# E-NeRV: Expedite Neural Video Representation with Disentangled Spatial-Temporal Context

Zizhang Li<sup>\*</sup><sup>(0)</sup>, Mengmeng Wang<sup>\*</sup> <sup>(0)</sup>, Huaijin Pi<sup>(0)</sup>, Kechun Xu<sup>(0)</sup>, Jianbiao Mei<sup>(0)</sup>, and Yong Liu<sup>†</sup><sup>(0)</sup>

Zhejiang University {zzli,mengmengwang,hjpi,kcxu,jianbiaomei}@zju.edu.cn yongliu@iipc.zju.edu.cn

Abstract. Recently, the image-wise implicit neural representation of videos, NeRV, has gained popularity for its promising results and swift speed compared to regular pixel-wise implicit representations. However, the redundant parameters within the network structure can cause a large model size when scaling up for desirable performance. The key reason of this phenomenon is the coupled formulation of NeRV, which outputs the spatial and temporal information of video frames directly from the frame index input. In this paper, we propose E-NeRV, which dramatically expedites NeRV by decomposing the image-wise implicit neural representation into separate spatial and temporal context. Under the guidance of this new formulation, our model greatly reduces the redundant model parameters, while retaining the representation ability. We experimentally find that our method can improve the performance to a large extent with fewer parameters, resulting in a more than 8× faster speed on convergence. Code is available at https://github.com/kyleleey/E-NeRV.

**Keywords:** implicit representation, neural video representation, spatialtemporal disentanglement

# 1 Introduction

Implicit neural representation (INR) have become popular in recent days. It presents a new manner to represent continuous signals as  $f_{\theta} : \mathbb{R}^m \to \mathbb{R}^n$ , which encodes the signal property as a function that maps the *m*-dimensional input (e.g. coordinates) to desired *n*-dimensional output (e.g. RGB values, occupancy, density), and the function is parameterized by deep neural networks with weight  $\theta$ . Unlike regular grid-wise representations, the compact INRs are proved to be suitable for complex scenes [30] and arbitrary scale sampling [6], as well as in lots of 3D tasks [30,24,37,34] and image representations [43,6,27,59,57,40,64]. Despite the prevalence of INRs, few works have studied the compatible INR for video signals.

Video has been treated as an additional supplement of the image in past INR works [43,27]. They usually take the 3-dimensional spatial-temporal coordinate

<sup>\*:</sup> Equal contributions. †: Corresponding author.



Fig. 1. Main motivation of our proposed method. We can greatly reduce the size of the parameters by introducing disentangled spatial-temporal representations with a light network, while maintaining the majority of performance. Furthermore, we distribute the saved parameters to increase channel dimensions in convolution blocks, resulting in an E-NeRV model with similar or fewer parameters but much better performance.

(x, y, t) as input and output RGB values. Most of the following works [41,63] focusing on video INRs adopt this configuration. However, the training and inference speed of this type of video INR will increase by the order of the third power when processing the video sequences with large resolution and numerous frames. In contrast, a recently proposed method, NeRV [2], reformulates the INR of video signals as  $f_{\theta} : \mathbb{R} \to \mathbb{R}^{3 \times H \times W}$ . Based on the concept that video is a tile of images, NeRV presents an image-wise video INR different from other pixelwise video INRs. With the frame index in the time axis as input, NeRV directly outputs the desired frame image. The training and inference speed is proved to be much faster than previous methods [43,47] experimentally in [2]. And NeRV combines the success of convolution architecture and GAN's network design for its NeRV Blocks, which endows the ability to reconstruct frames of large resolution with high fidelity. By changing the channel dimensions in NeRV Blocks, we can obtain a series of NeRV models with different sizes (NeRV models with more parameters will naturally perform better). However, as the channel dimensions increase, the model size will increase rapidly. This drawback mainly comes from the architecture of NeRV model, which brings lots of unnecessary and redundant parameters ( $2 \times$  larger model size when channel dimensions increase 25%). We ascribe it to the design motivation of NeRV: NeRV considers the spatial and temporal information that lies in each frame image in a hybrid manner and is directly generated from one particular temporal frame index, which results in the heavy model and sub-optimal performance.

Inspired by the video GAN researches [50,14,62] that decompose the content and motion information, we propose the image-wise video INR that explicitly disentangles the spatial-temporal context and fuses them for final prediction, and refactor the original NeRV's network accordingly. Based on this motivation (illustrated in Fig. 1), we can effectively lower the parameter size of the baseline model (from 12M to 5M) while maintaining the majority of performance. We further introduce the temporal embeddings in convolutional blocks to facilitate the representation ability. Besides, we spot the redundant design lying in the NeRV Block and upgrade it. We name our method E-NeRV since it *Expedites* the original NeRV from a disentangled perspective for video implicit representation. We systematically investigate multiple design choices and compare our method with the baseline NeRV model. Our contributions are summarized as follows:

- We identify the redundant structures in the image-wise video INR NeRV, which is its major limitation when scaling up for better performance, and attribute this disadvantage to its hybrid formulation.
- We propose E-NeRV, a novel image-wise video INR with disentangled spatialtemporal context.
- We demonstrate our method can consistently outperform the NeRV baseline in convergence speed (8×) and performance with fewer parameters. Moreover, the superior performance is consistent among different video INR's downstream applications.

# 2 Related Work

### 2.1 Implicit Neural Representation

Recently, Implicit Neural Representation (INR) has gained much popularity for its strong power in modeling a variety of signals. It parameterizes a specific signal by a function that outputs desired properties of the provided coordinate-like input and employs deep neural networks (usually Multi-Layer Perceptron, MLP) to approximate the function. Thus the signal is implicitly encoded in network's parameters. For instance, the images [6,27,43] can be defined as RGB values of each pixel location, and 3D objects or scenes can be represented as occupancy [29,38], signed distance [35] or radiance field [30] of each 3D point. INRs are primarily popular in 3D vision tasks like reconstruction [22,33,52,36,39,19,56] and novel-view synthesis [30,61,1,60,45].

Implicit Representations of Videos have not been thoroughly studied in this trend. Regular video implicit representations often take the spatial and temporal index of a pixel, i.e.  $(x, y, t) \in \mathbf{R}^3$ , as input and output the RGB values of the certain pixel in the certain frame. This simple definition suits short video clips with small image sizes, like  $7 \times 224 \times 224$  in [43,27].[41] further estimates optical flow for continuous video representation. But this setting is no longer suitable for videos containing hundreds of frames with image resolution at large scale, which requires a long time to optimize and inference [2] because of the increasing number of frames and pixels. In addition, the paradigm proposed in [27] for contextual embeddings also can not support the videos with a large amount of frames. Another line of research of video INRs focuses on generative 4 Z. Li et al.

adversarial networks (GAN) [9]. Instead of generating videos from latent code directly, DiGAN [62] generates parameters of video INRs from context and motion latent code. StyleGAN-V [44] further utilizes convolutional operators for large-scale image synthesis. However, in this work we focus on fitting INR to the particular video instead of generating diverse contents in GAN-based methods.

Recently proposed NeRV [2] employs an image-wise implicit representation for videos instead of pixel-wise representations before. By combining the implicit representation with advances in convolution for image synthesis, NeRV achieves promising results with less time in training and inference. Following NeRV, our E-NeRV further improves the architecture via a disentangled formulation for a superior performance and fast convergence.

#### 2.2 Optimization of INRs

Despite the success of INRs' expression ability, they naturally cost a long time to optimize for considerable performance. Many methods have been proposed to alleviate this problem and also acquire a better representation ability.

From the perspective of function characteristics, researches can be divided into studying an optimal encoding method and applying network regularization. Given that INRs tend to learn better mapping functions with a higher dimensional network input, many following works focus on a better encoding approach. Radial basis function (RBF) [7] utilizes the weighted sums of embedded RBF encodings. Positional encoding (PE) proposed in Fourier Feature Networks (FFN) [47] employs a set of Fourier functions to project inputs into high dimensions, and follow-up works [21,13] adopt a coarse-to-fine strategy on frequency for better convergence. Different from using existing functions, SPE [53] uses learnable spline functions and the latest instant-ngp [31] constructs a trainable hash map for shared embedding space. As for the regularization, many consistency constraints [16,8,32] regarding the 3D property have been studied in view synthesis. [46] regularize the in-domain initialization by a meta-learning approach. The distribution-based [40] and Lipschitz-based [23] regularizations can be applied on MLP regarding the smoothness prior for better convergence and generalization.

From the network architecture perspective, some recent works aim to accelerate the training and/or inference of 3D INRs with delicately designed architecture regarding the 3D sparsity. A common approach is to store the feature of MLP inside a pre-defined voxel space [23], point cloud [58] or octree structure [60], thus reducing the numbers of point query in both training and inference. To a greater range, SIREN [43] replaces the commonly used RELU activations in existing MLPs with sinusoidal activation functions and shows the solid fitting ability to complex signals. ACORN [26] and CoordX [20] aim at reducing the number of queries to coordinate-based models with different approaches: ACORN [26] adopts a hierarchical way to decompose the multi-scale coordinates while CoordX [20] designs a split MLP architecture to leverage the locality between input coordinate points. The following MINER [42] improves ACORN via a cross-scale similarity prior.

Our work expedites NeRV from an architecture perspective as we observe the existing unnecessary and redundant structures. By introducing our disentangled formulation, we demonstrate that the resulting model with much fewer parameters can keep the majority of performance or even exceed the NeRV baseline. When scaling up to the same size with baseline NeRV model, our E-NeRV shows greater performance and faster convergence speed.

# **3** Preliminaries

NeRV [2], as an image-wise representation, represents the video as a mapping function  $f_{\theta} : \mathbb{R} \to \mathbb{R}^{3 \times H \times W}$  parameterized by the network weight  $\theta$ . Given a video with **T** frames  $\mathbf{V} = \{\mathbf{v}_t\}_{t=1}^{\mathbf{T}}$ , the input is a scalar frame index value which is normalized to  $\mathbf{t} \in [0, 1]$ , and the output is the whole corresponding frame image  $\mathbf{v}_t \in \mathbb{R}^{3 \times H \times W}$ . By taking a closer look at its architecture, the formulation can be split into two parts:

$$\mathbf{f_{t}} = \text{RESHAPE}\left(F\left(\gamma\left(\mathbf{t}\right)\right)\right) \in \mathbb{R}^{C \times h \times w},$$
  
$$\mathbf{v_{t}} = G\left(\mathbf{f_{t}}\right).$$
(1)

The  $\gamma$  (t) means the regular frequency positional encoding proposed in [30]:

$$\gamma(\mathbf{t}) = (\sin(b^0 \pi \mathbf{t}), \cos(b^0 \pi \mathbf{t}), \dots, \sin(b^{l-1} \pi \mathbf{t}), \cos(b^{l-1} \pi \mathbf{t})), \tag{2}$$

where b and l are hyper-parameters. The function F stands for the MLP while the function G stands for the convolutional generator. To be more specific, it contains a sequence of NeRV Blocks with convolution and pixel-shuffle layers for the up-sample and image generation purpose. The network first maps the positional encoding of input frame index to a 1-d feature vector and then reshapes the vector to a 2-d feature map  $\mathbf{f_t} \in \mathbb{R}^{C \times h \times w}$ , where (h, w) = (9, 16) in NeRV's setting. The following convolution and pixel-shuffle operation gradually transform the feature map to the original image size. And the 1 × 1 convolution with sigmoid generates the desired three channels of normalized RGB values.

The success of NeRV comes from several reasons. It employs an image-wise representation, which avoids per-pixel training and inference. The quantitative comparison in [2] shows great training and inference speed improvement compared to pixel-wise representations. The NeRV Block containing convolution and pixel-shuffle is suitable for image generation and leads to about 40 PSNR of final performance, superior to other video implicit neural representations.

A series of models with different sizes and performances are provided in [2]. A larger model can obtain better performance, and the way to scale up the model size is to increase the channel dimensions within the NeRV Blocks. However, this paradigm remains drawbacks. First is the last layer of MLP. To generate a feature vector that can be reshaped to feature map of size  $C \times h \times w$  (~ 10<sup>5</sup>), the last layer of MLP can be extensive, and some naive solutions will cause a large performance drop (details in Section. 5.4). Then, the convolution kernel can also be vast because of the following large-scale factor pixel-shuffle layer. NeRV considers



Fig. 2. Architecture of proposed E-NeRV. Our spatial-temporal feature map  $f_t$  is generated from disentangled spatial and input temporal contexts with fewer parameters (Section. 4.1). The temporal information is also introduced to the convolution stages as a normalization procedure (Section. 4.1) for better performance. In addition, we re-design the NeRV Blocks to further remove the redundant structures (Section. 4.2).

the image-wise video implicit representation as an index-to-image formulation, while we consider it as a generation process with disentangled formulation, and the frame index only represents the temporal context. In Section. 4, we elaborate our attempt to upgrade the redundant structure with spatial-temporal disentanglement, and we quantitatively and qualitatively show the significant performance and convergence speed of our method in Section. 5.

# 4 Methodology

6

Z. Li et al.

The overall architecture of the proposed E-NeRV is illustrated in Fig. 2. This section will introduce our approach towards the redundant parameters and structures. More specifically, in Section. 4.1 we state how to disentangle spatial and temporal representation and the resulting formulation and architecture. And in Section. 4.2 we elaborate on our upgraded design of NeRV block.

#### 4.1 Disentangled image-wise video INR

The first redundant part in NeRV emerges at the last layer of MLP. For instance, the NeRV-L model with 12.5M parameters, almost 70% of its size comes from the last MLP layer, which outputs  $\mathbf{f_t} \in \mathbb{R}^{112 \times 9 \times 16}$ . Although the height and width of the feature map are relatively small, it requires large channel numbers to guarantee the final performance. In experiments (Section. 5.4), we show some trivial modifications that may ease the large size of parameters, but lead to a dramatic performance drop compared to ours. We claim that this structure needs to exist because NeRV generates frame feature map  $\mathbf{f_t}$  directly and only from the input  $\mathbf{t}$ , which means to derive the spatial and temporal information together from the temporal input.

As an alternative, we propose to disentangle the spatial-temporal information, and let temporal input become a feature vector to manipulate over the spatial space. In detail, we reformulate the generation of  $\mathbf{f}_t$  as follows:

$$\mathbf{f}_{\mathbf{t}} = F_{\theta}(F(\gamma(\mathbf{t})) \odot \mathbf{S}). \tag{3}$$

Here F still stands for an MLP network, but with a much less parameter size because the output of F in our method is only a *d*-dimensional vector, where  $d \ll C \times h \times w$ . We decompose the spatial-temporal information into the temporal one encoded in  $F(\gamma(\mathbf{t}))$ , and the spatial one encoded in the spatial context embeddings  $\mathbf{S} \in \mathbb{R}^{d \times h \times w}$ . Then a lightweight network  $F_{\theta}$  is employed to fuse the separated spatial and temporal information into the spatial-temporal embeddings.

Since **S** is expected to contain the spatial context, we initialize it using the normalized grid coordinates. Accordingly we get the initialized spatial context  $\mathbf{S}_0 \in \mathbb{R}^{2 \times h \times w}$ . First, we encode the  $\mathbf{S}_0$  into  $\hat{\mathbf{S}}_0$  using similar frequency positional encoding  $\gamma(\cdot)$  in Eq. 2. Then, we adopt a small transformer [49] with single-head self-attention and residual connection here to encourage the feature fusion among spatial locations to get the spatial context  $\mathbf{S}$ :

$$\begin{aligned} \mathbf{S} &= \Phi(\hat{\mathbf{S}}_0) = \texttt{softmax}(\mathbf{q}^T \mathbf{k}) \mathbf{v} + \hat{\mathbf{S}}_0 \\ &= \texttt{softmax}(f_q(\hat{\mathbf{S}}_0)^T f_k(\hat{\mathbf{S}}_0)) f_v(\hat{\mathbf{S}}_0) + \hat{\mathbf{S}}_0. \end{aligned}$$
(4)

where  $f_*$  stands for different projection networks to project input feature map's channel dimension to desired dimension  $d_t$ .  $\mathbf{q}, \mathbf{k}, \mathbf{v}$  denote for the query, key and value of the transformer. Now the **S** can be considered as embeddings containing the desired spatial context. And when representing different videos, the learnable parameters in  $\Phi$  are different. In other words, we parameterize the spatial information in a video in the weights of  $\Phi$ .

Next, after the disentangled procedure, we need to fuse the temporal vector  $F(\gamma(\mathbf{t})) \in \mathbb{R}^d$  with the spatial context  $\mathbf{S} \in \mathbb{R}^{d \times h \times w}$  to obtain the spatialtemporal information. First, we element-wise multiply the temporal vector with each feature vector from all the locations in  $\mathbf{S}$ . Then, we utilize  $F_{\theta}$  to further fuse the features together.  $F_{\theta}$  here can be any operations as long as it can encourage spatial and channel feature fusion. We employ a tiny multi-head attention transformer network similar to  $\Phi$  for its ability of long-range modeling and feature fusion. In experiments, we further compare this choice with other alternatives (Section. 5.4).

Besides, we observe that temporal information in NeRV is only related to the feature map at the beginning of function G in Eq. 1. Therefore, we further fuse temporal context to each NeRV Blocks in G to make sufficient and thorough use of the temporal embedding. In experiments we find this design can further boost the performance. In detail, we take inspiration from the design of GAN [15], and consider the temporal context as a concept of style vector. Unlike using element-wise multiplication to get coarse spatial-temporal feature map, here the temporal information only plays a role of distribution shift. As illustrated in the upper part of Fig. 2, we adopt a tiny MLP (~ 0.2M) to generate temporal feature  $\mathbf{l}_t \in \mathbb{R}^{d_0}$ . Then for the *i*-th block ( $i = 1, \ldots, 5$ ), a linear layer  $M_i$  generates perchannel mean  $\mu_i$  and standard deviation  $\sigma_i$  accordingly. We denote the input

feature map for the *i*-th block as  $\mathbf{f}_{\mathbf{t}}^{i}$ . This newly generated distribution shifts the feature map as an instance normalization with temporal context:

$$IN(\mathbf{f}_{\mathbf{t}}^{i}) = \sigma_{i} \left( \frac{\mathbf{f}_{\mathbf{t}}^{i} - \mu(\mathbf{f}_{\mathbf{t}}^{i})}{\sigma(\mathbf{f}_{\mathbf{t}}^{i})} \right) + \mu_{i}$$
(5)

where  $\mu(\mathbf{f}_{\mathbf{t}}^{i})$  and  $\sigma(\mathbf{f}_{\mathbf{t}}^{i})$  are computed across spatial dimensions. This operation is conducted at the beginning of each block to let the temporal information guide the generation of the corresponding frame.

#### 4.2 Upgraded NeRV Block

As stated in Section 3, another redundant structure lies in NeRV block. Because the convolution needs to generate enough channels for further pixel-shuffle operation, if the input feature map's channel dimension is  $C_1$ , desired output dimension is  $C_2$ , the up-sample scale factor is s and kernel size is  $3 \times 3$ , regardless of the bias, the size of trainable weight is  $C_1 \times C_2 \times s \times s \times 3 \times 3$ . When scale factor s is large, for example, s = 5 in the first NeRV block, the size can be enormous (up to 65% of the overall model) if we scale up the input and output channel dimension for better performance.

In order to tackle this problem, we modify the NeRV block with a subtle design: we replace the convolution kernel with two consecutive convolution kernels with small channels. Then we place the pixel-shuffle operation in the middle and introduce an intermediate dimension  $C_0$ . By using  $conv(\cdot, \cdot)$  to denote convolution kernel with corresponding input and output channel dimensions, our new architecture can be formulated as:

$$\operatorname{conv}(C_1, C_0 \times s \times s) \to \operatorname{pixel-shuffle}(s) \to \operatorname{conv}(C_0, C_2),$$
 (6)

and the parameters in this new formula are:  $3 \times 3 \times C_0 \times (C_1 \times s \times s + C_2)$ .

In practice, we set  $C_0 = \min(C_1, C_2)/4$ . If  $C_1 \leq C_2$ , the ratio of the parameters size is  $(C_1/4C_2 + 1/4s^2) \approx C_1/4C_2 \leq 1/4$ . We find replacing the first NeRV Block with this design can bring largely simplified size while maintaining most of the performance (see Section. 5.2). The reason is that the scale factor of the first block equals 5 and thus results in an oversize model. The following blocks with factor equaling 2 will not benefit much from this modification, so in our final setting, we replace the first NeRV block with our upgraded version.

## 5 Experiments

#### 5.1 Datasets and Implementation Details

We conduct quantitative and qualitative comparison experiments on 8 different video sequences collected from scikit-video and UVG [28] datasets, similar to experiment settings in [2]. Each video sequence contains about 150 frames and with a resolution of  $1280 \times 720$ .



Fig. 3. Process of gradually transforming original 12.57M NeRV-L with 39.63 PSNR to our E-NeRV with slightly fewer parameters but much better performance. PSNR results are tested on "Bunny" video. Please see Section. 5.2 for detailed descriptions.

We set up-scale factors 5, 2, 2, 2, 2 for each block of our model to reconstruct a  $1280 \times 720$  image from the feature map of size  $16 \times 9$ . We follow the training schedule of the original NeRV implementation for a fair comparison. We train the model using Adam optimizer [17]. Each model is trained for 300 epochs on each video sequence unless specified, with the batchsize of 1.

We adopt NeRV-L with 12.57M parameters as our baseline. For the part in our model that is orthogonal to our modification, we follow the same settings as in NeRV, like activation choice. We set  $d = d_t = 256$  for spatial and temporal feature fusion,  $d_0 = 128$  for temporal instance normalization. We set all the positional encoding layers in our model identical to NeRV's positional encoding formulated in Eq. 2, and we use b = 1.25 and l = 80 if not otherwise denoted. For training objective, we use the same combination of L1 and SSIM loss as [2]:

$$L = \frac{1}{T} \sum_{t=1}^{T} \alpha ||\mathbf{v}_t - \hat{\mathbf{v}}_t||_1 + (1 - \alpha)(1 - \text{SSIM}(\mathbf{v}_t, \hat{\mathbf{v}}_t)).$$
(7)

The  $\alpha$  is set to 0.7, T stands for the total number of frames,  $\mathbf{v}_t$  denotes the reconstructed frame image while  $\hat{\mathbf{v}}_t$  denotes its corresponding ground truth. Please refer to the supplementary material for more implementation details, experiments, results and visualizations.

### 5.2 Process of Removing Redundant Part and Scaling Up

In this section, we show how to replace the redundant structures and parameters with our proposed methods, and gradually distribute the saved parameters to

#### 10 Z. Li et al.

**Table 1.** PSNR (larger the better) comparison between NeRV-L and our method given similar model size under the same training schedule. The last row indicates performance improvement brought by our method. Our method consistently outperforms the base-line model on diverse kinds of video sequences.

		Bunny	Beauty	Bosphorus	Bee	Jockey	$\operatorname{Set}\operatorname{Go}$	Shake	Yacht
NeRV-L	$12.57 \mathrm{M}$	39.63	36.06	37.35	41.23	38.14	31.86	37.22	32.45
Ours	12.49M	42.87	36.72	40.06	41.74	39.35	34.68	39.32	35.58
ouis		$\uparrow$ 3.24	↑ <b>0.66</b>	$\uparrow$ 2.71	$\uparrow$ 0.51	$\uparrow$ 1.21	$\uparrow$ 2.82	↑ <b>2.10</b>	↑ <b>3.13</b>

the channels in convolution stages, which leads to E-NeRV with fewer parameters but much better performance at last.

The overall process is shown in Fig. 3. We first replace the heavy MLP with our disentangled formulation in Eq. 3 and corresponding structure. This step can decrease the parameters from 12.57M to 5.5M, while the obtained model can still get 38.04 PSNR. As a comparison, NeRV-M model in [2] with more parameters can only reach a much worse performance of 36.05 PSNR. Then we first scale up the channels in convolution blocks for a model with a size similar to NeRV-L, and the model after scaling can get 41.70 PSNR.

After the first scaling, another redundant structure emerges: the NeRV Block with up-scale factor 5 and large channel dimension can be overwhelming, so we replace it with our new design. As shown in Fig. 3, the resulting model reduces 37% parameters. It is notable that the obtained model already has less parameters (7.92M vs. 12.57M) but better performance (40.61M vs. 39.63M) compared to origination NeRV-L. Then we scale up channels again and add temporal instance normalization branch at last for our proposed E-NeRV.

#### 5.3 Main Results

We provide a comparison of our method and NeRV in Table. 1. We refer to [2] for further comparison to pixel-wise video INRs such as SIREN [43] and FFN [30], which demonstrates that NeRV surpasses these methods in both performance and speed. Although our proposed E-NeRV has similar speed and parameters, it consistently performs better than the NeRV on different video sequences.

Because the design of our proposed E-NeRV does not employ any kinds of data prior, we claim this improvement exists when using E-NeRV to represent any video sequences. Notably, our method can bring larger promotion for the videos with more dynamic contents in Table 1, for example, the "Bunny" and "Yacht" videos. We assume this is because our disentangled implicit representation can better model the spatial and temporal variations for the videos with more dynamic contents.

Since training the INRs to fit a video sequence is an over-fitting process, longer schedule naturally leads to better performance. In other words, if the proposed method's performance surpasses another method with the same schedule,



**Fig. 4.** PSNR *vs.* Epochs. Comparison of NeRV-L and our method on "Bunny" and "Yacht" videos. Our method's performance within 300 epochs is better than NeRV's at 2400 epochs, which shows better performance and faster convergence.

it guarantees a better performance and faster convergence speed simultaneously. In Fig. 4, we provide a comparison between our method and NeRV on "Bunny" and "Yacht" videos with different training schedules. Our method's performance at 300 epochs exceeds the baseline's with a large margin. It also surpasses the baseline's performance at 2400 epochs, as  $8 \times$  faster on convergence. Actually, our method's performances at 300 epochs beat the baseline's at 2400 epochs on all the diverse videos. We provide the detailed results in supplementary.

#### 5.4 Comparison with Alternatives

We compare our method with four alternative approaches attempting to remove the redundant parameters or conduct fusion of  $F_{\theta}$  in Eq. 3:

**NeRV-** $C_S$ : Since the last layer of MLP with an output size of  $C \times h \times w$  causes overwhelming parameters, we add an intermediate channel dimension  $C_S$  which is lower than C. The MLP outputs feature map at size  $C_S \times h \times w$ , and a following  $1 \times 1$  convolution will increase the channel dimension to C as original setting before the NeRV Blocks.

**NeRV-Split:** Inspired by split architecture in [20], we redesign the MLP structure and let it output the tensor with size  $C \times (h + w)$ , then split it into two parts with size  $C \times h$  and  $C \times w$  respectively. A tensor product is conducted to generate the desired  $C \times h \times w$  feature map  $\mathbf{f}_t$  accordingly.

**E-NeRV-MLP:** Since the function  $F_{\theta}$  is responsible for feature fusion of spatial and temporal context, any fusion operation is applicable. We replace our original setting, a small transformer block with attention mechanism, with two successive MLPs on spatial channels  $(h \times w)$  and feature channels (C).

#### 12Z. Li et al.

	Size	$PSNR \uparrow$	MS-SSIM $\uparrow$	$\Phi F_{\theta}$ IN PSNR $\uparrow$ MS-SSIM $\uparrow$
NeRV- $C_S$	$5.8 \mathrm{M}$	33.72	0.9562	NeBV-L 36.74 0.9802
Ours-1†	$5.8 \mathrm{M}$	35.91	0.9738	
NeRV-Split	$7.2 \mathrm{M}$	35.78	0.9724	Variant 1 × × × 37.89 0.9837
Ours-2†	7.2M	36.74	0.9782	Variant 2 $\checkmark$ $\checkmark$ $\checkmark$ 38.09 0.9843
E-NeRV-MLF	<b>P</b> 12M	38.54	0.9861	Variant $3 \checkmark \checkmark X 38.45 = 0.0855$
E-NeRV-Conv	7 12.5M	38.67	0.9865	
E-NeRV	12.5M	38.79	0.9866	E-NeRV $\checkmark \checkmark \checkmark 38.79$ 0.9866

Table 2. Alternative comparison Table 3. Ablation study on components

**E-NeRV-Conv:** We use  $3 \times 3$  convolution block to replace the transformer block. The convolution block fuses the features within a window region and scans over the entire feature map in a sliding window manner.

The results are shown in Table. 2. For fair comparison on how to lower the size of the parameter, we establish two versions of our method: we remove the part of the structure to introduce the temporal context in convolution blocks stage as described in 4.1 since it can further boost the performance, and decrease the convolution's channel dimensions to make the resulting models' parameter sizes identical to the sizes of two alternatives. It can be seen that our method surpasses these alternatives given similar parameters settings.

As for the alternatives for feature fusion in  $F_{\theta}$ , transformer can bring incremental performance growth compared with MLP or Conv. However, all three models can beat NeRV-L with a large margin. The disentangled representation and structure itself can significantly lower the size, so that we can distribute the saved parameters to the convolution for much better performance. With the rapid growth of vision transformer research [10], any other more complicated structures, like the combination of transformer and convolution, are also welcomed and may further raise the performance. We claim that on some videos with almost still content, like "Beauty" and "Bee", the differences between each alternative are slight compared to the improvement on more dynamic videos. Since metrics are averaged over all videos, the difference between alternatives may also seem incremental in Table. 2, but the partial ordering relation is the same over 8 videos.

#### 5.5**Ablation Studies**

In this section, we study the effects of three novel components of our proposed method: the spatial fusion function  $\Phi$  at the beginning of the network, the spatial and temporal fusion  $F_{\theta}$  and the temporal instance normalization method to introduce temporal context in each convolution blocks. The ablation experiments are executed on all the video sequences and obtained metrics are averaged.

As shown in Table. 3, E-NeRV obtains better performance as gradually adding these modules, and this increasing property exists on all the experiment video sequences. It is notable that the "Variant 1", without the fusion and



Table 4. Denoising(left) and compression(right) results comparison of our method and NeRV.

temporal context in convolution stages, can still outperform baseline on both metrics among different video sequences. To be more specific, simply using the proposed disentanglement formulation to reduce the redundant parameters and distributing them to the following convolution blocks, the obtained model with similar parameters can already surpass the NeRV-L. We claim this further proves the effectiveness of our disentanglement motivation to some extent.

#### 5.6 Downstream application results

In addition to the representation ability, we also compare the E-NeRV's performance to NeRV's on different downstream tasks for video INR, including video denoising and compression. The results are shown in Table. 4.

Both experiments follow the NeRV's pipeline and we further conduct an ablation on different prune ratios for compression. The PSNR metrics are average among all the video sequences. In denoising results the "Noise" refers to noisy frames before any denoising. Here we only compare to NeRV since they beat other filter-based and learning-based methods in their paper. The denoising results of E-NeRV also prove the advantage of our disentangled spatial representation, which serves as a spatial prior in video denoising.

For the compression experiments, the performances of both methods drop as increasing the compression ratio (Sparsity in the figure), but E-NeRV remains better performance at all different compression ratios. The results also show that the compression ability of frame-wise video INR, i.e. the pipeline of pruning the network weight as compressing the video sequence, remains intact for the proposed E-NeRV. The detailed overall results can be found in the supplementary.

#### 5.7 Temporal frequency analysis

The frequency of Fourier feature mapping can greatly influence the INR's representation ability [47]. A minor frequency may lead to smoothness among input and suitable for interpolation but also degrade INR's fitting on training points.

In this section, we study the effect of different frequency in our disentangled representation. We devide the video at a ratio of 3:1 into seen and unseen



**Fig. 5.** Frequency variation for different encodings of: (a) NeRV's input  $\mathbf{t}$ , (b) our spatial map  $\mathbf{S}_0$ , (c) input  $\mathbf{t}$  in E-NeRV and (d)  $\mathbf{t}$  for IN.

frames, and adjust the frequency which is 1.25 in our general setting. The results are provided in Fig. 5.

Starting from NeRV's interpolation (39.3/28.58) at frequency 1.25, we can see that since NeRV consider spatial and temporal in a coupled manner, lowering the frequency can boost the interpolation but also cause a performance drop on seen frames (Fig. 5 (a)). On the contrary, our disentangled representation allows manipulating frequency in three encodings: spatial grid coordinates, temporal input **t** and **t** used in temporal IN. In detail, adjusting the frequency in IN module from 1.25 to 1.05 leads to the optimal interpolation while perserving the performance on training points (Fig. 5 (d)), which can be considered as another advantage from our disentangle structure. More dataset partition details, interpolation results and visualization are available in supplementary.

# 6 Conclusion

In this paper we present E-NeRV, an image-wise video implicit representation with disentangled spatial and temporal context. Following previous image-wise video INR [2], our method retains its advantages on training and inference speed compared to pixel-wise video INRs [47,43,27], but boosts the performance and convergence speed with a large margin. We quantitatively show that our proposed disentanglement structures together with other modifications can greatly reduce the original unnecessary and redundant parameters. By reallocating the saved parameters, our method with fewer parameters can perform much better, with an  $8 \times$  faster convergence speed. We experimentally analyze the function of each component in our method on diverse video sequences.

Finally, we remark our method can be further improved by applying a more effective and sophisticated feature fusion method for our disentangled representations. In future work, we plan to apply our image-wise video INR to other downstream tasks like optical flow estimation and video super-resolution.

Acknowledgments. We thank all authors and reviewers for the contributions. This work is supported by a Grant from the National Natural Science Foundation of China (No. U21A20484).

# References

- Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5855–5864 (2021)
- Chen, H., He, B., Wang, H., Ren, Y., Lim, S.N., Shrivastava, A.: Nerv: Neural representations for videos. Advances in Neural Information Processing Systems 34 (2021)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
- Chen, X., Ding, M., Wang, X., Xin, Y., Mo, S., Wang, Y., Han, S., Luo, P., Zeng, G., Wang, J.: Context autoencoder for self-supervised representation learning. arXiv preprint arXiv:2202.03026 (2022)
- Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021)
- Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8628–8638 (2021)
- Dash, C.S.K., Behera, A.K., Dehuri, S., Cho, S.B.: Radial basis function neural networks: a topical state-of-the-art survey. Open Computer Science 6(1), 33–63 (2016)
- Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised nerf: Fewer views and faster training for free. arXiv preprint arXiv:2107.02791 (2021)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems 27 (2014)
- Han, K., Wang, Y., Chen, H., Chen, X., Guo, J., Liu, Z., Tang, Y., Xiao, A., Xu, C., Xu, Y., et al.: A survey on vision transformer. arXiv preprint arXiv:2012.12556 (2020)
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16000–16009 (2022)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
- Hertz, A., Perel, O., Giryes, R., Sorkine-Hornung, O., Cohen-Or, D.: Sape: Spatially-adaptive progressive encoding for neural optimization. Advances in Neural Information Processing Systems 34 (2021)
- Hsieh, J.T., Liu, B., Huang, D.A., Fei-Fei, L.F., Niebles, J.C.: Learning to decompose and disentangle representations for video prediction. Advances in neural information processing systems **31** (2018)
- Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE international conference on computer vision. pp. 1501–1510 (2017)
- Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5885–5894 (2021)

- 16 Z. Li et al.
- 17. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Lai, Z., Liu, S., Efros, A.A., Wang, X.: Video autoencoder: self-supervised disentanglement of static 3d structure and motion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9730–9740 (2021)
- Li, Z., Niklaus, S., Snavely, N., Wang, O.: Neural scene flow fields for space-time view synthesis of dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6498–6508 (2021)
- 20. Liang, R., Sun, H., Vijaykumar, N.: Coordx: Accelerating implicit neural representation with a split mlp architecture. arXiv preprint arXiv:2201.12425 (2022)
- Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5741–5751 (2021)
- Littwin, G., Wolf, L.: Deep meta functionals for shape representation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1824– 1833 (2019)
- Liu, H.T.D., Williams, F., Jacobson, A., Fidler, S., Litany, O.: Learning smooth neural functions via lipschitz regularization. arXiv preprint arXiv:2202.08345 (2022)
- 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)
- Liu, Y., Wang, K., Liu, L., Lan, H., Lin, L.: Tcgl: Temporal contrastive graph for self-supervised video representation learning. IEEE Transactions on Image Processing **31**, 1978–1993 (2022)
- Martel, J.N., Lindell, D.B., Lin, C.Z., Chan, E.R., Monteiro, M., Wetzstein, G.: Acorn: Adaptive coordinate networks for neural scene representation. arXiv preprint arXiv:2105.02788 (2021)
- Mehta, I., Gharbi, M., Barnes, C., Shechtman, E., Ramamoorthi, R., Chandraker, M.: Modulated periodic activations for generalizable local functional representations. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14214–14223 (2021)
- Mercat, A., Viitanen, M., Vanne, J.: Uvg dataset: 50/120fps 4k sequences for video codec analysis and development. In: Proceedings of the 11th ACM Multimedia Systems Conference. pp. 297–302 (2020)
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4460– 4470 (2019)
- 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: European conference on computer vision. pp. 405–421. Springer (2020)
- 31. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. arXiv preprint arXiv:2201.05989 (2022)
- Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. arXiv preprint arXiv:2112.00724 (2021)
- Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3504–3515 (2020)

- Oechsle, M., Peng, S., Geiger, A.: Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5589–5599 (2021)
- 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)
- Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5865–5874 (2021)
- 37. Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9054–9063 (2021)
- Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: European Conference on Computer Vision. pp. 523–540. Springer (2020)
- Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10318–10327 (2021)
- 40. Ramasinghe, S., MacDonald, L., Lucey, S.: On regularizing coordinate-mlps. arXiv preprint arXiv:2202.00790 (2022)
- Rho, D., Cho, J., Ko, J.H., Park, E.: Neural residual flow fields for efficient video representations. arXiv preprint arXiv:2201.04329 (2022)
- Saragadam, V., Tan, J., Balakrishnan, G., Baraniuk, R.G., Veeraraghavan, A.: Miner: Multiscale implicit neural representations. arXiv preprint arXiv:2202.03532 (2022)
- Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. Advances in Neural Information Processing Systems 33, 7462–7473 (2020)
- 44. Skorokhodov, I., Tulyakov, S., Elhoseiny, M.: Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. arXiv preprint arXiv:2112.14683 (2021)
- Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P.P., Barron, J.T., Kretzschmar, H.: Block-nerf: Scalable large scene neural view synthesis. arXiv preprint arXiv:2202.05263 (2022)
- 46. Tancik, M., Mildenhall, B., Wang, T., Schmidt, D., Srinivasan, P.P., Barron, J.T., Ng, R.: Learned initializations for optimizing coordinate-based neural representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2846–2855 (2021)
- 47. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in Neural Information Processing Systems **33**, 7537–7547 (2020)
- Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. arXiv preprint arXiv:2203.12602 (2022)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)

- 18 Z. Li et al.
- 50. 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)
- Wang, G., Wang, K., Wang, G., Torr, P.H., Lin, L.: Solving inefficiency of selfsupervised representation learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9505–9515 (2021)
- Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689 (2021)
- Wang, P.S., Liu, Y., Yang, Y.Q., Tong, X.: Spline positional encoding for learning 3d implicit signed distance fields. arXiv preprint arXiv:2106.01553 (2021)
- Wang, X., Gupta, A.: Unsupervised learning of visual representations using videos. In: Proceedings of the IEEE international conference on computer vision. pp. 2794–2802 (2015)
- 55. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via nonparametric instance discrimination. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3733–3742 (2018)
- Xian, W., Huang, J.B., Kopf, J., Kim, C.: Space-time neural irradiance fields for free-viewpoint video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9421–9431 (2021)
- 57. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. arXiv preprint arXiv:2111.11426 (2021)
- 58. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. arXiv preprint arXiv:2201.08845 (2022)
- Xu, X., Wang, Z., Shi, H.: Ultrasr: Spatial encoding is a missing key for implicit image function-based arbitrary-scale super-resolution. arXiv preprint arXiv:2103.12716 (2021)
- Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenoctrees for real-time rendering of neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5752–5761 (2021)
- 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)
- 62. Yu, S., Tack, J., Mo, S., Kim, H., Kim, J., Ha, J.W., Shin, J.: Generating videos with dynamics-aware implicit generative adversarial networks. In: International Conference on Learning Representations (2021)
- 63. Zhang, Y., van Rozendaal, T., Brehmer, J., Nagel, M., Cohen, T.: Implicit neural video compression. arXiv preprint arXiv:2112.11312 (2021)
- Chuang, Y.: Filtering in implicit neural networks. arXiv preprint arXiv:2201.13013 (2022)