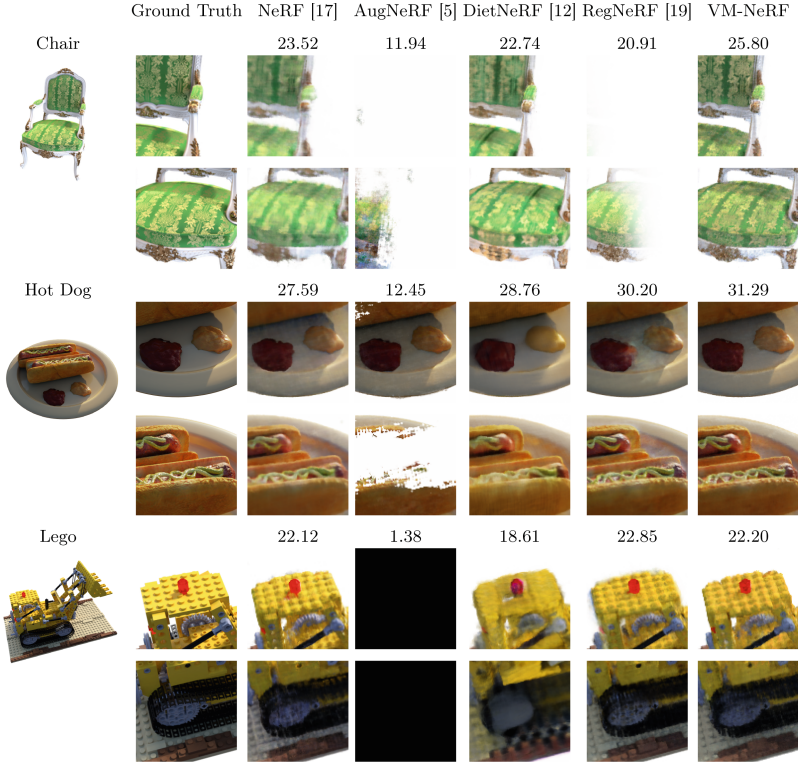


Fig. 3. Comparisons on test-set views of scenes of NeRF realistic synthetic 360°. Unlike AugNeRF [5], VM-NeRF is an effective method that can be used for few-shot view synthesis problems. Unlike DietNeRF [12], VM-NeRF enables NeRF to learn scenes with a higher definition. VM-NeRF produce less artefacts than RegNeRF during rendering [19]. We report the PSNR that we measured for each method and for each rendered image. AugNeRF unsuccessfully learns Chair and Lego (white and black outputs).

mation from two views to produce a new input view, without losing information through the encoding of the low-dimensional representation vector. Figure 3 shows that our approach compared to RegNeRF produces fewer artefacts by correlating the nearby views.

Ablation Study. We assess the stability of VM-NeRF by evaluating the rendering quality when different combinations of views are used to train NeRF. Table 2 shows that the performance is fairly stable throughout different view configurations. We also observed that the algorithm is robust to variations in the distance between view pairs. As long as a view pair is not singular and the distance between cameras is adequate to create acceptable 3D projective transformations of the object, we can successfully synthesise new views with VM-NeRF.

6 Conclusions

We presented a novel method for few-shot view synthesis that blends NeRF and the View Morphing technique [24]. View morphing requires no prior knowledge of the 3D shape and it is based on general principles of projective geometry. We evaluated our approach using the conventional dataset employed by NeRF-based methods, demonstrating that VM-NeRF more effectively learns 3D scenes across various few-shot view synthesis setups. VM-NeRF can interpolate only along the line that connects the optical centres of each camera pair. Therefore, it cannot reconstruct the whole object if only a part of it is viewed during training. Lastly, we designed our approach to be fully differentiable, so an attractive research direction is to integrate our approach into an end-to-end training pipeline.

Acknowledgements. This work was supported by the PNRR project FAIR - Future AI Research (PE00000013), under the NRRP MUR program funded by the NextGenerationEU. This research is partially supported by the project Future Artificial Intelligence Research (FAIR) - PNRR MUR Cod. PE0000013 - CUP: E63C2200194 0006 and the framework “RAISE - Robotics and AI for Socio-economic Empowerment” supported by European Union - NextGenerationEU.

References

1. Aanæs, H., Jensen, R.R., Vogiatzis, G., Tola, E., Dahl, A.B.: Large-scale data for multiple-view stereopsis. *Int. J. Comput. Vision* **120**(2), 153–168 (2016)
2. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-NeRF: a multiscale representation for anti-aliasing neural radiance fields. In: *ICCV* (2021)
3. Chaitin, G.J.: Register allocation & spilling via graph coloring. *ACM Sigplan Not.* **17**(6), 98–101 (1982)
4. Chen, D., Liu, Y., Huang, L., Wang, B., Pan, P.: GeoAug: data augmentation for few-shot NeRF with geometry constraints. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *ECCV 2022*. LNCS, vol. 13677, pp. 322–337. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19790-1_20
5. Chen, T., Wang, P., Fan, Z., Wang, Z.: Aug-NeRF: training stronger neural radiance fields with triple-level physically-grounded augmentations. In: *CVPR* (2022)
6. Devernay, F., Peon, A.R.: Novel view synthesis for stereoscopic cinema: detecting and removing artifacts. In: *Workshop on 3D Video Processing (ACMMM)* (2010)
7. Fusiello, A., Trucco, E., Verri, A.: A compact algorithm for rectification of stereo pairs. *Mach. Vis. Appl.* **12**(1), 16–22 (2000)
8. Gallo, O., Troccoli, A., Jampani, V.: Novel View Synthesis: From Depth-Based Warping to Multi-Plane Images and Beyond (2020). <https://nvlabs.github.io/nvs-tutorial-cvpr2020/>. *Conference on Computer Vision and Pattern Recognition*
9. Hartley, R.I., Zisserman, A.: *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge (2004)
10. Hedman, P., Srinivasan, P.P., Mildenhall, B., Barron, J.T., Debevec, P.: Baking neural radiance fields for real-time view synthesis. In: *ICCV* (2021)
11. Ichnowski, J., Avigal, Y., Kerr, J., Goldberg, K.: Dex-NeRF: using a neural radiance field to grasp transparent objects. In: *CRL* (2022)

12. Jain, A., Tancik, M., Abbeel, P.: Putting NeRF on a diet: semantically consistent few-shot view synthesis. In: ICCV (2021)
13. Kajiyaya, J., Herzen, B.: Ray tracing volume densities. In: SIGGRAPH (1984)
14. Long, X., Lin, C., Wang, P., Komura, T., Wang, W.: SparseNeuS: fast generalizable neural surface reconstruction from sparse views. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) ECCV 2022. LNCS, vol. 13692, pp. 210–227. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-19824-3_13
15. Max, N.: Optical models for direct volume rendering. *IEEE Trans. Vis. Comput. Graph.* **1**(2), 99–108 (1995)
16. Mildenhall, B., et al.: Local light field fusion: practical view synthesis with prescriptive sampling guidelines. *ACM Trans. Graph.* **38**(29), 1–14 (2019)
17. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
18. Müller, T., Rousselle, F., Novák, J., Keller, A.: Real-time neural radiance caching for path tracing. *ACM Trans. Graph.* **40**(4), 1–16 (2021)
19. Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S.M., Geiger, A., Radwan, N.: RegNeRF: regularizing neural radiance fields for view synthesis from sparse inputs. In: CVPR (2022)
20. Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3D object detection in point clouds. In: ICCV (2019)
21. Radford, A., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
22. Rematas, K., Martin-Brualla, R., Ferrari, V.: ShaRF: shape-conditioned radiance fields from a single view. In: ICML (2021)
23. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR (2016)
24. Seitz, S.M., Dyer, C.R.: View morphing. In: Conference on Computer Graphics and Interactive Techniques (1996)
25. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. *J. Big Data* **6**(1), 1–48 (2019)
26. Tewari, A., et al.: Advances in neural rendering. In: Computer Graphics Forum, vol. 41, no. 2, pp. 703–735 (2022)
27. Wang, J., et al.: Generalizing to unseen domains: a survey on domain generalization. *IEEE Trans. Knowl. Data Eng.* **35**(08), 8052–8072 (2023)
28. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* **13**(4), 600–612 (2004)
29. Xie, Y., et al.: Neural fields in visual computing and beyond. In: Computer Graphics Forum, vol. 41, no. 2, pp. 641–676 (2022)
30. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)