

Forecasting Future Videos from Novel Views via Disentangled 3D Scene Representation

Sudhir Yarram¹ and Junsong Yuan¹

State University of New York at Buffalo
<https://skrya.github.io/projects/ffn-dsr>

Abstract. Video extrapolation in space and time (VEST) enables viewers to forecast a 3D scene into the future and view it from novel viewpoints. Recent methods propose to learn an entangled representation, aiming to model layered scene geometry, motion forecasting and novel view synthesis together, while assuming simplified affine motion and homography-based warping at each scene layer, leading to inaccurate video extrapolation. Instead of entangled scene representation and rendering, our approach chooses to disentangle scene geometry from scene motion, via lifting the 2D scene to 3D point clouds, which enables high quality rendering of future videos from novel views. To model future 3D scene motion, we propose a disentangled two-stage approach that initially forecasts ego-motion and subsequently the residual motion of dynamic objects (e.g., cars, people). This approach ensures more precise motion predictions by reducing inaccuracies from entanglement of ego-motion with dynamic object motion, where better ego-motion forecasting could significantly enhance the visual outcomes. Extensive experimental analysis on two urban scene datasets demonstrate superior performance of our proposed method in comparison to strong baselines.

1 Introduction

The Video Extrapolation in Space and Time (VEST) task, as introduced in [55], combines the tasks of novel view synthesis and video prediction. VEST involves extrapolating scenes to novel views while simultaneously being able to forecast future videos. Addressing VEST requires solving intricate challenges such as estimating scene geometry, forecasting future motion, and synthesizing disoccluded content due to object movement and novel viewpoints.

Modeling the complex relationship between view changes (*geometry*) and future scene (*motion*) for an entire scene presents a significant challenge. To address this complexity, existing approaches like VEST-MPI [55] leverage generalized multi-plane image (MPI) representation. This technique simplifies the task by breaking down images into layers based on fixed depth, facilitating the modeling of view changes and future motion on a per-layer basis.

Despite the efficacy of this MPI-based approach, it has two major limitations, as illustrated in Fig. 1. First, the learned representation of each layer tends to be highly entangled. It attempts to simultaneously model scene geometry, future scene motion and the synthesis of disoccluded content. This entanglement arises from the assumptions that all pixels within a layer can render novel

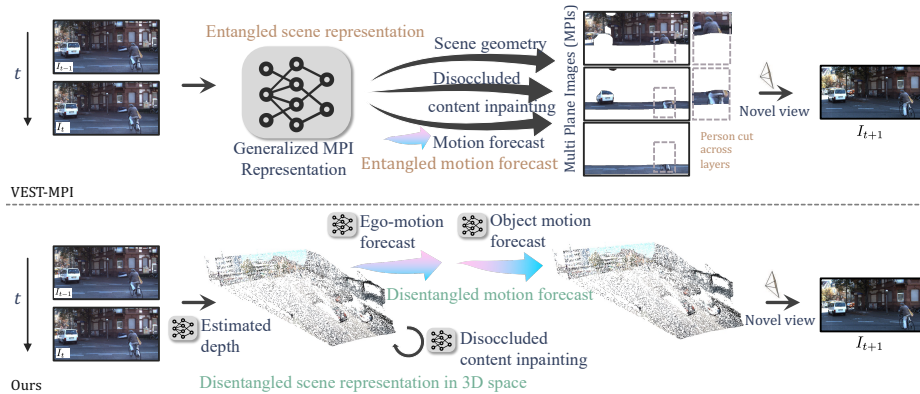


Fig. 1: Comparisons of VEST approaches. Compared to VEST-MPI [55], our method features: **(1) Disentangled 3D geometry and motion representation:** While VEST-MPI [55] relies on an entangled layered MPI representation with simplified affine motion and homography-based warping, we employ depth maps to transform 2D images into 3D point clouds, disentangling scene geometry from motion for high-quality rendering from novel viewpoints; **(2) Disentangled ego-motion and object motion forecast:** Departing from VEST-MPI’s simultaneous modeling of ego-motion and object motion, we adopt a disentangled two-stage forecasting approach. Our approach first predicts ego-motion, then addresses residual object motion. This separation allows our model to predict 3D motion more accurately, improving the accuracy of 3D motion forecasts.

view through simple homography-based warping, and future scene motions can be approximated via affine transformations alongside flow-based warping. However, these assumptions often fall short in complex scenes, such as urban driving scenarios, where the sophisticated geometry and the diverse motions of objects within each layer challenge these simplifications. Consequently, this can result in blurred visual content, especially in long-term forecasts.

Second, the layered MPI representation faces challenges in accurately modeling future 3D motion, primarily due to the complexity inherent in predicting both ego-motion and dynamic object motion within each layer. Ego-motion affects all observed elements in the scene and can become entangled with the irregular motion of dynamic objects, potentially causing inaccuracies in the motion prediction. The precise ego-motion forecasting is crucial since even slight errors can lead to significant deviations in the visual outcome. It thus underscores the importance of disentangling these motion components to achieve more accurate motion predictions.

To address these challenges, we propose an approach that models 3D scene geometry within a continuous space of point clouds rather than a layered MPI. This continuous scene representation aids in the disentanglement of scene geometry from scene motion, enabling the rendering of scene from novel camera perspectives, thereby achieving high-fidelity motion forecasting with photorealistic synthesis of future videos. The key to our approach is to leverage estimated

depth map and lift the scene representation from the 2D RGB images to 3D point clouds. Each 3D point is defined by its 3D location and a learned appearance feature. These point clouds can effectively account for 3D scene motion by displacing 3D location of the corresponding point based on its physical motion flow. Finally, we can use sophisticated 3D-to-2D rendering techniques [31] to process these point clouds, to synthesize the future frame from a novel view, instead of simple homography-based warping.

To accurately forecast future motion, moving away from simultaneously estimating ego-motion and object motions used in VEST-MPI [55], we propose a disentangled two-stage 3D motion forecasting approach. Our two-stage approach begins by forecasting camera ego-motion and subsequently addresses the residual motion of dynamic objects (*e.g.* car, people), thereby modeling future scene motion. Since static background motion in past frames can estimate future ego-motion, we introduce an ego-motion forecasting module that analyzes static background elements to forecast ego-motion. Subsequently, we introduce a multi-scale object forecasting module to address the residual motion of objects. This module analyzes the scene elements across past frames to predict residual object motion. Utilizing the forecasted 3D motion, we relocate the point clouds, creating a future 3D scene and then synthesizing the future frame.

In summary, our contributions include: (1) a novel approach that leverages explicit 3D scene geometry to address VEST by disentangling 3D scene geometry from scene motion, enabling the rendering of scene from novel views, and thereby achieving high-fidelity scene forecasting with photorealistic synthesis of future videos. (2) Disentangled two-stage approach for future motion forecast that first forecasts ego-motion and then the residual motion of dynamic objects, in order to achieve high quality video extrapolation in space and time.

Our method outperforms baselines in VEST, video prediction, and novel view synthesis on benchmark datasets such as KITTI [9] and Cityscapes [5].

2 Related Work

Video prediction. Previous video prediction (VP) methods vary from unconditional synthesis [4, 6–8, 18, 33, 40] to conditional prediction tasks [1, 10, 13, 19, 34, 47, 50]. These conditional prediction methods use diverse strategies to forecast motion. Some decompose videos using object-centric [50], semantic-aware [1] and motion-aware [44] approaches to forecast motion. DMVFN [13] proposes to handle diverse motion scales of objects through a pyramidal approach [30, 38], while [19] forecasts the trajectory of object through sparse control points. However, these methods work in either the 2D or quantized 3D domain, facing challenges in modeling occlusions. In contrast, our approach employs 3D motion forecasting, overcoming these limitations by incorporating the depth dimension. This integration enables for accurate estimates of the object motion patterns as they move in and out of each other’s paths.

Novel view synthesis. Novel view synthesis aims to reconstruct unseen viewpoints from a set of input 2D images and their corresponding camera poses. Recent neural rendering methods have achieved impressive synthesis results [23, 24, 27,

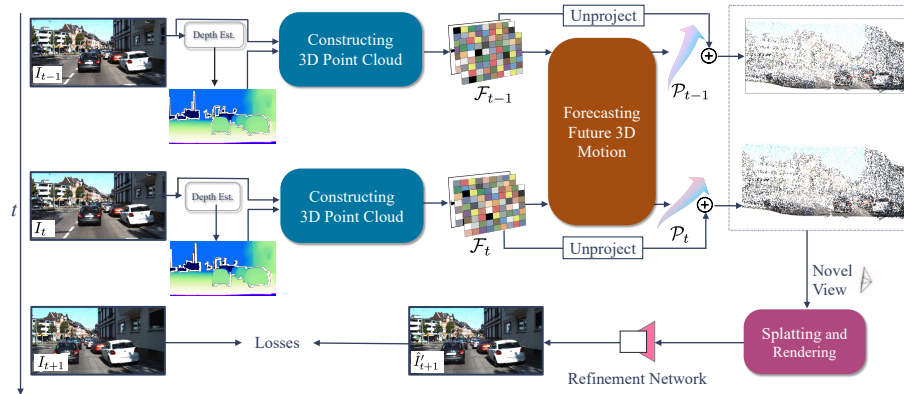


Fig. 2: Method overview. Our framework aims to forecast a 3D scene into the future and view it from novel viewpoints. It comprises three primary steps **(1) Constructing 3D point clouds:** Starting with two past frames as the input, we construct per-frame 3D point clouds. (i) The process for each frame involves depth estimation, dis-occlusion handling via inpainting, and feature extraction to finally generate what we refer to as feature layer. (ii) The point-wise features in this feature layer are then lifted into 3D space using corresponding depth values, generating 3D point clouds. This process is performed on both $I_{(t-1)}$ and $I_{(t)}$ to obtain feature layers $\mathcal{F}_{(t-1)}$ and $\mathcal{F}_{(t)}$ and point clouds $\mathcal{P}_{(t-1)}$ and $\mathcal{P}_{(t)}$. **(2) Forecasting future 3D motion:** We leverage the feature layers $\mathcal{F}_{(t-1)}$ and $\mathcal{F}_{(t)}$ to forecast future 3D motion for each of the point clouds. This forecasted 3D motion allows us to update the positions of point clouds $\mathcal{P}_{(t-1)}$ and $\mathcal{P}_{(t)}$ to their new, forecasted locations. **(3) Splating and Rendering:** A point-based renderer processes these motion-adjusted point clouds through 3D-to-2D splating to generate feature maps. Finally, refinement network takes these rendered feature maps and decodes them to synthesize a novel view $I'_{(t+1)}$ based on the target viewpoint.

28, 37, 39, 46, 46]. However, these methods typically assume the availability of dense views as input, a condition rarely met in practical scenarios. Moreover, their primary focus lies in interpolating across frames rather than extrapolation. Therefore, we consider methods capable of handling sparse views as input.

Among these, several works [11, 20, 22, 41, 42, 52, 53] aim to infer the 3D structure of scenes by learning to predict a scene representation from a single image. Yet, these methods lack the ability to forecast into the future.

Depth-based 3D scene representation. Recent “3D photography” methods [15, 16, 21, 35] demonstrate how estimated depth can bring 2D images to life for a 3D viewing experience. For instance, 3D Photo [35] employs monocular depth mapping to generate layered depth images (LDIs) [29], enabling contextually aware color and depth inpainting. These methods benefit by disentangling 3D scene geometry from camera pose. However, unlike our approach, these methods assume a static scene. Among these, 3D Moments [45] and Li et al.’s work [21] are most related to our approach as they achieve space-time view synthesis. However, 3D Moments [45] focuses on space-time interpolation across frames and Li et al [21] focuses on animating a static scene. In contrast, our approach conditions on past frames to forecast future scene motion.

3 Our Method

Given two consecutive frames from a video sequence, $\mathbf{I}_{(t-1)}$ and $\mathbf{I}_{(t)}$ as the input, our goal is to forecast a future frame $\hat{\mathbf{I}}'_{(t+1)}$ rendered from a novel viewpoint. Our framework is illustrated in Fig. 2. The output of our framework is a future video featuring high-fidelity video forecast with high-quality novel view synthesis.

Initially, we generate an estimated depth map for the input frames. Subsequently, we construct 3D point clouds by first addressing potential dis-occlusions in each frame through inpainting, followed by feature map extraction and their unprojection into 3D space, based on depth values (Sec. 3.1). For future 3D motion prediction, since the 3D motion would result from both camera motion and dynamic object motions, we disentangle these 3D motions via a two-stage approach that initially forecasts ego-motion by leveraging static background regions across frames, and subsequently predicting residual motion for dynamic objects (e.g., cars, persons) (Sec. 3.4). With these forecasts, we update the point cloud’s location to reflect the anticipated 3D motion. Finally, we project these modified 3D point clouds into 2D space, according to the target viewpoint, to synthesize the future frame (Sec. 3.3).

Our framework operates in a fully end-to-end differentiable manner, incorporating ego-motion forecasting modules and object motion forecasting modules, and directly supervises the training process based on the $L1$ loss and perceptual losses [36], which can be obtained by comparing the predicted future frame $\hat{\mathbf{I}}'_{(t+1)}$ and the ground truth future frame $\mathbf{I}_{(t+1)}$.

3.1 Disentangled 3D Scene Representation

Instead of quantizing a 3D scene into multiple depth layers as in [55], we introduce an approach to model the geometry of a 3D scene as a continuous space of point clouds. This is achieved via lifting scene from 2D RGB images to 3D through estimated depth map. This simple yet effective treatment provides feasibility to disentangle scene geometry from scene motion, enabling better rendering of the scene from novel camera perspectives. Moreover, For each input frame, we build a 3D point cloud $\mathcal{P} = \{(\mathbf{x}^i, \mathbf{f}^i)\}$, where \mathbf{x}^i represents the 3D location of each 3D point and \mathbf{f}^i signifies its appearance feature. Furthermore, we disentangle the process of forecasting future 3D motion to first forecast ego-motion and subsequently predict residual motion for objects. Such disentanglement enables the model to more accurately forecast 3D motion by focusing on the crucial ego-motion forecasting, where even minor deviation to precisely forecast can lead to significant deviations in visual outcome, then focuses on object motion at various scales. The following sections detail the process of constructing these point clouds and the methodology for forecasting future 3D motion \mathbf{u}^i , aiming to displace each 3D point to where it is forecasted to be in the future.

3.2 Constructing 3D Point Cloud

As shown in Fig. 3, we start by estimating the 3D geometry of the input frame \mathbf{I} using the metric depth estimation methods such as ZoeDepth [2], which can

estimate depth in absolute physical units, e.g. meters. Accurate depth map is crucial for 3D scene representation. The resulting depth map is denoted by \mathbf{D} .

Semantic segmentation. A major challenge in synthesizing photorealistic future frames is handling dis-occlusions caused by dynamic object motion and transitions to novel views. These dis-occlusions appear as “holes” in the future frames, where visual information is missing. Many 3D photography methods [15, 21, 35, 45] employ depth-based layering [20] to address dis-occlusions during view transitions. Despite their efficacy, these methods often struggle with dis-occlusions due to the motion of objects within each depth layer.

To better address dis-occlusions, we leverage semantic segmentation on the input image to identify potential regions for “holes” by segmenting both dynamic category objects and likely foreground category objects. Semantic segmentation is performed on \mathbf{I} to generate binary mask \mathbf{M} . This mask distinguishes static/background categories (e.g., building, roads, sidewalk, vegetation) from dynamic/foreground categories (e.g., car, bus, pedestrian, pole). For simplicity, we refer to the former category regions as the background and the later as foreground. Note that binary mask \mathbf{M} specifically identifies the background regions. For a detailed categorization, please check supplementary material.

Image and Depth Inpainting. To mitigate dis-occlusions, we leverage \mathbf{M} to inpaint the areas where foreground regions are masked out in both \mathbf{I} and its depth map \mathbf{D} using the surrounding background as context. Importantly, we ensure that the depth values assigned to the inpainted regions are farther in depth than those pixels belonging to the foreground objects. This is given by:

$$\mathbf{I}^{\overline{\text{BG}}}, \mathbf{D}^{\overline{\text{BG}}} = \text{Inpaint}(\mathbf{I} \odot \mathbf{M}, \mathbf{D} \odot \mathbf{M}, \mathbf{M}), \quad (1)$$

where \odot is element-wise multiplication. Here, $\mathbf{I}^{\overline{\text{BG}}}$ refers to the inpainted image frame, and $\mathbf{D}^{\overline{\text{BG}}}$ is the inpainted depth map.

Feature Encoding. To enhance rendering quality and minimize artifacts in the synthesized future frames, we employ a 2D feature extraction network. This network processes both the original frame, \mathbf{I} and inpainted frame, $\mathbf{I}^{\overline{\text{BG}}}$, generating features \mathbf{F} , and $\mathbf{F}^{\overline{\text{BG}}}$, respectively.

Subsequently, we unproject the features \mathbf{F} from the original 2D frame into 3D space using the corresponding depth map \mathbf{D} . To ensure that visual content is present behind the foreground objects in 3D space, we unproject the inpainted regions, using the corresponding features $\mathbf{F}^{\overline{\text{BG}}}$ and depth map $\mathbf{D}^{\overline{\text{BG}}}$ dedicated to these regions. This unprojection process facilitates the creation of point clouds, represented as \mathcal{P} :

$$\mathcal{P} = \text{Unproject}(\mathbf{F}, \mathbf{D}) \cup \text{Unproject}(\mathbf{F}^{\overline{\text{BG}}} \diamond [1 - \mathbf{M}], \mathbf{D}^{\overline{\text{BG}}} \diamond [1 - \mathbf{M}]), \quad (2)$$

where $1 - \mathbf{M}$ inverts the binary mask and \diamond operator selectively retains pixel-wise features from \mathbf{F} for which the corresponding value in $1 - \mathbf{M}$ is 1. For simplicity, we refer to the set $\{\mathbf{F}, \mathbf{D}, \mathbf{M}\}$ as *original feature layer*, denoted by \mathcal{F} and the set

$\{\mathbf{F}^{\overline{\text{BG}}}, \mathbf{D}^{\overline{\text{BG}}}, \mathbf{M}\}$ as *inpainted feature layer* $\mathcal{F}^{\overline{\text{BG}}}$. Applying the above process to both input frames $\mathbf{I}_{(t-1)}$ and $\mathbf{I}_{(t)}$ results in:

$$\begin{aligned}\mathcal{F}_{(t-1)} &= \{\mathbf{F}_{(t-1)}, \mathbf{D}_{(t-1)}, \mathbf{M}_{(t-1)}\}, \mathcal{F}_{(t-1)}^{\overline{\text{BG}}} = \{\mathbf{F}_{(t-1)}^{\overline{\text{BG}}}, \mathbf{D}_{(t-1)}^{\overline{\text{BG}}}, \mathbf{M}_{(t-1)}\}, \\ \mathcal{F}_{(t)} &= \{\mathbf{F}_{(t)}, \mathbf{D}_{(t)}, \mathbf{M}_{(t)}\}, \mathcal{F}_{(t)}^{\overline{\text{BG}}} = \{\mathbf{F}_{(t)}^{\overline{\text{BG}}}, \mathbf{D}_{(t)}^{\overline{\text{BG}}}, \mathbf{M}_{(t)}\}.\end{aligned}\quad (3)$$

And the resultant 3D point clouds $\mathcal{P}_{(t-1)}$ and $\mathcal{P}_{(t)}$ can be obtained from feature layers $\{\mathcal{F}_{(t-1)}, \mathcal{F}_{(t-1)}^{\overline{\text{BG}}}\}$ and $\{\mathcal{F}_{(t)}, \mathcal{F}_{(t)}^{\overline{\text{BG}}}\}$ using Unproject operation as mentioned in Eq. 2. Next, we outline the methodology for generating a future frame using the point clouds $\mathcal{P}_{(t-1)}$ and $\mathcal{P}_{(t)}$, and the forecasted 3D motion for these point clouds. This is followed by a detailed description of forecasting the 3D motion.

3.3 Splatting and Rendering

In order to generate the future frame given a novel camera pose, we leverage the differentiable point-based rendering technique [31] that ‘‘splats’’ the 3D points in motion adjusted point clouds onto a 2D image plane, synthesizing novel views.

Given two input frames \mathbf{I}_{t-1} and \mathbf{I}_t , we have point clouds $\mathcal{P}_{(t-1)}$ and $\mathcal{P}_{(t)}$. Let $\mathbf{u}_{(t-1)\rightarrow(t+1)}$ refer to the forecasted 3D motion flow for 3D points in $\mathcal{P}_{(t-1)}$ and $\mathbf{u}_{(t)\rightarrow(t+1)}$ to the forecasted 3D motion flow for 3D points in $\mathcal{P}_{(t)}$. The method for calculating these motion flows will be explained later. The process of generating future frames from a novel viewpoint \mathbf{K} can be summarized in three steps:

1. Displacement by 3D Motion Flow:

$$\mathcal{P}_{(t-1)\rightarrow(t+1)} = \mathcal{P}_{(t-1)} + \mathbf{u}_{(t-1)\rightarrow(t+1)}, \mathcal{P}_{(t)\rightarrow(t+1)} = \mathcal{P}_{(t)} + \mathbf{u}_{(t)\rightarrow(t+1)}, \quad (4)$$

2. Scene Rendering via Point-based Splatting:

$$\mathbf{F}_{(t-1)\rightarrow(t+1)} = \text{Render}(\mathcal{P}_{(t-1)\rightarrow(t+1)}, \mathbf{K}), \mathbf{F}_{(t)\rightarrow(t+1)} = \text{Render}(\mathcal{P}_{(t)\rightarrow(t+1)}, \mathbf{K}), \quad (5)$$

3. Future Frame Synthesis:

$$\hat{\mathbf{I}}_{(t+1)} = \text{Refine}(\text{Concat}[\mathbf{F}_{(t-1)\rightarrow(t+1)}, \mathbf{F}_{(t)\rightarrow(t+1)}]). \quad (6)$$

where ‘‘Concat’’ refers to the concatenation of two forecasted features $\mathbf{F}_{(t-1)\rightarrow(t+1)}$, $\mathbf{F}_{(t)\rightarrow(t+1)}$. Here, the Render function employs a point-based renderer [31], while the Refine function uses a refinement network based on 2D U-Net architecture [32]. And $\hat{\mathbf{I}}_{(t+1)}$ is the forecasted future frame. For long-term forecasts, we can recursively predict future frames $\{\hat{\mathbf{I}}_{t+2}, \hat{\mathbf{I}}_{t+3}, \dots\}$ in a similar manner.

3.4 Forecasting Future 3D Motion

Now, we detail the process of forecasting future 3D motion flows, $\mathbf{u}_{(t-1)\rightarrow(t+1)}$ and $\mathbf{u}_{(t)\rightarrow(t+1)}$, as depicted in Fig. 3.

Most previous video prediction methods focused on forecasting 2D scenes [1, 13, 25, 51], or simplified layered 3D scene. They often rely on pixel-wise backward warping [14] to forecast the next frame from previous frames. However,

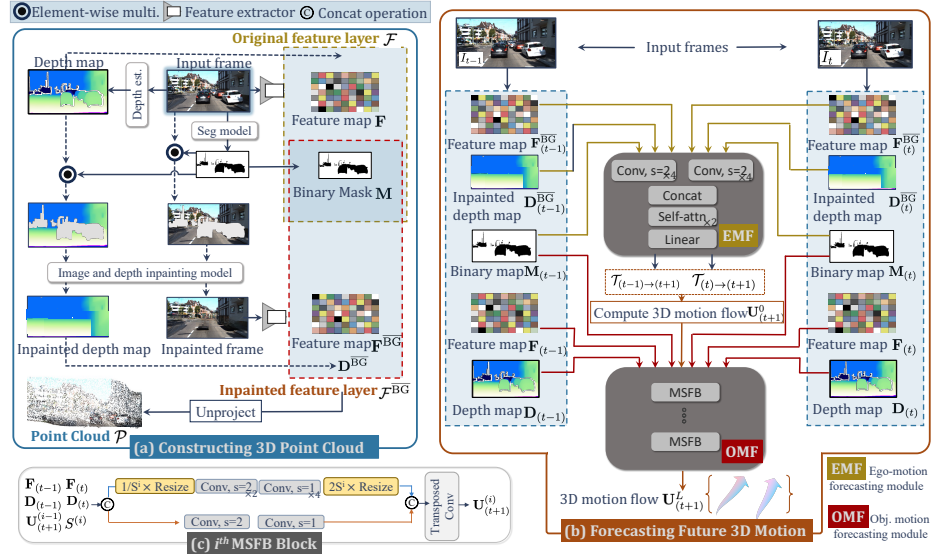


Fig. 3: (a) Constructing 3D point cloud. (1) Estimate the depth map D from the input image I . (2) Address “holes” in future frames caused by dis-occlusions from dynamic object motion: (i) segment dynamic category (foreground) objects to produce a binary mask M , identifying potential regions for “holes”. (ii) mask these foreground regions in both input image and depth map, then inpaint them using the background context. (3) Extract features from both original and inpainted frames to produce F and F^{BG} . (4) Create 3D point cloud \mathcal{P} by unprojecting the 2D features F , F^{BG} into 3D, using depth maps D , D^{BG} , respectively. For simplicity, we refer to the set $\{F, D, M\}$ as *original feature layer*, denoted by \mathcal{F} and the set of $\{F^{BG}, D^{BG}, M\}$ as *inpainted feature layer* \mathcal{F}^{BG} . **(b) Forecasting future 3D motion.** Given feature layers from past frames, our method forecasts future 3D motion flow in two stages: (1) ego-motion forecasting using the EMF module, which processes the background (static category) across frames using inpainted feature layers $\mathcal{F}_{(t-1)}^{BG}$ and $\mathcal{F}_{(t)}^{BG}$, yielding two relative ego-pose transformations, $\mathcal{T}_{(t-1) \rightarrow (t+1)}$ and $\mathcal{T}_{(t) \rightarrow (t+1)}$. These transformations lead to initial 3D motion flows $\mathbf{u}_{(t-1) \rightarrow (t+1)}^0$ and $\mathbf{u}_{(t) \rightarrow (t+1)}^0$, referred to as $\mathbf{U}_{(t+1)}^0$. (2) The OMF module then refines the initial 3D motion flow $\mathbf{U}_{(t+1)}^0$ by accounting for foreground object motion, using original and inpainted feature layers to derive the final forecasted 3D motion flow, $\mathbf{U}_{(t+1)}^L$, after L MMFB blocks. **(c) Multi-scale motion flow block (MMFB).** We illustrate the design of a MMFB block here.

2D motion forecasting may encounter ambiguities when dealing with occlusions. Additionally, backward warping inherently lacks one-to-one pixel correspondence across frames, resulting in content stretching. In contrast, our 3D motion forecasting mitigates the ambiguities around occluded regions by considering the extra depth information, enabling more accurate estimations of how objects move in and out of each other’s paths. To avoid content stretching, we opt for a one-to-one correspondence among 3D points across the future and the past frames.

To accurately forecast future 3D motion, we propose a **dis-entangled two-stage** approach that first forecasts relative ego-motion and then addresses the residual 3D motion of dynamic objects (e.g., cars, people). Ego-motion impacts every observed element within the scene, and its interaction with the irregular motion of dynamic objects can lead to inaccuracies in motion prediction. By disentangling these two components, we can achieve more reliable and detailed forecasts. Moreover, ego-motion forecasting is crucial as it influences all observed elements in the scene. Failure to precisely forecast ego-motion can lead to significant inaccuracies in the predicted scene dynamics, as even minor deviations in camera trajectory prediction can cause large deviations in the visual outcome.

Ego-motion forecasting (EMF). To address this, we design an ego-motion forecasting (EMF) module that infers relative pose changes of the camera from previous frames to future frames. Our EMF module E takes inpainted features and inpainted depth along with binary mask from inpainted feature layers $\mathcal{F}_{(t-1)}^{\text{BG}}$ and $\mathcal{F}_{(t)}^{\text{BG}}$ as input, and forecasts the two separate pose transformations $\mathcal{T}_{(t-1)\rightarrow(t+1)}$ and $\mathcal{T}_{(t)\rightarrow(t+1)}$. The transformation \mathcal{T} is parameterized as a rotation matrix \mathbf{R} in the quaternion form (q_w, q_x, q_y, q_z) and translation vector \mathbf{t} is (t_x, t_y, t_z) . The EMF module can be defined as:

$$\mathcal{T}_{(t-1)\rightarrow(t+1)}, \mathcal{T}_{(t)\rightarrow(t+1)} = E\left(\mathbf{F}_{(t-1)}^{\text{BG}}, \mathbf{D}_{(t-1)}^{\text{BG}}, \mathbf{M}_{(t-1)}, \mathbf{F}_{(t)}^{\text{BG}}, \mathbf{D}_{(t)}^{\text{BG}}, \mathbf{M}_{(t)}\right). \quad (7)$$

EMF processes past inpainted features $\mathbf{F}_{(t-1)}^{\text{BG}}$ and $\mathbf{F}_{(t)}^{\text{BG}}$ with corresponding inpainted depths $\mathbf{D}_{(t-1)}^{\text{BG}}$ and $\mathbf{D}_{(t)}^{\text{BG}}$ through several convolutional layers. After that, these features are concatenated and passed through a series of self-attention blocks [43], using masks $\mathbf{M}_{(t-1)}$ and $\mathbf{M}_{(t)}$ to focus on non-inpainted background spatial positions. The self-attention blocks allows the module to focus on relevant features across the temporal sequence, enhancing its ability to predict future ego-motion by identifying patterns and dependencies in the movement of the background across frames. Following the self-attention blocks, the resultant features are input to a linear layer to output the 7 (4 rotation and 3 translation) parameters for each of the transformations $\mathcal{T}_{(t-1)\rightarrow(t+1)}$ and $\mathcal{T}_{(t)\rightarrow(t+1)}$.

Computing 3D motion flow. Finally, for each of the predicted ego-pose change $\mathcal{T} = (\mathbf{R}, \mathbf{t})$ from Eq. 7, we compute the 3D motion flow \mathbf{u} for every 3D point $p \in \mathcal{P}$ as follows:

$$\mathbf{u} = \mathbf{R}\mathbf{x} + \mathbf{t} - \mathbf{x}, \quad (8)$$

where \mathbf{x} is the 3D location of point p .

Object motion forecasting (OMF). Dynamic objects display a varied nature of motions within a given scene. To handle this diversity, we propose an object motion forecasting module. This module is specifically tailored to predict the residual motion of objects after taking into account the ego-motion, while accounting for the object motions. By doing so, we can achieve a more nuanced and accurate prediction of each object’s 3D motion, enhancing the overall realism and coherence of the forecasted 3D scene.

Our object motion forecasting module (OMF) is inspired by the 2D motion forecasting module in DMVFN [13], which preserves spatial resolution and captures significant motion. We thus propose new multi-scale motion flow block (MMFB), constructed using a series of convolutional blocks, with its architecture depicted in Fig. 3. Unlike DMVFN, which focuses on appearance features to predict 2D optical flow, our MMFB leverages both appearance features and depth information to estimate future 3D motion.

Our OMF module consists of L MMFB blocks. As shown in Fig. 3, each MMFB block contains two branches for processing the input frame along with the depth map. One branch operates at downsampled scale factor S that focuses on capturing larger-scale motion, while the other branch operates at higher image resolution, enhancing spatial details. Let $\mathbf{u}_{(t) \rightarrow (t+1)}$ encompass the motion field representing the motion of each 3D point. For brevity, let $\mathbf{U}_{(t+1)} = \{\mathbf{u}_{(t-1) \rightarrow (t+1)}, \mathbf{u}_{(t) \rightarrow (t+1)}\}$ represent two predicted 3D motion flows from two separate frames. The i^{th} MMFB block, denoted as F_{MMFB}^i , learns to estimate the target 3D motion flow $\mathbf{U}_{(t+1)}^i$ by utilizing the original feature layers $\mathcal{F}_{(t-1)}$ and $\mathcal{F}_{(t)}$, as well as the estimated previous blocks motion flow $\mathbf{U}_{(t+1)}^{i-1}$ and the scale factor S^i as input:

$$\mathbf{U}_{(t+1)}^i = F_{MMFB}^i(\mathbf{F}_{(t-1)}, \mathbf{D}_{(t-1)}, \mathbf{M}_{(t-1)}, \mathbf{F}_{(t)}, \mathbf{D}_{(t)}, \mathbf{M}_{(t)}, \mathbf{U}_{(t+1)}^{i-1}, S^i). \quad (9)$$

Note that the initial pair of 3D motion flows $\mathbf{U}_{(t+1)}^0$ is obtained from ego-motion forecasting using $\mathcal{T}_{(t-1) \rightarrow (t+1)}, \mathcal{T}_{(t) \rightarrow (t+1)}$, serving as a foundation for subsequent residual motion forecasting of dynamic objects. Finally, after processing through L MMFB layers, OMF obtains the final forecasted 3D motion flow $\mathbf{U}_{(t+1)}^L = \{\mathbf{u}_{(t-1) \rightarrow (t+1)}, \mathbf{u}_{(t) \rightarrow (t+1)}\}$.

4 Experiments

Dataset and Metric. We experiment with two urban datasets. **KITTI** [9] comprises 28 driving videos with a resolution of 375×1242 , using 24 for training and the remaining 4 for testing. **Cityscapes** [5] contains 3,475 videos with resolution 1024×2048 . We use 2,945 driving videos for training and 500 videos for testing, all at a resolution of 512×1024 . Our evaluations extend to forecasting up to 5 frames into the future for the KITTI dataset and up to 10 frames for the Cityscapes dataset. For evaluation, we employ standard error metrics: the Multi-Scale Structural Similarity Index Measure (SSIM) [49] and the perceptual metric LPIPS [54], following the conventions of [19, 55].

Implementation Details. We use the following networks in our framework.

- Pre-trained Networks: we utilize the ZoeDepth [2] method for per-frame depth estimation due to its zero-shot performance across datasets, employing the pre-trained ZoeD-M12-NK model. For segmentation, we incorporate DeepLabV3 [3] pre-trained on the Cityscapes dataset, and the 3D Photo [35] method for RGB-D inpainting on a per-frame basis.

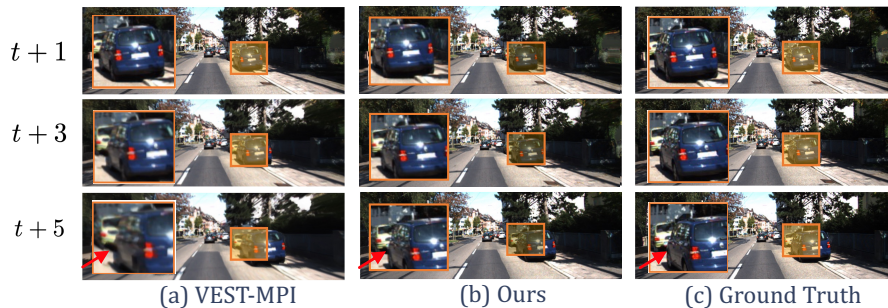


Fig. 4: Qualitative comparison with VEST-MPI [55] on video prediction task (VEST-[S,T]). The results show that our method produces sharper frames with high-quality motion forecasts, particularly over the long term.

Method	Publication	Inputs	SSIM ($\times 10^{-2}$) \uparrow			LPIPS ($\times 10^{-2}$) \downarrow			
			$t+1$	$t+3$	$t+5$	$t+1$	$t+3$	$t+5$	
DMVFN [13]	\rightarrow 3D Photo [35]	CVPR'23	R+D	85.4	72.9	62.6	13.0	22.7	29.2
WALDO [19]	\rightarrow 3D Photo [35]	ICCV'23	R+L+F+D	85.7	75.0	68.3	12.9	19.5	24.9
Ours			R+L+D	86.2	76.1	69.3	12.3	17.1	21.4

Table 1: Quantitative results for concurrent video extrapolation in space and time (VEST-[S+T]) on KITTI [9] dataset. We use the train-test splits from [19] and compare our method with state-of-the-art video prediction techniques, in combination with the 3D Photography method, 3D Photo [35], as baselines. ‘R’, ‘F’, ‘L’ and ‘D’ denote the video frames, optical flow, segmentation map, and depth map, respectively. Our method achieves comparable performance with better predictions farther into the future.

- Trainable Networks: the feature extraction and refinement networks are implemented using the ResNet34 [12] and 2D U-net architectures [32], respectively. The Ego Motion Forecasting (EMF) module leverages a network architecture elaborated in Section 3.4. Furthermore, the 9 (L) MMFB blocks within the OMF module, depicted in Figure 3, employs a decreasing scaling factor sequence of [4, 4, 4, 2, 2, 2, 1, 1, 1]. We provide more details on the architectures in the supplementary material.

Training utilizes the Adam optimizer [17] with a cosine annealing strategy and a base learning rate of 10^{-4} . The training utilizes equal weights of one for both $L1$ and perceptual losses [36]. We train on two A6000 GPUs with a batch size of 4 for 300 epochs. Further, to address issues such as foreground objects becoming semi-transparent due to gaps between samples when the camera zooms in, we use adaptive point-based rendering [31] where each 3D point can be rendered by adjusting its point radius proportional to its depth.

4.1 Baseline Methods

We compare our approach with several state-of-the-art methods in video prediction (VP) task, novel view synthesis (NVS) and video extrapolation in space and time (VEST). Additionally, we closely analyze the following baselines for further comparison.

Extrapolation	Method Publication	In space only		In time only		
		LPIPS↓	SSIM↑	LPIPS ($\times 10^{-2}$)↓		
				$t+1$	$t+3$	$t+5$
LDI [42]	ECCV'18	N/A	57.2	N/A		
MINE [20]	ICCV'21	10.8	82.2	N/A		
Tucker et al. [41]	CVPR'20	N/A	73.3	N/A		
PredRNNv2 [48]	TPAMI'22	N/A	N/A	30.8	45.7	54.2
VEST-MPI [55]	ECCV'22	<u>8.5</u>	<u>82.5</u>	<u>11.5</u>	<u>28.8</u>	<u>39.1</u>
Ours		5.2	94.6	8.1	18.6	20.4

Table 2: Results for independent VEST (VEST-[S,T]) on KITTI dataset. We follow [55] and use the LDI [42] train-test splits.

WALDO [55], a video prediction method that employs layered decomposition to split videos into object-centric layers and a background layer. Then, it generates sparse control points for these layers and forecasts these points into the future, and employs warping and occlusion handling to generate the next frame.

DMVFN [13], a video prediction method that employs dynamic 2D flow estimation via multi-scale voxel flow blocks (MVFB) to capture diverse motion cues between adjacent frames. This estimated flow is utilized for pixel-wise backward warping [14] the past frames to generate next frame.

VEST-MPI [55] method addresses both Video Prediction (VP) and Novel View Synthesis (NVS) tasks concurrently by introducing a generalized multi-plane image representation. It decomposes images into RGBA planes and parameterizes each plane to effectively model the spatial and temporal dynamics.

Experimental setups. We compare our approach with state-of-the-art (SOTA) methods using three experimental setups, following the conventions established in prior works [19, 55]. Additionally, we present a novel experimental setup and corresponding baselines.

- Independent video extrapolation in space and time (VEST-[S,T])
- Concurrent video extrapolation in space and time (VEST-[S+T])
- Video Prediction (VP)

4.2 Independent VEST (VEST-[S,T])

Setup. Following VEST-MPI [55], we conduct experiments by independently extrapolating in space (novel view synthesis) and in time (video prediction) using the KITTI dataset [9], denoted as VEST-[S,T].

Baselines. We compare our approach with VEST-MPI [55] and three leading novel view synthesis methods: LDI [42], MINE [20], and Tucker et. al [41].

Results. As shown in Table 2, our method outperforms the baselines on error metrics. Additionally, comparisons in the video prediction task are provided in Table 4. In Fig. 4, our method is compared with VEST-MPI [55] on video prediction task. VEST-MPI struggles with in accurate motion forecasting and rendering details of moving objects, leading to blurred dynamic content (denoted by the red arrow), especially in long-term forecasts. This limitation arises from its highly entangled internal representation. In contrast, our approach synthesizes sharper future frames with high-fidelity motion forecasting by disentangled modeling of 3D scene geometry, camera pose and scene motion.

Method	SSIM ($\times 10^{-2}$)↑			LPIPS ($\times 10^{-2}$)↓		
	$t+1$	$t+3$	$t+5$	$t+1$	$t+3$	$t+5$
A) w/ fixed depth	83.2	73.1	66.4	15.1	20.2	25.8
B) w/o EMF	85.6	75.5	68.4	12.9	18.1	22.4
C) w/o OMF	84.6	74.8	67.9	13.5	19.1	24.3
Full	86.2	76.1	69.3	12.3	17.1	21.4

Table 3: Ablation study on VEST-[S+T] on KITTI [9].

		Cityscapes (512 × 1024)						KITTI (256 × 832)							
Method	Publication	Inputs	$t+1$		$t+5$		$t+10$		$t+1$		$t+3$		$t+5$		
			SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	
PredNet	[26]	NeurIPS'17	R	84.0	26.0	75.2	36.0	66.3	52.2	56.3	55.3	51.4	58.6	47.5	62.9
MCNet	[44]	ICLR'17	R	89.7	18.9	70.6	37.3	59.7	45.1	53.0	24.0	63.5	31.7	55.4	37.3
VFlow	[25]	CVPR'17	R	83.9	17.4	71.1	28.8	63.4	36.6	53.9	32.4	46.9	37.4	42.6	41.5
OMP	[50]	CVPR'20	R+L	89.1	8.5	75.7	16.5	67.4	23.3	79.2	18.5	67.6	24.6	60.7	30.4
VPVFI	[51]	CVPR'22	R	94.5	6.4	80.4	17.8	70.0	27.8	82.7	12.3	69.5	20.3	61.1	26.4
CorrWise	[10]	CVPR'22	R	92.8	8.5	83.9	15.0	75.1	21.7	82.0	17.2	73.0	22.0	66.7	25.9
SADM	[1]	CVPR'21	R+L+F	<u>95.9</u>	7.6	83.5	14.9	N/A	N/A	83.1	14.4	72.4	24.6	64.7	31.2
DMVFN	[13]	CVPR'23	R	95.7	5.6	83.5	14.9	N/A	N/A	88.5	<u>10.7</u>	78.0	19.3	70.5	26.0
WALDO	[19]	ICCV'23	R+L+F	95.7	<u>4.9</u>	<u>85.4</u>	<u>10.5</u>	<u>77.1</u>	<u>15.8</u>	86.7	10.8	76.6	<u>16.3</u>	70.2	<u>20.6</u>
VEST-MPI	[55]	ECCV'22	R	N/A	N/A	N/A	N/A	N/A	N/A	15.6	N/A	34.4	N/A	44.7	
Ours			R+L+D	96.4	4.6	86.2	9.8	78.0	14.9	<u>87.7</u>	10.1	77.6	15.4	71.3	19.8

Table 4: Comparisons to state-of-the-art video prediction methods on Cityscapes [5] and KITTI [9] datasets. We compute multi-scale SSIM ($\times 10^{-2}$) and LPIPS ($\times 10^{-2}$) for the future frame evaluation. ‘R’, ‘F’, ‘L’, ‘I’ and ‘D’ denote the video frames, optical flow, segmentation map, instance map and depth map, respectively. ‘N/A’ means not available.

4.3 Concurrent VEST (VEST-[S+T])

Setup. We proposed to evaluate concurrent video extrapolation in space and time on the KITTI dataset [9], referred to as VEST-[S+T]. We use train-test splits from [19] with an image resolution of 256×832 .

Baselines. While VEST-MPI [55] provides a promising baseline for comparison, the unavailability of pre-trained models and challenges in reproducibility have limited its application in our study. Instead, we use alternate strong baselines:

WALDO [19] \rightarrow **3D Photo** [35]: Combining video prediction and single-image novel view synthesis methods, we first use WALDO [19], to generate a future frame. Then, we employ 3D Photo [35] to convert the future frame into a layered depth image (LDI), rendering it from a desired viewpoint through a constructed mesh. Additionally, we enhance 3D Photo [45] with the state-of-the-art monocular depth estimator ZoeDepth [2].

DMVFN [13] \rightarrow **3D Photo** [35]: Following the same procedure, we use DMVFN instead of WALDO.

Results. Quantitative results in Table 1 show that our approach consistently outperforms prior SOTA methods in error metrics. Notably, our method consistently improves SSIM and LPIPS over long-term forecasts. As depicted in Fig. 5, DMVFN exhibits stretch artifacts due to backward warping and inconsistent occlusion handling resulting from computing flow in the 2D domain. WALDO faces challenges with inaccuracies in layered decomposition, especially for objects moving in different directions but merged into the same layer, creating inconsistent artifacts. Our method effectively addresses these issues through the incorporation of 3D scene geometry and 3D motion modeling using a disentangled approach.

4.4 Video Prediction (VP)

In Table 4, we present a quantitative comparison with state-of-the-art (SOTA) video prediction approaches, using train/test splits from [1, 47]. Our approach surpasses SOTA methods across all error metrics.




Fig. 5: Qualitative Results for concurrent video extrapolation in space and time (VEST-[S+T]). (a) The DMVFN [13] \rightarrow 3D Photo [35] baseline produces stretching artifacts around the car due to disocclusions caused by the use of 2D flow-based backward warping. (b) In WALDO [19] \rightarrow 3D Photo [35], \rightarrow indicates the inconsistent motion of the road. This occurs due to the layered approach of WALDO [19], where the road and car are mistakenly assigned to the same layer and have similar motion. (c) Our approach mitigates these issues, achieving high-fidelity motion forecasting (indicated by \rightarrow).

4.5 Ablation Study

We conduct an ablation study on the KITTI dataset for the VEST-[S+T] task to validate the effectiveness of various proposed system components. Table 3 presents comparisons between our full system and variants: A) Using a fixed depth for both input images, determined by the optimal distance for 3D-2D splatting. Using a fixed depth transforms 3D motion forecasting task into 2D optical flow forecasting task; B) Omitting the ego-motion forecasting (EMF) module and using only the object motion forecasting (OMF) module for forecasting 3D motion; C) using only ego-motion forecasting (EMF) module without object motion forecasting (OMF) module. Without OMF module, synthesized future frame quality degrades as shown in Table 3. In the absence of EMF module, OMF module can partially compensate for it. However, having an explicit EMF module improves visual quality.

5 Discussion and Conclusion

Limitations. Our approach relies on physical depth estimates and faces challenges with inaccuracy in depth, particularly with thin structures and inconsistencies in depths across frames. These inaccuracies can lead to the forecasted results that feature unnatural distortions of thin structures, impacting overall realism of the forecasted frames. Utilizing RGB-D video can be a solution.

Conclusion. We present a novel approach for video extrapolation in space and time. By leveraging 3D scene geometry, our method disentangles 3D scene geometry from scene motion. It further disentangles the forecasting of ego-motion from the motion of specific objects, overcoming limitations found in recent approaches. Our approach shows significant improvements over prior state-of-the-art methods in forecasting future videos from novel views.

6 Acknowledgements

This material is based upon work supported under the AI Research Institutes program by National Science Foundation and the Institute of Education Sciences, U.S. Department of Education through Award #2229873 - AI Institute for Transforming Education for Children with Speech and Language Processing Challenges. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, the Institute of Education Sciences, or the U.S. Department of Education.

References

1. Bei, X., Yang, Y., Soatto, S.: Learning semantic-aware dynamics for video prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 902–912 (2021)
2. Bhat, S.F., Birkel, R., Wofk, D., Wonka, P., Müller, M.: Zoedepth: Zero-shot transfer by combining relative and metric depth. arXiv preprint arXiv:2302.12288 (2023)
3. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 801–818 (2018)
4. Clark, A., Donahue, J., Simonyan, K.: Adversarial video generation on complex datasets. arXiv preprint arXiv:1907.06571 (2019)
5. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
6. Denton, E., Fergus, R.: Stochastic video generation with a learned prior. In: International conference on machine learning. pp. 1174–1183. PMLR (2018)
7. Finn, C., Goodfellow, I., Levine, S.: Unsupervised learning for physical interaction through video prediction. *Advances in neural information processing systems* **29** (2016)
8. Franceschi, J.Y., Delasalles, E., Chen, M., Lamprier, S., Gallinari, P.: Stochastic latent residual video prediction. In: International Conference on Machine Learning. pp. 3233–3246. PMLR (2020)
9. 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)
10. Geng, D., Hamilton, M., Owens, A.: Comparing correspondences: Video prediction with correspondence-wise losses. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3365–3376 (2022)
11. Han, Y., Wang, R., Yang, J.: Single-view view synthesis in the wild with learned adaptive multiplane images. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–8 (2022)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
13. Hu, X., Huang, Z., Huang, A., Xu, J., Zhou, S.: A dynamic multi-scale voxel flow network for video prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6121–6131 (2023)

14. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. *Advances in neural information processing systems* **28** (2015)
15. Jampani, V., Chang, H., Sargent, K., Kar, A., Tucker, R., Krainin, M., Kaeser, D., Freeman, W.T., Salesin, D., Curless, B., et al.: Slide: Single image 3d photography with soft layering and depth-aware inpainting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12518–12527 (2021)
16. Khan, N., Xiao, L., Lanman, D.: Tiled multiplane images for practical 3d photography. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 10454–10464 (2023)
17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
18. Kumar, M., Babaeizadeh, M., Erhan, D., Finn, C., Levine, S., Dinh, L., Kingma, D.: Videoflow: A conditional flow-based model for stochastic video generation. *arXiv preprint arXiv:1903.01434* (2019)
19. Le Moing, G., Ponce, J., Schmid, C.: Waldo: Future video synthesis using object layer decomposition and parametric flow prediction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 23229–23241 (2023)
20. 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: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12578–12588 (2021)
21. Li, X., Cao, Z., Sun, H., Zhang, J., Xian, K., Lin, G.: 3d cinemagraphy from a single image. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4595–4605 (2023)
22. Li, X., Hong, C., Wang, Y., Cao, Z., Xian, K., Lin, G.: Symmnerf: Learning to explore symmetry prior for single-view view synthesis. In: *Proceedings of the Asian conference on computer vision*. pp. 1726–1742 (2022)
23. Li, Z., Wang, Q., Cole, F., Tucker, R., Snavely, N.: Dynibar: Neural dynamic image-based rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4273–4284 (2023)
24. Liu, L., Gu, J., Zaw Lin, K., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. *Advances in Neural Information Processing Systems* **33**, 15651–15663 (2020)
25. Liu, Z., Yeh, R.A., Tang, X., Liu, Y., Agarwala, A.: Video frame synthesis using deep voxel flow. In: *Proceedings of the IEEE international conference on computer vision*. pp. 4463–4471 (2017)
26. Lotter, W., Kreiman, G., Cox, D.: Deep predictive coding networks for video prediction and unsupervised learning. *arXiv preprint arXiv:1605.08104* (2016)
27. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
28. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 165–174 (2019)
29. Peng, J., Zhang, J., Luo, X., Lu, H., Xian, K., Cao, Z.: Mpib: An mpi-based bokeh rendering framework for realistic partial occlusion effects. In: *European Conference on Computer Vision*. pp. 590–607. Springer (2022)
30. Ranjan, A., Black, M.J.: Optical flow estimation using a spatial pyramid network. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4161–4170 (2017)

31. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3d deep learning with pytorch3d. arXiv preprint arXiv:2007.08501 (2020)
32. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
33. Saito, M., Saito, S., Koyama, M., Kobayashi, S.: Train sparsely, generate densely: Memory-efficient unsupervised training of high-resolution temporal gan. *International Journal of Computer Vision* **128**(10–11), 2586–2606 (2020)
34. Schmeckpeper, K., Georgakis, G., Daniilidis, K.: Object-centric video prediction without annotation. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 13604–13610. IEEE (2021)
35. Shih, M.L., Su, S.Y., Kopf, J., Huang, J.B.: 3d photography using context-aware layered depth inpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8028–8038 (2020)
36. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
37. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems* **32** (2019)
38. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8934–8943 (2018)
39. Tewari, A., Fried, O., Thies, J., Sitzmann, V., Lombardi, S., Sunkavalli, K., Martin-Brualla, R., Simon, T., Saragih, J., Nießner, M., et al.: State of the art on neural rendering. In: *Computer Graphics Forum*. vol. 39, pp. 701–727. Wiley Online Library (2020)
40. Tian, Y., Ren, J., Chai, M., Olszewski, K., Peng, X., Metaxas, D.N., Tulyakov, S.: A good image generator is what you need for high-resolution video synthesis. arXiv preprint arXiv:2104.15069 (2021)
41. Tucker, R., Snavely, N.: Single-view view synthesis with multiplane images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 551–560 (2020)
42. Tulsiani, S., Tucker, R., Snavely, N.: Layer-structured 3d scene inference via view synthesis. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 302–317 (2018)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
44. Villegas, R., Yang, J., Hong, S., Lin, X., Lee, H.: Decomposing motion and content for natural video sequence prediction. arXiv preprint arXiv:1706.08033 (2017)
45. Wang, Q., Li, Z., Salesin, D., Snavely, N., Curless, B., Kontkanen, J.: 3d moments from near-duplicate photos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3906–3915 (2022)
46. Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2021)
47. Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. arXiv preprint arXiv:1808.06601 (2018)

48. Wang, Y., Wu, H., Zhang, J., Gao, Z., Wang, J., Philip, S.Y., Long, M.: Predrnn: A recurrent neural network for spatiotemporal predictive learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(2), 2208–2225 (2022)
49. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*. vol. 2, pp. 1398–1402. Ieee (2003)
50. Wu, Y., Gao, R., Park, J., Chen, Q.: Future video synthesis with object motion prediction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5539–5548 (2020)
51. Wu, Y., Wen, Q., Chen, Q.: Optimizing video prediction via video frame interpolation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17814–17823 (2022)
52. Xu, D., Jiang, Y., Wang, P., Fan, Z., Shi, H., Wang, Z.: Sinnerf: Training neural radiance fields on complex scenes from a single image. In: *European Conference on Computer Vision*. pp. 736–753. Springer (2022)
53. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4578–4587 (2021)
54. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
55. Zhang, Y., Wu, J.: Video extrapolation in space and time. In: *European Conference on Computer Vision*. pp. 313–333. Springer (2022)