Video Extrapolation in Space and Time

Yunzhi Zhang and Jiajun Wu

Stanford University {yzzhang,jiajunwu}@cs.stanford.edu

Abstract. Novel view synthesis (NVS) and video prediction (VP) are typically considered disjoint tasks in computer vision. However, they can both be seen as ways to observe the spatial-temporal world: NVS aims to synthesize a scene from a new point of view, while VP aims to see a scene from a new point of time. These two tasks provide complementary signals to obtain a scene representation, as viewpoint changes from spatial observations inform depth, and temporal observations inform the motion of cameras and individual objects. Inspired by these observations, we propose to study the problem of Video Extrapolation in Space and Time (VEST). We propose a model that leverages the self-supervision and the complementary cues from both tasks, while existing methods can only solve one of them. Experiments show that our method achieves performance better than or comparable to several state-of-the-art NVS and VP methods on indoor and outdoor real-world datasets. *

1 Introduction

Novel view synthesis (NVS) and video prediction (VP) are both widely studied computer vision tasks. The former extrapolates a scene to a different camera viewpoint, while the latter extrapolates to a future timestamp. NVS focuses on the scene geometry revealed from discrete camera positions, whereas VP extracts information from the moving trajectories of both cameras and objects. The self-supervision signals from these two tasks can be jointly used to extract a scene representation. To this end, we call attention to the problem of Video Extrapolation in Space and Time (VEST) that considers both tasks.

We solve the problem of VEST by first developing a representation that incorporates both the spatial and temporal consistency from video data as inductive biases. We generalize Multiplane Images (MPIs) [44], which is originally a layered representation that decomposes images into RGBA planes, by additionally parameterizing the flow field of each plane in order to model the temporal dynamics. Images from a future timestamp and a novel viewpoint can be rendered from flow-based and homography-based warping, respectively.

Compared with previous MPI-based NVS approaches [4,29,44,32], our generalized MPI representation leverages learning signals derived from both the

^{*} Project page: https://cs.stanford.edu/~yzzhang/projects/vest/.



Fig. 1: We propose the task of Video Extrapolation in Space and Time (VEST) that exploits both spatial and temporal consistency in video data. (a) During inference time, the model takes in two consecutive frames as inputs, colored in orange, and outputs images extrapolated along the temporal axis for future frame prediction and along the spatial axis for novel view synthesis, colored in green. The black curve denotes the camera trajectory. (b) and (c) are examples of model inputs; (d) and (e) are model outputs. Images from the bottom row are zoom-in views.

spatial and temporal coherence in video data, while previous NVS methods utilize frame tuples randomly sampled from the training video sequences without considering the temporal information. Compared with other optical-flow-based VP methods [5], our method predicts the motion field individually for each MPI plane instead of the full scene. Since each MPI plane captures a relatively simple structure, we can effectively estimate the motion field of each plane with affine transformations [35].

We instantiate a model that performs the spatial-temporal extrapolation with the generalized MPI representation. As shown in Fig. 1, our model predicts MPI planes from monocular inputs and leverages historical frames for motion inference. Experiments show that our problem formulation and model are generally applicable to diverse scenarios: indoor and outdoor scenes, and videos taken by a single and multiple static or moving camera(s). Our method achieves favorable performance on *both* tasks compared to baselines designed for either.

Our main contributions are as follows:

- We propose VEST as a self-supervision task for learning a generalized MPI representation from video data.
- We instantiate a model that learns the proposed representation for simultaneous video extrapolation in space and time.
- Our VEST model produces realistic results in both space and time extrapolation on a diverse range of datasets.

2 Related Works

Novel view synthesis. Synthesizing novel views based on 2D images is a challenging problem, as it requires the reasoning of the 3D structure of the perceptual world. As shown in Fig. 2(a) and Fig. 2(b), monocular view synthesis approaches [32,38,10] take I_{11} as inputs and outputs the synthesized image for a



Fig. 2: Overview of vision tasks related to our work. I_{ij} denotes the RGB frame for the *i*-th timestamp in a video sequence taken from the *j*-th camera viewpoint. We restrict $i, j \in \{1, 2, 3\}$ for illustration. In (a) and (b), monocular or binocular NVS methods only extrapolate in space, while in (c), VP methods only extrapolate in time. Video-based NeRF methods in (d) train on a collection of images for each scene, and perform interpolation instead of extrapolation for inference. Finally, (e) illustrates the problem we focus on in this work. Given historical frames, the task consists of both predicting the future and extrapolating to a novel viewpoint. During inference, novel views I_{21}, I_{23} can be synthesized the same way as training time, while I_{22}, I_{23} can be inferred by duplicating I_{11} as input (VEST-single from 4.1).

query view $j \neq 1$. Stereo view synthesis approaches [44] are similar but take in two input views I_{11}, I_{12} . Most of these approaches synthesize novel views based on camera parameters estimated from Structure-from-Motion (SfM) techniques, except for Lai et al. [10] which uses camera parameters predicted by the model. These methods focus exclusively on spatial extrapolation and do not take the temporal axis into account. Previous works [45,44,27,33] use layered representation and apply per-layer warping to obtain the novel view based on camera poses. SynSin [38] and Worldsheet [8] predict the depth from a monocular image input and warps a 3D representation of the scene in the features space and pixel space, respectively, to synthesize novel views. In comparison, our proposed VEST encapsulates NVS but also has the ability to perform temporal extrapolation.

Closer to ours, Lin et al. [13] and Yoon et al. [41] tackle the problem of NVS for dynamic scenes. Lin et al. [13] addresses the temporal inconsistency of the MPI representation when applied to scenes with moving objects by identifying the error-prone regions with a learned 3D mask volume. There are three key differences between our model and Lin et al. First, they require two images from synchronized stereo cameras as inputs during inference, while our method takes in two consecutive frames from a single camera. Second, they assume a static camera and static background while we do not. Third, they rely on the static background computed from the full video sequences from two source views, while ours directly predicts the future video sequence based on two input frames. Yoon et al. [41] proposes to fuse the depth estimated from a single view and multiple views with a depth fusion network, and predicts the novel view on top. It requires segmentation labels and optical flow inputs, while our method is fully self-supervised.

4 Y. Zhang and J. Wu

Video prediction. Video prediction methods typically take in two or more frames taken by one camera and only extrapolate along the temporal axis as shown in Fig. 2(c). Deep neural networks have been widely used in video prediction [22,30,26,5]. Previous works studying temporal extrapolation for videos include hallucination-based methods [16,34] and warping-based methods [15]. More recently, Wang et al. [36] propose PredRNN-V2, integrating convolutional recurrent units with a pair of decoupled memory cells. Several methods also proposed to decompose high-dimensional videos into object-centric [39] or semanticaware [1] regions, and model the motion of each region independently. These methods consider only monocular input and often require models pre-trained on large datasets for semantic and depth information, while our method factorizes videos into depth-aware components by using the geometric cue from view synthesis, and it requires no external supervision.

Dynamic neural radiance fields. Neural Radiance Fields (NeRF) [19] shows impressive results in synthesizing novel views with high fidelity. Not surprisingly, many follow-up papers have attempted to extend it to video for image synthesis in novel space and time point [3]. Our method differs from these methods in two ways. First, though on static scenes, methods like PixelNeRF can already learn a neural scene representation conditioned on one or few input images [42,20,40,12,21,31], all video-based NeRF methods still employ single-video learning and need to be retrained for every new scene, while our method performs 4D synthesis conditioned on the input images. Second, NeRF-based methods require hundreds of images for training, and essentially perform interpolation among input time- and viewpoints as shown in Fig. 2(d), while our method focuses on video extrapolation.

Layered representations. The idea of decomposing images into RGBA layers is effective in multiple problem domains. For view synthesis, Shade et al. [25] proposed layered depth images to represent a scene with multi-layer depth and color images. Viewers can then see the scene from different points in space. For video prediction, there has also been a line of work inspired by the classic research on layered motion representations [35]. A notable example is a recent work by Lu et al. [17,18] on the problem of video manipulation, where they decompose videos into semantic-aware RGBA layers. Our layered representation builds upon all these ideas and aims to tackle both NVS and VP.

3 Method

3.1 Multiplane Images

Before introducing the representation we use, we first review the classical Multiplane Images (MPIs) [44]. An image $I \in \mathbb{R}^{H \times W \times 3}$ is represented by D RGBA planes, $\{(c_i, \alpha_i)\}_{i=1}^{D}$, where $c_i \in \mathbb{R}^{H \times W \times 3}$ are RGB values and $\alpha_i \in \mathbb{R}^{H \times W \times 1}$ are alpha values. Each plane corresponds to fixed depth d_i . The plane is frontoparallel to the camera and can be written as $\mathbf{n}^T \mathbf{x} - d_i = 0$, where $\mathbf{n} = [0, 0, 1]^T$ is the normal vector of the plane. Let (R, t) be the rotation and translation matrix from target to source view, and K, K' be the camera intrinsics for source and target views. The transformation for the *i*-th plane from target to source view, denoted as $\mathcal{W}_{i}^{R,t,K,K'}$, is defined as

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} \sim \mathcal{W}_i^{R, \boldsymbol{t}, K, K'} \begin{bmatrix} u' \\ v' \\ 1 \end{bmatrix} := K \left(R - \frac{\boldsymbol{t} \boldsymbol{n}^T}{d_i} \right) (K')^{-1} \begin{bmatrix} u' \\ v' \\ 1 \end{bmatrix}, \qquad (1)$$

where (u, v) and (u', v') are the coordinates from the source and target views, respectively.

The MPI representation for the target view $\{(c'_i, \alpha'_i)\}_{i=1}^D$ is computed as

$$c'_{u',v'} = c_{u,v},$$
 (2)

$$\alpha'_{u',v'} = \alpha_{u,v}.\tag{3}$$

Here (u, v) are sampled according to Equation (1). Similar to Zhou et al. [44], we apply bilinear sampling from the neighboring grid corners when (u, v) is not aligned with the coordinate grid.

Finally, an MPI renderer synthesizes the target-view image with

$$\hat{I}' = \sum_{i=1}^{D} c'_i \alpha'_i \prod_{j=i+1}^{D} (1 - \alpha'_j).$$
(4)

3.2 Generalized Multiplane Images

We generalize the MPI representation to model the motion from image $I = I_t$ to $I' = I_{t+1}$, the next frame in a video sequence. Note that different from Section 3.1, here we use I' to refer to the frame for timestamp t + 1 which is not necessarily from a different camera viewpoint.

Formally, the generalized MPI representation for image I is denoted as $\{(c_i, \alpha_i, \mathcal{T}_i)\}_{i=1}^{D}$, where $\mathcal{T} = \mathcal{T}_i$ is an operator defined by

$$\begin{bmatrix} u \\ v \end{bmatrix} \sim \mathcal{T} \begin{bmatrix} u' \\ v' \end{bmatrix} := \begin{bmatrix} u' + \Delta u' \\ v' + \Delta v' \end{bmatrix}.$$
 (5)

Let θ be a parameterization for \mathcal{T} . When \mathcal{T} is fully parameterized, $\theta \in \mathbb{R}^{H \times W \times 2}$ is the pixel displacement field, and $\Delta u' = \theta_{u',v',0}$, $\Delta v' = \theta_{u',v',1}$ correspond to the number of pixels to be shifted in the u, v coordinates, respectively. However, since each MPI plane has a relatively simple structure, in practice we restrict \mathcal{T} to be in the class of affine transformations. \mathcal{T} can now be parameterized by $\theta \in \mathbb{R}^{2 \times 3}$, and

$$\mathcal{T}\begin{bmatrix}u'\\v'\end{bmatrix} = \begin{bmatrix}\theta_{1,1} & \theta_{1,2} & \theta_{1,3}\\\theta_{2,1} & \theta_{2,2} & \theta_{2,3}\end{bmatrix}\begin{bmatrix}u'\\v'\\1\end{bmatrix}.$$
(6)

We denote the generalized MPI representation for image I as

$$\{(c_i, \alpha_i, \theta_i)\}_{i=1}^D.$$
(7)

Finally, we predict the next frame \hat{I}' based on Equation (2)-(4).



Fig. 3: Model architecture. The MPI encoder receives monocular frames I_{t-1} and I_t as inputs and outputs the generalized MPI representation for I_t , which consists of D RGBA planes and the motion parameters for each plane. Then the MPI renderer renders the target view image I'_t based on camera parameters and the RGBA planes. The next-frame prediction I_{t+1} is generated by first warping each RGBA plane with motion parameters, and then composing the planes. The training objective is to match \hat{I}'_t and \hat{I}_{t+1} with the ground truth. Images are licensed under CC BY-NC-SA 3.0.

3.3 Video Extrapolation in Space and Time

We now introduce our VEST model that predicts the generalized MPIs given monocular video frames, and performs spatial and temporal extrapolation. An overview of our model is shown in Fig. 3.

Training. The model takes in two consecutive frames from a video sequence, I_{t-1} and I_t , and outputs the generalized MPI representation (Equation (7)) for I_t . The target-view image I'_t is synthesized following Section 3.1, and the next-frame prediction following Section 3.2.

Inference. During inference, with inputs I_{t-1} and I_t , the model can be queried to extrapolate to other space-time coordinates. Future frames with longer horizon I_{t+2}, I_{t+3}, \cdots can be inferred by iteratively forwarding the model in an autoregressive manner.

Even when there is only one input frame I_t available, our model can synthesize the novel view I'_t by having I_t replicated twice as inputs and still produce realistic NVS results, corresponding to VEST-single from Table 1 and Table 2 as we will show later.

Losses. The training loss is the sum of spatial and temporal extrapolation errors, namely

$$\mathcal{L}^{\text{total}} = \mathcal{L}^{\text{space}}(\hat{I}'_t, I'_t) + \mathcal{L}^{\text{time}}(\hat{I}_{t+1}, I_{t+1}), \tag{8}$$

where

$$\mathcal{L}^{\text{space}} = \lambda_1^{\text{space}} \mathcal{L}_1 + \lambda_{\text{perc}}^{\text{space}} \mathcal{L}_{\text{perc}}, \qquad (9)$$

$$\mathcal{L}^{\text{time}} = \lambda_1^{\text{time}} \mathcal{L}_1 + \lambda_{\text{perc}}^{\text{time}} \mathcal{L}_{\text{perc}}.$$
 (10)

 \mathcal{L}_1 is the ℓ_1 loss, and \mathcal{L}_{perc} is the perceptual loss using pretrained VGG-19 [28] features.

3.4 Implementation Details

We adopt a model architecture similar to Tucker and Snavely [32], specified in Appendix A. The network outputs a tensor in $\mathbb{R}^{D \times H \times W \times 7}$, which is split channelwise into $f^{\alpha} \in \mathbb{R}^{D \times H \times W \times 1}$ for alpha values and $f^{\theta} \in \mathbb{R}^{D \times H \times W \times 6}$ for motion parameters. We set the RGB values for each plane (c_i from Equation (7)) to be the RGB values of the source image. The final motion parameter θ for each MPI plane is computed as a weighted spatial average of f^{θ} :

$$w_i = \alpha_i \prod_{j=i+1}^{D} (1 - \alpha_j), \tag{11}$$

$$\theta_i = \text{SpatialAverage}(w_i \otimes f^{\theta}),$$
(12)

where \otimes denotes the element-wise multiplication. Note that w_i are the same weights for RGB values used in Equation (4).

Camera parameters are estimated with SfM [23,24] and are used to render novel views. Since SfM models have ambiguous depth scales, we compute a depth scale factor σ similar to Tucker and Snavely [32] such that MPI planes are associated with scaled depth values $\{\sigma d_i\}_{i=1}^{D}$ instead of $\{d_i\}_{i=1}^{D}$. Here σ is computed to minimize the log-squared error of the predicted depth map \hat{Z} and a set of sparse 3D points P_s ,

$$\sigma = \exp\left[\frac{1}{\mid P_s \mid} \sum_{(u,v,z) \in P_s} (\ln z - \ln \hat{Z}_{u,v})\right],\tag{13}$$

where the depth map is computed with Equation (11) and

$$\hat{Z} = \sum_{i=1}^{D} d_i w_i. \tag{14}$$

For each element $(u, v, z) \in P_s$, (u, v) is the pixel coordinate in the source image, and z is the depth value of the corresponding 3D point. We run COLMAP [23] to obtain P_s for each video sequence.

4 Experiments

We conduct extensive experiments to validate that our method compares favorably to state-of-the-art methods on both monocular NVS and VP in diverse scenarios: during training, our model may learn from single- and multi-view videos, static and moving cameras, and indoor and outdoor scenes. In particular, our method is the only one that performs both tasks simultaneously. Finally in Section 4.5, we provide an analysis of intermediate model outputs for an intuitive understanding of our method.



Fig. 4: Training setups. Training inputs are colored in orange and outputs are in green. Black curves denote camera trajectories. Our method can be trained on multi-view or single-view datasets. I_{33} is the ground truth for future frame prediction. The NVS supervision comes from a stereo frame when available (I_{23} in (a)), or from a temporally nearby frame otherwise (I_{44} in (b)).

4.1 Learning from Multi-View Videos with Moving Cameras

Setup. We first test our model on learning from videos with binocular views, as shown in Fig. 4(a). KITTI [6] is a benchmark dataset widely used for both NVS and VP. It contains street scenes captured by two stereo cameras mounted on a moving car. Following prior methods [33,32,11], we use the 28 city scenes from the KITTI-raw dataset and split them into 20 sequences for training, 4 for validation, and 4 for testing. This is denoted as LDI [33] split in Table 1. To be directly comparable with prior works on VP [39,1], we also evaluate on the split with 24 sequences for training and 4 for testing, denoted as FVS [39] split.

Evaluation. For NVS evaluation, images are cropped by 5% from the boundary of all sides and resized to 128×384 . We report the similarity of synthesized images compared to the ground truth using LPIPS [43] computed with VGG features, SSIM [37] and PSNR. We report LPIPS with AlexNet features for VP to be directly comparable to prior methods [39,1].

Baselines. For NVS, we compare our results with Tucker and Snavely [32] and LDI [33]. LDI [33] derives a 2-layer representation from single-view inputs. We also compare with MINE [11], a recent work that extends Tucker and Snavely [32] to the continuous depth domain using implicit functions. MINE [11]-32 and MINE [11]-64 from Table 1 refer to two model variants with 32 and 64 planes, respectively, as reported in the original paper.

For VP, we compare our method with PredRNN-V2 [36], a recently proposed method using a pair of decoupled RNN memory cells. We retrain their non-actionconditioned model on the dataset with 2 input frames and a prediction length of 5. We also compare with two VP methods which decompose scenes for better dynamics modeling. FVS [39] predicts the next frame by decomposing frames in a video into object-centric layers and modeling the dynamics of each layer. Similar to ours, FVS also assumes the motion of each layer to be affine. Another baseline is SADM [1] which decomposes frames into semantically consistent regions and predicts the motion of each region. Both FVS and SADM require instance maps and optical flow inputs obtained from pretrained models.

Results. Results are shown in Table 1. All our models are trained with the same configuration. For NVS, our model achieves a significant improvement across all metrics compared to the NVS baselines including variants of MINE [11] that are pre-trained on ImageNet [2]. We further show a qualitative comparison in Fig. 5

Split	Method	Extrapo	plation i	n Space	Extrapolation in Time		
		LPIPS↓	$SSIM\uparrow$	$PSNR\uparrow$	$\rm LPIPS\text{-}AlexNet \downarrow$		
					t+1	t+3	t+5
LDI [33] split train 256×768 test 128×384	MINE [11]-32*	0.112	0.822	21.4		N/A	
	MINE [11]-64*	0.108	0.820	21.3		N/A	
	LDI [33]	N/A	0.572	16.5		N/A	
	Tucker et al. [32]	N/A	0.733	19.5		N/A	
	PredRNN-V2 [36]		N/A		0.3085	0.4573	0.5422
	VEST-no-perc	0.111	0.819	21.4	0.1129	0.2902	0.3943
	VEST-single	0.091	0.816	21.2		N/A	
	VEST (ours)	0.085	0.825	21.6	0.1154	0.2881	0.3911
LDI [33] split train 128×384 test 128×384	MINE [11]*	0.112	0.828	21.9		N/A	
	MINE [11]	0.129	0.812	21.4		N/A	
	PredRNN-V2 [36]		N/A		0.2153	0.3946	0.4984
	VEST-single	0.105	0.806	20.7		N/A	
	VEST (ours)	0.097	0.818	21.1	0.0798	0.2348	0.3384
FVS [39] split train 256×832 test 256×832	FVS [39]*		N/A		0.1848	0.2461	0.3049
	SADM [1]*		N/A		0.1441	0.2458	0.3116
	PredNet [16]		N/A		0.5535	0.5866	0.6295
	MCNet [34]		N/A		0.2405	0.3171	0.3739
	VoxelFlow [15]		N/A		0.3247	0.3743	0.4159
	VEST (ours)	0.150	0.739	19.9	0.1560	0.3441	0.4467

Table 1: Results on KITTI [6] with three train-test splits used in previous methods. Our method achieves better performance than all NVS baselines, including those pre-trained on ImageNet [2] (denoted by *). We also achieve competitive performance compared with VP baselines. LDI [33] and Tucker et al. [32] do not report LPIPS values, denoted by N/A. Baselines can do only one task while ours solves both, indicated by N/A.



Fig. 5: View synthesis results on KITTI. Left: (a) results from MINE [11]; (b) results from VEST (ours); (c) ground truth images. Right: zoom-in views. In the first two examples, our model produces fine details for small objects, such as windows and printed texts. The third example shows that our model performs more accurate extrapolation even for challenging, thin objects, such as trees with small distortion.

where MINE [11] uses ImageNet pretraining. Please refer to the project page for better visualization.

We additionally test our model with a different inference procedure (VESTsingle as shown in Table 1). Instead of I_{t-1} and I_t , the model receives a single source frame I_t repeated twice as inputs to predict I'_t . In this way, our model receives no additional information from historical frames. VEST and VEST-single



Fig. 6: Video prediction results on KITTI. Left: (a) predictions from PredRNN-V2 [36]; (b) predictions from VEST (ours); (c) ground truth images. Right: zoom-in views. Each row corresponds to frame prediction for t + 1, t + 3, t + 5. Our VEST produces much sharper predictions compared to the baseline.

from Table 1 use the same checkpoint and both outperform all non-pretrained baselines across all metrics.

For VP, our approach leads to a significant improvement compared to PredRNN-V2 [36]. Fig. 6 shows that our model makes prediction with sharper edges and higher fidelity. In these examples, street and cars are moving towards the camera, which is effectively a scale transformation in 2D. Because scale transformation is a subclass of affine transformation, our method can accurately capture such perpendicular motions in the scenes. All qualitative results from the KITTI dataset are used under license CC BY-NC-SA 3.0.

Ablation. We ablate the importance of using Perceptual loss for NVS. In Table 1, VEST-no-perc corresponds to a variant with $\lambda_{\text{perc}}^{\text{space}} = 0$, suggesting that perceptual loss in the spatial loss term largely improves performance. We also include an ablation study on the effect of the number of MPI planes in Section C.3.

4.2 Learning from Single-View Videos with a Moving Camera

Setup. We evaluate our method on video data with a single view, as shown in Fig. 4(b). RealEstate10K [44] is a standard NVS benchmark dataset consisting of 80K videos filming mostly indoor scenes. We follow the training-test split from Lai et al. [10], which uses a randomly sampled subset of the full dataset. It includes 10K video sequences for training and 5K for testing. The training and evaluation resolution is 256×256 .

Evaluation. Following MINE [11], for evaluation, we randomly sample 5 source frames from each testing sequence and sample target frames that are 5, 10, or at most 30 frames apart for each of the source frames. We evaluate on the intersection of the testing frames used in MINE and our testing split. This results in 372 testing sequences in total. The similarity scores are measured with LPIPS [43] with VGG features, SSIM [37], and PSNR.

Baselines. We run the publicly released checkpoint from MINE [11] with 64 planes, which is pre-trained on ImageNet and trained on the full training split of RealEstate10K with a resolution of 384×256 . We first run their model with

Method	LPIPS↓		$\rm SSIM\uparrow$		$\mathrm{PSNR}\uparrow$		LPIPS SSIM	PSNR
	n = 5	$n = \operatorname{rand}$	n = 5	$n = \operatorname{rand}$	n = 5	$n = \operatorname{rand}$	t + 1	
MINE [11]*	0.0986	0.1774	0.9018	0.8221	27.9837	24.3112	N/A	
SynSin [38]	N/A	N/A	N/A	0.7400	N/A	22.3100	N/A	
Tucker et al. [32]	0.0967	0.1761	0.8699	0.7851	27.0500	23.5200	N/A	
VEST-single	0.0944	0.1736	0.8700	0.7688	26.6599	24.6906	N/A	
Zhou et al. [44]	0.0816	0.1667	0.8943	0.8014	27.5788	24.1531	N/A	
VEST (ours)	0.0841	0.1596	0.8987	0.8003	28.2078	25.7553	$0.0436 \ \ 0.9334$	31.6831

Table 2: Results on RealEstate10K [44]. Our model achieves a better or comparable performance compared to the baselines, and only ours can predict future frames. SynSin [38] does not report LPIPS values, denoted by N/A. n specifies the number of frames between source and target images in the video sequence, and n = rand means that n is uniformly sampled between 1 and 30. A larger n indicates larger range of spatial extrapolation. * indicates that the model requires pretraining.



Fig. 7: Novel view synthesis results on RealEstate10K. (a)-(c) correspond to results from MINE [11], SynSin [38], and VEST (ours). (d) corresponds to the ground truth target image I_{22} . These examples show that VEST produces sharper details than MINE [11]. It also predicts object positions more accurately compared to SynSin [38].

resolution 384×256 , and then downsample to 256×256 to compare with the ground truth. Both our model and MINE receive the same set of sparse points for scale-invariant analysis in Equation (13). We also compare with Synsin [38] and Tucker and Snavely [32], both taking single-view inputs, and Zhou et al. [44], which takes in binocular inputs.

Results. We show quantitative results in Table 2. We obtain similar scores as reported in the original paper.^{**} Results for Tucker and Snavely [32] and

^{**} The LPIPS scores of MINE [11] computed are slightly worse compared to the original paper due to a bug in the evaluation script in their public codebase, where tensors in range [0, 1] are fed into an LPIPS package which expects inputs in range [-1, 1].

12 Y. Zhang and J. Wu



Fig. 8: Results on the ACID dataset filmed by a monocular moving camera. (a)-(c) are VP results and (d)-(f) NVS results.

Test-time Input	Method	LPIPS↓	$\mathrm{SSIM}\uparrow$	$\mathrm{PSNR}\uparrow$
Multi-view	Lin et al. [13]	0.1558	0.8667	21.1988
Single-view	Lin et al. [13] VEST-single (ours) VEST (ours)	0.3719 0.2600 0.2591	0.4929 0.6242 0.6249	16.7794 21.6957 21.7664

Table 3: Results on the dataset from Lin et al. [13] filmed by multi-view static cameras. Our method achieves better performance than the baseline method under the same view setting.

SynSin [38] are taken from the original paper for reference, and they are evaluated on a different test split. Our method outperforms MINE [11], a state-of-the-art single-view NVS method on this dataset, despite that the baseline uses ImageNet pre-training and trains on the full dataset. We show qualitative comparisons in Fig. 7.

On this dataset, modeling the motion is beneficial as our method uses the motion parallax from input frames by modeling the dynamics as opposed to explicit plane sweep volume construction as done in the stereo-based baselines [44,13]. Indeed, our method is comparable to the binocular baseline Zhou et al. [44].

Additionally, we train our model on Aerial Coastline Imagery Dataset (ACID) [14], a single-view dataset mostly filming outdoor natural scenes. Our method can synthesize images corresponding to the queried viewpoint and timestamp, as shown in Fig. 8. Both RealEstate10K and ACID datasets include YouTube videos under the Creative-Commons license.

4.3 Learning from Multi-View Videos with Static Cameras

Setup. We evaluate our method on learning from a multi-view dataset with static cameras. We use the dataset from Lin et al. [13], containing videos captured by 10 synchronized, static cameras, where scenes have a static background and human body movement in the foreground. The dataset is split into 86 scenes for training and 10 scenes for testing. The training and evaluation resolution is 256 \times 256. We follow Lin et al. [13] and pretrain our model on RealEstate10K.



Fig. 9: Results on dataset from Lin et al. [13]. (a)-(b) VP results; (c)-(d) NVS results. Our method assigns the moving human figure and the background scene into different MPI layers.



(g) RGBA plane (h) RGBA plane (i) RGBA plane (j) RGBA plane (k) Predicted flow (l) Predicted depth

Fig. 10: Video prediction results on the cloud dataset. (a)-(b) are two input frames, and (c)-(f) are predicted future frames. (g)-(j) are RGBA planes ordered from far to near, and (k)-(l) are the flow and depth maps predicted by the model. In this example, regions of clouds and the building are assigned to different planes, corresponding to different motions.

Results. We compare with Lin et al. [13], a state-of-the-art multi-view NVS method, which incorporates a learned 3D mask volume into the MPI representation to improve the temporal consistency of MPI planes.

As shown in Table 3, our model outperforms the baseline on novel view synthesis from a single view. Note that Lin et al. [13] was originally designed to perform view synthesis from stereo inputs and works well in their setup; in contrast, our model is designed to take single-view input during inference. We further show qualitative results in Fig. 9.

4.4 Learning from Single-View Videos with a Static Camera

In the last camera setting, videos are filmed by a single, static camera, and NVS task becomes non-applicable due to a lack of viewpoint changes throughout a video sequence. While our method is motivated by leveraging the spatial and temporal cues from scenes, we evaluate our method under this camera setting for completeness, and show that our method can decompose scenes into layers based on different motion patterns.

We collect a dataset of 175 moving clouds videos from YouTube, which will be made publicly available. Our method produces realistic VP results as shown in Fig. 10. Since each affine transformation layer is weighted by the predicted alpha map, the overall scene dynamics is not restricted to be affine. We defer further analysis on modeling scene dynamics in such camera setting to Appendix C.4.



Fig. 11: Visualization of intermediate predictions on KITTI dataset. (a) The second input frame, which is the source image for MPI representations. (b) and (c) are depth and flow predictions. In (d-f), we show 3 out of 16 RGBA planes predicted by the model, ordered from far to near.

4.5 Qualitative Analysis of the Generalized MPIs

To understand how the model learns to solve NVS and VP simultaneously, we qualitatively analyze the intermediate model outputs as shown in Fig. 11. The depth map is computed from Equation (14), and the flow map is composed from per-layer flow maps similarly. In the shown example, the camera is moving forward and objects closer to the camera tend to have more dominant motions. With the generalized MPI representation, the scene dynamics is decomposed into plane-wise motion fields. Since each plane is depth-aware, this representation helps introduce an inductive bias for the model to learn the depth-motion correspondence which facilitates the learning of the two extrapolation tasks.

5 Conclusion

In this work, we view NVS and VP as extrapolation along two axes for the spatial-temporal coordinates of videos. NVS utilizes camera viewpoint changes in a video sequence to discover depth, while VP considers both camera and object motions. The two tasks can be jointly learned to develop a scene representation from video data with complementary learning signals coming from each of the tasks. We propose a generalized MPI representation to tackle both tasks, and develop a model that achieves superior or comparable performance compared to previous methods that tackle only one of the tasks, on natural datasets for indoor and outdoor scenes. Please see Appendix C.5 for more discussions.

Acknowledgement. We thank Angjoo Kanazawa, Hong-Xing (Koven) Yu, Huazhe (Harry) Xu, Noah Snavely, Ruohan Zhang, Ruohan Gao, and Shangzhe (Elliott) Wu for detailed feedback on the paper, and Kaidi Cao for collecting the cloud dataset. This work is in part supported by the Stanford Institute for Human-Centered AI (HAI), the Stanford Center for Integrated Facility Engineering (CIFE), the Samsung Global Research Outreach (GRO) Program, and Amazon, Autodesk, Meta, Google, Bosch, and Adobe.

15

References

- Bei, X., Yang, Y., Soatto, S.: Learning semantic-aware dynamics for video prediction. In: CVPR (2021)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: CVPR (2009)
- Du, Y., Zhang, Y., Yu, H.X., Tenenbaum, J.B., Wu, J.: Neural radiance flow for 4D view synthesis and video processing. In: ICCV (2021)
- Flynn, J., Broxton, M., Debevec, P., DuVall, M., Fyffe, G., Overbeck, R., Snavely, N., Tucker, R.: DeepView: View synthesis with learned gradient descent. In: CVPR (2019)
- 5. Gao, H., Xu, H., Cai, Q.Z., Wang, R., Yu, F., Darrell, T.: Disentangling propagation and generation for video prediction. In: ICCV (2019)
- Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The KITTI dataset. The International Journal of Robotics Research 32(11), 1231–1237 (2013)
- Girdhar, R., Ramanan, D.: CATER: A diagnostic dataset for compositional actions and temporal reasoning. In: ICLR (2020)
- 8. Hu, R., Ravi, N., Berg, A.C., Pathak, D.: Worldsheet: Wrapping the world in a 3d sheet for view synthesis from a single image. In: ICCV (2021)
- 9. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
- Lai, Z., Liu, S., Efros, A.A., Wang, X.: Video autoencoder: self-supervised disentanglement of static 3D structure and motion. In: ICCV (2021)
- 11. Li, J., Feng, Z., She, Q., Ding, H., Wang, C., Lee, G.H.: MINE: Towards continuous depth mpi with nerf for novel view synthesis. In: ICCV (2021)
- 12. Li, Z., Niklaus, S., Snavely, N., Wang, O.: Neural scene flow fields for space-time view synthesis of dynamic scenes. In: CVPR (2021)
- 13. Lin, K.E., Xiao, L., Liu, F., Yang, G., Ramamoorthi, R.: Deep 3D mask volume for view synthesis of dynamic scenes. In: ICCV (2021)
- Liu, A., Tucker, R., Jampani, V., Makadia, A., Snavely, N., Kanazawa, A.: Infinite nature: Perpetual view generation of natural scenes from a single image. In: ICCV (2021)
- Liu, Z., Yeh, R., Tang, X., Liu, Y., Agarwala, A.: Video frame synthesis using deep voxel flow. In: ICCV (2017)
- 16. Lotter, W., Kreiman, G., Cox, D.: Deep predictive coding networks for video prediction and unsupervised learning. In: ICLR (2017)
- Lu, E., Cole, F., Dekel, T., Xie, W., Zisserman, A., Salesin, D., Freeman, W.T., Rubinstein, M.: Layered neural rendering for retiming people in video. In: SIGGRAPH Asia (2020)
- Lu, E., Cole, F., Dekel, T., Zisserman, A., Freeman, W.T., Rubinstein, M.: Omnimatte: Associating objects and their effects in video. In: CVPR (2021)
- 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)
- Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: ICCV (2021)
- 21. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: CVPR (2021)
- Ranzato, M., Szlam, A., Bruna, J., Mathieu, M., Collobert, R., Chopra, S.: Video (language) modeling: a baseline for generative models of natural videos. arXiv preprint arXiv:1412.6604 (2014)

- 16 Y. Zhang and J. Wu
- 23. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR (2016)
- 24. Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise view selection for unstructured multi-view stereo. In: ECCV (2016)
- Shade, J., Gortler, S., He, L.w., Szeliski, R.: Layered depth images. In: SIGGRAPH (1998)
- Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In: NeurIPS (2015)
- 27. Shih, M.L., Su, S.Y., Kopf, J., Huang, J.B.: 3D photography using context-aware layered depth inpainting. In: CVPR (2020)
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
- Srinivasan, P.P., Tucker, R., Barron, J.T., Ramamoorthi, R., Ng, R., Snavely, N.: Pushing the boundaries of view extrapolation with multiplane images. In: CVPR (2019)
- Srivastava, N., Mansimov, E., Salakhutdinov, R.: Unsupervised learning of video representations using LSTMs. In: ICML (2015)
- Tretschk, E., Tewari, A., Golyanik, V., Zollhöfer, M., Lassner, C., Theobalt, C.: Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In: ICCV (2021)
- Tucker, R., Snavely, N.: Single-view view synthesis with multiplane images. In: CVPR. pp. 551–560 (2020)
- Tulsiani, S., Tucker, R., Snavely, N.: Layer-structured 3D scene inference via view synthesis. In: ECCV (2018)
- Villegas, R., Yang, J., Hong, S., Lin, X., Lee, H.: Decomposing motion and content for natural video sequence prediction. In: ICLR (2017)
- Wang, J.Y.A., Adelson, E.H.: Layered representation for motion analysis. In: CVPR (1993)
- Wang, Y., Wu, H., Zhang, J., Gao, Z., Wang, J., Yu, P., Long, M.: Predrnn: A recurrent neural network for spatiotemporal predictive learning. IEEE TPAMI (2022)
- 37. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. IEEE TIP **13**(4), 600–612 (2004)
- Wiles, O., Gkioxari, G., Szeliski, R., Johnson, J.: SynSin: End-to-end view synthesis from a single image. In: CVPR (2020)
- Wu, Y., Gao, R., Park, J., Chen, Q.: Future video synthesis with object motion prediction. In: CVPR (2020)
- Xian, W., Huang, J.B., Kopf, J., Kim, C.: Space-time neural irradiance fields for free-viewpoint video. In: CVPR (2021)
- Yoon, J.S., Kim, K., Gallo, O., Park, H.S., Kautz, J.: Novel view synthesis of dynamic scenes with globally coherent depths from a monocular camera. In: CVPR. pp. 5336–5345 (2020)
- 42. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelNeRF: Neural radiance fields from one or few images. In: CVPR (2021)
- 43. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep networks as a perceptual metric. In: CVPR (2018)
- 44. Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. In: SIGGRAPH (2018)
- Zhou, T., Tulsiani, S., Sun, W., Malik, J., Efros, A.A.: View synthesis by appearance flow. In: ECCV (2016)