Scalable Neural Video Representations with Learnable Positional Features

Anonymous Author(s)
Affiliation
Address
email

Abstract

Succinct representation of complex signals using coordinate-based neural representations (CNRs) has seen great progress, and several recent efforts focus on extending them for handling videos. Here, the main challenge is how to (a) alleviate a compute-inefficiency in training CNRs to (b) achieve high-quality video encoding while (c) maintaining the parameter-efficiency. To meet all requirements (a), (b), and (c) simultaneously, we propose *neural video representations with learnable positional features* (NVP), a novel CNR by introducing "learnable positional features" that effectively amortize a video as latent codes. Specifically, we first present a CNR architecture based on designing 2D latent keyframes to learn the common video contents across each spatio-temporal axis, which dramatically improves the all of those three requirements. Then, we propose to utilize existing powerful image and video codecs as a compute-/memory-efficient compression procedure of latent codes. We demonstrate the superiority of NVP on the popular UVG benchmark; compared with prior arts, NVP not only trains 2 times faster (less than 5 minutes) but also exceeds their encoding quality as 34.00→34.43 (measured with the PSNR metric), even using >8 times fewer parameters. We also show intriguing properties of NVP, *e.g.*, video inpainting, video frame interpolation, etc.[1]

1 Introduction

Recent advances in coordinate-based neural representations (CNRs) [8, 12, 16, 37, 43] have shown great promise in the field as a new paradigm for representing complex signals, including gigapixel images [26, 31], audios [37], 3D scenes [27, 30, 32], and even large city-scale street views [44]. Instead of storing signal outputs as a coordinate grid (*e.g.*, image pixels), whose memory requirement scales unfavorably in terms of resolution and dimension, CNRs represent each signal as a compactly parameterized, continuous neural network; they interpret a signal as a coordinate-to-value function and train a neural network to approximate this mapping. CNRs enjoy numerous appealing properties, including data compression [10, 11, 54], super-resolution [5], novel view synthesis [17, 21, 30, 51], and generative modeling of complex, high-dimensional data [13, 23, 38, 39, 52, 55], while being parameter-efficient interpretation of a given signal in various scenarios.

In particular, several works have attempted to exploit CNRs to interpret *video signals* [4, 22, 39, 52] by learning a neural network $f : \mathbb{R}^3 \to \mathbb{R}^3$ with $f(x, y, t) = (r, g, b)$ and exhibited their potential as a succinct representation of videos, as well as providing numerous applications including video generative modeling [39, 52], video compression [4, 54], and video super-resolution [6]. They observe that conventional CNR architectures [37, 43] often fail to encode large-scale videos due to their complex temporal dynamics accompanied by large spatial variations, but the problem can be remarkably mitigated by designing CNR architectures specialized for videos. For instance, Chen

[1] Videos are available at the anonymized site https://neurips2022-nvp.github.io.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

Figure 1: Reconstruction results on Jockey in UVG-HD after training each model for a "1 minute" with a single NVIDIA V100 32GB GPU. NVP can capture the detail of a video containing dynamic motions, *e.g.*, the legs of a running horse, while the prior methods generate blurry artifacts.

et al. [4] proposes a CNR structure that focuses on the continuous modeling of the video signal only along the temporal dimension, allowing for more radical variations along spatial axes; it exhibits a comparable encoding quality to existing powerful video codecs (*e.g.*, H.264 [48], HEVC [40]) while enjoying lots of intriguing properties (*e.g.*, denoising and video frame interpolation).

However, CNRs suffer from a severe compute-inefficiency,[2] limiting their scalability to encode real-world, large-scale videos despite their advantages. To alleviate this issue, several works [24, 31, 34] have proposed new CNR architectures by separating a CNR f into two parts; $f = h_\phi \circ g_\theta$ for a coordinate-to-latent mapping $g_\theta(x, y, t) = \mathbf{z}$ and a latent-to-RGB mapping $h_\phi(\mathbf{z}) = (r, g, b)$. They construct g_θ as an embedding function defined with *latent grids* $\mathbf{U}_\theta$ in which the shape resembles the grid interpretation of a given signal (*e.g.*, a 2D array of C-dimensional latent codes $\mathbf{U}_\theta \in \mathbb{R}^{H \times W \times C}$ for image pixels) rather than as a neural network. These approaches have shown a promise in compute-efficiency due to the strong locality induced by a grid structure of $\mathbf{U}_\theta$; however, they result in another problem: these architectures severely sacrifice the parameter-efficiency since the parameter size of θ can be very large, growing proportionally to both the input coordinate dimension and the signal resolution. In this paper, we focus on developing video CNRs that are the best of both worlds: achieving high-quality encoding and the compute-/parameter-efficiency simultaneously.

Contribution. We introduce *neural video representations with learnable positional features* (NVP), a novel CNR for videos. NVP avoids requiring a single giant full-dimensional 3D array in g_θ by presenting *learnable positional features* that effectively amortize a given video as "2D and 3D" latent grids with succinct parameters. Specifically, we decompose the coordinate-to-latent mapping as

$$g_\theta = \underbrace{g_{\theta_{\texttt{xy}}} \times g_{\theta_{\texttt{xt}}} \times g_{\theta_{\texttt{yt}}}}_{\text{2D keyframes}} \times \underbrace{g_{\theta_{\texttt{xyt}}}}_{\text{3D sparse features}}, \qquad \text{for} \quad \theta := (\theta_{\texttt{xy}}, \theta_{\texttt{xt}}, \theta_{\texttt{yt}}, \theta_{\texttt{xyt}}),$$

where we present two types of latent grids for constructing these mappings (see Figure 2).

- *Latent keyframes*: We first design "image-like" 2D latent grids $\mathbf{U}_{\theta_{\texttt{xy}}}, \mathbf{U}_{\theta_{\texttt{xt}}}, \mathbf{U}_{\theta_{\texttt{yt}}}$ for $g_{\theta_{\texttt{xy}}}, g_{\theta_{\texttt{xt}}}, g_{\theta_{\texttt{yt}}}$ (respectively) that learn the representative video contents across *each* spatio-temporal axis and dramatically improve the parameter efficiency of NVP.[3]
- *Sparse positional features*: We then introduce a "video-like" 3D latent grid $\mathbf{U}_{\theta_{\texttt{xyt}}}$ for $g_{\theta_{\texttt{xyt}}}$, whose size is much smaller than the original video pixels, but effectively encodes video details locally.

Moreover, we propose a compute-/memory-efficient compression procedure to further reduce the parameter θ by incorporating existing image and video codecs, *e.g.*, JPEG [46] for images and HEVC [40] for videos, respectively. In particular, we treat 2D and 3D latent grids like the image or

[2] Measured with a single NVIDIA V100 32GB GPU, NeRV [4] takes at least 15 GPU hours to encode a single video of 600 frames with a 1920×1080 resolution for the desired quality.

[3] Such a spatio-temporal consideration is different from conventional approaches for specifying keyframes deterministically across only the temporal direction.

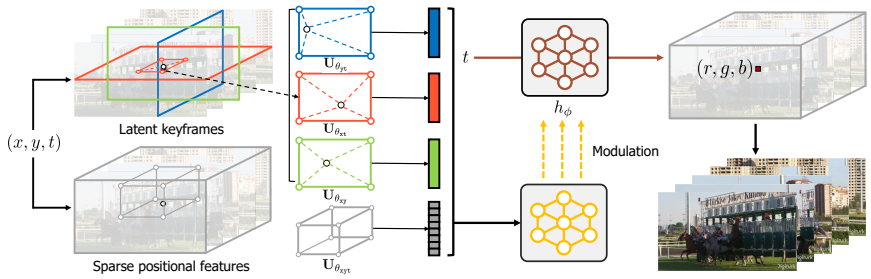

Figure 2: Overall illustration of our NVP. At a given space-time coordinate, NVP computes a latent vector from latent keyframes (see Figure 17 in Appendix for details) and sparse positional features. The latent vector is passed through a neural network to compute the corresponding RGB output.

video pixels and utilize powerful compression pipelines of them. Our compression scheme does not require any re-training of trained parameters, which significantly increases compute-efficiency than prior approaches to compressing CNR parameters but remarkably maintains the encoding quality. We also remark that such a compression approach is not applicable to the hashing-based latent grid of the prior method [31], while ours is "collision-free" and maintains the video- or image-like structures. Finally, for the choice of h_ϕ, we suggest using a modulated network (with respect to the temporal coordinate) to improve the encoding quality of videos that contain dynamic motions.

We verify the effectiveness of our method on the popular UVG benchmark [29]. In particular, NVP achieves the peak signal-to-noise ratio (PSNR; higher is better) metric of 34.43 in 5 minutes (with a single NVIDIA V100 32GB GPU): it is achieved $>$2 times faster, even with using $>$8 times fewer parameters than the state-of-the-art on compute-efficiency that reaches 34.00 in 10 minutes. Moreover, compared with prior arts on encoding quality, our method improves the learned perceptual image patch similarity (LPIPS [53]; lower is better) as 0.163$\rightarrow$0.131 ($+$19.6%) with a similar number of parameters while requiring $\sim$20% less training time. We also show numerous compelling properties of NVP, *e.g.*, video inpainting, video frame interpolation, super-resolution, compression, and consistent frame-wise encoding results without deviating quality.

2 Related work

Coordinate-based neural representations. Coordinate-based neural representations (CNRs), also known as implicit neural representations or neural fields, have emerged as a new paradigm for representing complex, continuous signals. They propose to encode signals through a neural network, typically a multilayer perceptron (MLP) combined with high-frequency sinusoidal activations [37] or Gaussian activations [8]. Prior works have focused on utilizing neural fields on various complicated data, *e.g.*, gigapixel images [26, 31], 2D videos [4, 39, 52], 3D static scenes, [27, 30, 32], and 3D dynamics scenes [22, 36, 49]. In particular, most approaches have focused on constructing CNR architectures for encoding 3D scenes [2, 14, 44] and exhibited how employing a specialized prior knowledge for a given signal domain in the architecture can remarkably boost the encoding quality. Despite these successes, extending CNRs for videos is yet under-explored. In this paper, we aim to move toward developing a CNR for videos by exploiting their unique temporal properties.

Hybrid CNRs. Rather than solely designing a neural network of coordinate-to-RGB mapping, hybrid CNRs incorporate learnable latent codes that follow a grid structure, *e.g.*, image CNRs combined with 2D latent spatial grids [5, 28], 3D scene (or shape) CNRs with latent cubic grids [3, 7, 18, 24, 26, 34]. Specifically, they compute a latent vector using the grid-structured latent code and pass it through a neural network to compute the signal output at a given input coordinate. Such approaches have shown significant efficiencies in training time and the encoding quality due to the powerful locality induced by grid-shaped latent codes. However, the number of parameters required for latent grid-based representations easily grows proportionally to the input coordinate dimension or data resolution, limiting the scalability of hybrid CNRs. Remarkably, some of the recent approaches have exhibited this inefficiency can be significantly mitigated by considering multi-level (or progressive) structures for latent grids [24, 26, 41, 42]. While prior works primarily focus on encoding images or 3D scenes, we aim to design parameter-efficient hybrid CNRs for videos.

3 NVP: Neural video representations with learnable positional features

We first formulate our problem setup as follows. Given a video signal $\mathbf{v} := (\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_T)$ consisting of T video frames, the goal is to find a compact neural representation $f_\mathbf{w}$ with parameters $\mathbf{w}$, from which the original video $\mathbf{v}$ can be reconstructed with high quality. Here, the quality can be defined using various distortion metrics, *e.g.*, peak signal-to-noise ratio (PSNR) [15] and LPIPS [53], for evaluating a reconstruction quality and a perceptual similarity, respectively.

To achieve this goal, we take an approach based on *coordinate-based neural representations* (CNRs)—a paradigm where each datum (*e.g.*, video) is parameterized as a neural network of a coordinate mapping. In particular, we aim to represent the given video using a neural network $f_\mathbf{w} : \mathbb{R}^3 \to \mathbb{R}^3$, which maps the space-time coordinates (x, y, t) of the video to corresponding RGB values (r, g, b), where $f_\mathbf{w}$ is optimized with reconstruction objectives, *e.g.*, mean-squared error. Such an approach has significant potential, as CNRs have shown to encode other continuous, complex signals (*e.g.*, 3D scenes) [37, 43] compactly while enjoying lots of intriguing properties, *e.g.*, super-resolution [5] and denoising [4]. However, CNRs have suffered from tremendous time costs for training, and even in the case of videos, it is difficult to achieve the high-quality encodings if one utilizes conventional CNR architectures that overlook the complex spatio-temporal dynamics of videos [4]. Our contribution lies in resolving these issues by designing "learnable positional features" that succinctly encode a video as latent codes with high quality and keeping their compute-/parameter-efficiency intact.

In the rest of this section, we provide a detailed description of each component in NVP. In Section 3.1, we explain the architecture of NVP. We then describe our compression procedure in Section 3.2.

3.1 Architecture

We design our video CNR $f_\mathbf{w}$ as a composition of two functions with a parameterization $\mathbf{w} := (\theta, \phi)$: a coordinate-to-latent mapping g_θ and a latent-to-RGB mapping h_ϕ. Here, we decompose the coordinate-to-latent-mapping as $g_\theta = g_{\theta_{\mathtt{xy}}} \times g_{\theta_{\mathtt{xt}}} \times g_{\theta_{\mathtt{yt}}} \times g_{\theta_{\mathtt{xyt}}}$ with $g_\theta(x, y, t) = (\mathbf{z}_{\mathtt{xy}}, \mathbf{z}_{\mathtt{xt}}, \mathbf{z}_{\mathtt{yt}}, \mathbf{z}_{\mathtt{xyt}})$ (for $\theta := (\theta_{\mathtt{xy}}, \theta_{\mathtt{xt}}, \theta_{\mathtt{yt}}, \theta_{\mathtt{xyt}})$), where each $g_{\theta_{\mathtt{xy}}}$, $g_{\theta_{\mathtt{xt}}}$, $g_{\theta_{\mathtt{yt}}}$ is formalized with image-like 2D latent spatial grids $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$, $\mathbf{U}_{\theta_{\mathtt{yt}}}$(respectively) and $g_{\theta_{\mathtt{xyt}}}$ is designed with a video-like sparse 3D latent grid $\mathbf{U}_{\theta_{\mathtt{xyt}}}$. We then present the latent-to-RGB mapping $h_\phi(\mathbf{z}_{\mathtt{xy}}, \mathbf{z}_{\mathtt{xy}}, \mathbf{z}_{\mathtt{yt}}, \mathbf{z}_{\mathtt{xyt}}) = (r, g, b)$ to be a multi-layer perceptrion (MLP) modulated by another neural network. To explain our architecture, we assume all the input coordinate (x, y, t) of g_θ (and $f_\mathbf{w}$) is in $[0, 1]^3 \subset \mathbb{R}^3$ without loss of generality.

Learnable latent keyframes. At a high level, learnable latent keyframes $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$, $\mathbf{U}_{\theta_{\mathtt{yt}}}$are image-like 2D latent grids learned to capture the common representative contents in a given video $\mathbf{v}$ across each t-, y-, x-axis, respectively. For a given input coordinate (x, y, t), we compute latent vectors $\mathbf{z}_{\mathtt{xy}}$, $\mathbf{z}_{\mathtt{xt}}$, $\mathbf{z}_{\mathtt{yt}}$ from $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$, $\mathbf{U}_{\theta_{\mathtt{yt}}}$individually. We explain our keyframe only with $\mathbf{U}_{\theta_{\mathtt{yt}}}$by letting $\mathbf{U} := \mathbf{U}_{\mathtt{yt}}$ for simplicity, but note that other keyframes $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$ operate in the similar manner.

Formally, $\mathbf{U}$ is L 2D spatial grids of C-dimensional latent codes u_{ij}, whose resolution is $H_l \times W_l$:

$$\begin{aligned} &\mathbf{U} := (U_1, \ldots, U_L), \\ &U_l := (u^l_{ij}) \in \mathbb{R}^{H_l \times W_l \times C} \quad \text{for } l = 1, \ldots, L, \\ &u_{ij} \in \mathbb{R}^C \qquad \text{for } 1 \le i \le H_l,\ 1 \le j \le W_l. \end{aligned}$$

Here, the keyframe follows an L-level multi-resolution structure, *i.e.*, for each level l, the height H_l and the width W_l become different as $H_l = \lfloor \gamma^{l-1} H_1 \rfloor$ and $W_l = \lfloor \gamma^{l-1} W_1 \rfloor$ with fixed $\gamma > 1$ and $H_1, W_1 > 0$, where $\lfloor \cdot \rfloor$ indicates the floor function of the input. Since $\mathbf{U} = \mathbf{U}_{\theta_{yt}}$ is shared over the x-axis, we compute the latent vector $\mathbf{z}_{\mathtt{yt}} := (z^1_{\mathtt{yt}}, \ldots, z^L_{\mathtt{yt}})$ by considering only the value of y and t at a given coordinate (x, y, t). Specifically, for $l = 1, \ldots, L$, each $z^l_{\mathtt{yt}}$ is a linearly interpolated vector of four vectors in the spatial grid U_l, where the indices of these vectors are chosen as the closest ones to the relative position of the input coordinate $(y, t) \in [0, 1]^2$:

$$\begin{aligned} (m, n) &= (\lfloor y H_l \rfloor, \lfloor t W_l \rfloor) \quad \text{for } l = 1, \ldots, L, \\ z^l_{\mathtt{yt}} &= \mathtt{lerp}\Big((y H_l - m, t W_l - n); u^l_{mn}, u^l_{m,n+1}, u^l_{m+1,n}, u^l_{m+1,n+1} \Big), \end{aligned}$$

where $\mathtt{lerp}$ indicates a linear interpolation operation at the input coordinate between given vectors.

Note that we *learn* the keyframe as the latent codes, unlike conventional approaches that specify the keyframe in a deterministic manner from T video frames $(\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_T)$; it encourages to capture

the representative contents of the video over each spatio-temporal direction better. We also remark that our architecture involves two keyframes $\mathbf{U}_{\theta_{\mathtt{xt}}}, \mathbf{U}_{\theta_{\mathtt{yt}}}$ that are considered across spatial axes. Considering these keyframes may not be beneficial in the RGB space as the spatial variation of video pixels is often large. However, in our approach, these keyframes are learned under a more flexible, continuous latent space, and thus the representative frames can even be found in these spatial directions in such a space while encoding the RGB outputs of the video accurately.

Finally, recall that $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$, $\mathbf{U}_{\theta_{\mathtt{yt}}}$consist of multi-resolution spatial grids, *i.e.*, the resolution of spatial grids in each latent keyframe grows from coarse to fine. Since natural scenes often include repeating patterns in various scales, *e.g.*, a scene of flowers of different sizes, this multi-resolution architecture promotes learning common patterns with a reduced memory and computation costs, which is also validated in Müller et al. [31].

Sparse positional features. Given a space-time coordinate (x, y, t), we compute the latent vector $\mathbf{z}_{\mathtt{xyt}}$with $\mathbf{U}_{\theta_{\mathtt{xyt}}}$that represents the local details of the video at this input position. Here, $\mathbf{U}_{\theta_{\mathtt{xyt}}} \coloneqq (u_{ijk}) \in \mathbb{R}^{H\times W\times S\times D}$ is a 3D *sparse* grid of D-dimensional latent codes, *i.e.*, the 3D grid size $H \times W \times S$ is much smaller than the size of the video pixels of a 3D RGB grid:

$$\mathbf{U}_{\theta_{\mathtt{xyt}}} \coloneqq (u_{ijk}) \in \mathbb{R}^{H\times W\times S\times D},\ u_{ijk} \in \mathbb{R}^D \quad \text{for } 1 \le i \le H,\ 1 \le j \le W,\ 1 \le k \le S.$$

To evaluate the latent vector $\mathbf{z}_{\mathtt{xyt}}$, we concatenate $h \times w \times s$ latent codes in $\mathbf{U}_{\theta_{\mathtt{xyt}}}$that their indices are near the relative position of the input (x, y, t):

$$\begin{aligned}(m, n, k) &= (\lfloor xH \rfloor, \lfloor yW \rfloor, \lfloor tS \rfloor),\\ \mathbf{z}_{\mathtt{xyt}} &= (u_{mnk}, \ldots, u_{m+h-1,n+w-1,k+s-1})\end{aligned}$$

where $h, w, s > 0$ are given as hyperparameters.

The locality of $\mathbf{U}_{\theta_{\mathtt{xyt}}}$as 3D latent codes dramatically alleviates the compute-inefficiency in fitting videos, in contrast to conventional CNRs where the entire parameters are shared (as a neural network) for arbitrary input coordinates (x, y, t) and thus require a significant training cost. Moreover, recall that we construct $\mathbf{U}_{\theta_{\mathtt{xyt}}}$as a sparse 3D grid of latent codes; remarkably, $\mathbf{U}_{\theta_{\mathtt{xyt}}}$efficiently captures the video details even if the size is smaller than the number of video pixels since the common contents of a given video are effectively encoded with the latent keyframes $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$, $\mathbf{U}_{\theta_{\mathtt{yt}}}$.

Modulated implicit function. With the latent vector $\mathbf{z} \coloneqq [\mathbf{z}_{\mathtt{xy}}, \mathbf{z}_{\mathtt{xt}}, \mathbf{z}_{\mathtt{yt}}, \mathbf{z}_{\mathtt{xyt}}]$ evaluated from g_θ, a naïve design choice of the latent-to-RGB mapping h_ϕ is to utilize a MLP that maps $\mathbf{z}$ to the corresponding RGB output (r, g, b). However, if a video contains temporally dynamic motions, we found such a simple MLP architecture lacks an expressive power and fails to capture the complex dynamics of the given video, even with the large network size of h_ϕ.

To circumvent this issue, we design h_ϕ to be a K-layer MLP (coined as a synthesizer network) modulated by another modulator network [28]: where the latent vector $\mathbf{z}$ and the time coordinate t are passed through the modulator and the synthesizer, respectively. Here, the modulator network utilizes piecewise linear activations (*e.g.*, ReLU), where the synthesizer uses sinusoidal activations. Specifically, an RGB output (r, g, b) of the latent vector $\mathbf{z}$ (from (x, y, t)) is computed as follows:

$$\begin{aligned}\boldsymbol{\alpha}_0 &= t,\\ \boldsymbol{\alpha}_k &= \mathbf{z}_k \odot \sin(\boldsymbol{A}_k \boldsymbol{\alpha}_{k-1} + \boldsymbol{b}_k) \quad \text{for } k = 1, \ldots, K-1,\\ (r, g, b) &= \boldsymbol{A}_K \boldsymbol{\alpha}_{K-1} + \boldsymbol{b}_K\end{aligned}$$

where $\boldsymbol{A}_k$, $\boldsymbol{b}_k$ are weights and biases of k-th layer of the synthesizer, $\boldsymbol{z}_k$ is k-th hidden feature of the modulator, and $\odot$ denotes an element-wise multiplication. Intriguingly, this modulation notably improves the encoding quality of videos that contains dynamic motions.

3.2 Compression procedure

Recall that we aim to find "compact" video CNRs; several works have focused on reducing the number of coordinate-based neural representations parameters (or bits) after training while maintaining their performance. In particular, they have relied on exploiting existing well-known techniques for neural network compression, *e.g.*, exploiting magnitude pruning [4, 20] or quantization [4, 54], and exhibited considerable results. However, these approaches mainly involve a re-training of CNR parameters, requiring severe computation costs, and thus are not suitable for practical scenarios.

Table 1: PSNR, FLIP, and LPIPS of different CNRs to encode videos in UVG-HD under each encoding time. ↑ and ↓ denote higher and lower values are better, respectively. Subscripts denote standard deviations, and bolds indicate the best results. * indicates applying the method without the corresponding compression scheme. We report the BPP values of NeRV without compressing parameters if the encoding time is $\leq$ 1 hour since the NeRV's compression requires a longer time. On the other hand, the compression procedure of NVP only takes less than 1 minute, but for a fair comparison, we do not apply it to NVP as well whenever NeRV is not compressed.

Encoding time	Method	BPP	PSNR (↑)	FLIP (↓)	LPIPS (↓)
∼5 minutes	Instant-ngp [31]	7.416	33.52±2.97	0.090±0.034	0.228±0.121
	NeRV-S* [4]	1.072	24.16±5.17	0.219±0.097	0.542±0.180
	NVP-S* (ours)	0.901	**34.43±2.80**	**0.075±0.022**	**0.187±0.093**
∼10 minutes	Instant-ngp [31]	7.416	34.00±2.55	0.090±0.029	0.200±0.109
	NeRV-S* [4]	1.072	26.53±5.92	0.176±0.088	0.460±0.184
	NVP-S* (ours)	0.901	**35.73±2.25**	**0.065±0.015**	**0.162±0.100**
∼1 hour	Instant-ngp [31]	7.416	36.14±2.50	0.070±0.023	0.157±0.091
	NeRV-S* [4]	1.072	33.26±4.31	0.094±0.038	0.240±0.112
	NVP-S* (ours)	0.901	**37.42±2.16**	**0.054±0.013**	**0.145±0.107**
∼15 hours	SIREN [37]	0.284	27.20±3.77	0.169±0.059	0.409±0.124
	FFN [43]	0.284	28.18±3.62	0.153±0.055	0.442±0.126
	Instant-ngp [31]	0.229	28.81±3.48	0.155±0.057	0.390±0.135
	NeRV-S [4]	0.201	36.14±3.97	**0.067±0.023**	0.163±0.101
∼12 hours	**NVP-S (ours)**	0.214	**36.34±2.19**	**0.067±0.017**	**0.128±0.073**
>40 hours	SIREN [37]	0.284	26.09±3.88	0.175±0.082	0.486±0.188
	FFN [43]	0.284	29.53±3.44	0.135±0.052	0.391±0.124
	Instant-ngp [31]	0.436	29.98±3.39	0.138±0.051	0.358±0.140
	NeRV-L [4]	0.485	35.00±3.31	0.079±0.020	0.145±0.100
∼17 hours	**NVP-L (ours)**	0.412	**37.71±1.80**	**0.061±0.013**	**0.110±0.085**

Instead, we propose a compression pipeline for NVP, which does not require re-training, yet significantly reduces the number of bits while preserving the video quality. Our main idea is to incorporate existing image and video codecs that have shown their promises for the compression of given pixels. In particular, we focus on compressing keyframes $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$, $\mathbf{U}_{\theta_{\mathtt{yt}}}$, and sparse positional features $\mathbf{U}_{\theta_{\mathtt{xyt}}}$, as the parameter size of the modulated implicit function h_ϕ is neglectable compared with those. Specifically, we quantize $\mathbf{U}_{\theta_{\mathtt{xyt}}}$ and $\mathbf{U}_{\theta_{\mathtt{xy}}}$, $\mathbf{U}_{\theta_{\mathtt{xt}}}$, $\mathbf{U}_{\theta_{\mathtt{yt}}}$ as 3D/2D grids of 8-bit latent codes and regard them as video and image pixel grids, where the number of the channel becomes the dimension of latent codes. Based on these interpretations, we compress these latent codes by utilizing existing video and image codecs, *e.g.*, HEVC [40] for videos and JPEG [46] for images. Intriguingly, we found this procedure can significantly reduce the parameters while notably maintaining the video quality without any fine-tuning of the latent-to-RGB mapping h_ϕ (See Section 4.3).

4 Experiments

We verify the effectiveness of our framework on UVG-HD [29], a representative benchmark for evaluating video encodings. Experimental results demonstrate that our neural video representations with learnable positional features (NVP) simultaneously improves the overall performance by (a) alleviating a compute-inefficiency in training, (b) achieving high-quality video encoding, and (c) maintaining parameter-efficiency. We also show applications of our NVP, including video inpainting and spatio-temporal interpolation. Finally, we conduct ablation studies to validate each component.

Evaluation. We follow similar setups on prior work [4] proposing coordinate-based neural representations (CNRs) for videos. All reported numbers are averaged over 7 videos in UVG-HD: Beauty, Bosphorus, HoneyBee, Jockey, ReadySetGo, ShakeNDry, and Yachtride, along with the standard deviations unless otherwise specified. We also use the Big Buck Bunny video, which is used in NeRV [4]. For quantitative evaluation, we use the following metrics: peak signal-to-noise ratio (PSNR) for reconstruction quality, LPIPS [53], FLIP [1], and SSIM [47] for a perceptual similarity, where all metrics are evaluated in a frame-wise manner and averaged over the whole video. We evaluate these metrics on different bits-per-pixel (BPP; lower is better) to evaluate the parameter efficiency; see Appendix A.1 for more description of the evaluation.

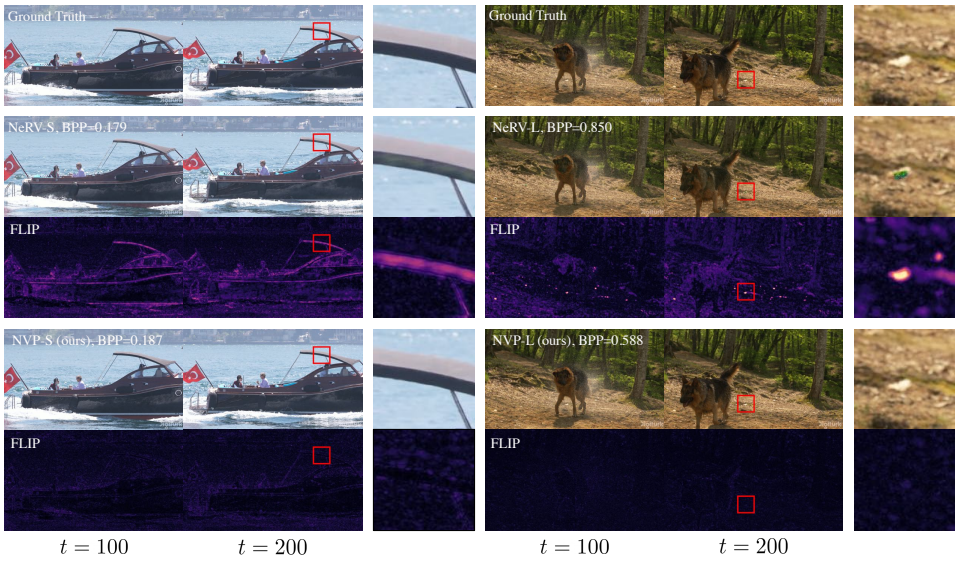

Figure 3: Illustration of reconstructions on Yachtride (left), and ShakeNDry (right) in UVG-HD. FLIP indicates the output of evaluating its metric. The red box is zoomed-in as the image at the right.

Implementation details. All main experiments, including baselines, are processed with a single GPU (NVIDIA V100 32GB) and 28 instances from a virtual CPU (Intel® Xeon® Platinum 8168 CPU @ 2.70GHz), where it takes at most ∼17 hours to run our method and ∼3 days to run other baselines. Moreover, we denote NVP-S and NVP-L as the model with latent code dimensions of sparse positional features to be 2 and 4, respectively; see Appendix A.2 for more details.

Baselines. We compare our method with SIREN [37], FFN [43], which are well-known signal-agnostic CNR architectures, Instant-ngp [31] for state-of-the-art CNRs on compute-efficiency, and NeRV [4], which is a CNR specialized for videos. For all of the baseline methods, we follow their reported experimental setups. In particular, for NeRV [4], we use two configurations provided in the official implementation: NeRV-S and NeRV-L for a small and a large model, respectively. See Appendix B for a detailed description of baseline methods.

4.1 Main results

Figure 1, Figure 3, and Table 1 summarize quantitative and qualitative results of NVP and baselines. Remarkably, as shown in Figure 1, NVP can accurately capture dynamic motions and high-frequency details of a video *e.g.*, the legs of a running horse, while prior state-of-the-art CNRs fail to achieve and shows a blurry artifact. We also emphasize such a high-quality encoding is accomplished in "less than 1 minute", which supports the superior compute-efficiency of NVP in training.

Moreover, Table 1 verifies the effectiveness of NVP with quantitative evaluations, which outperforms all other baselines at varying encoding times from ∼1 minute to >40 hours. In particular, NVP significantly improves LPIPS compared with previous methods, demonstrating how the encoded videos are perceptually similar to ground-truth videos. Such a result is also confirmed in Figure 3; NeRV shows a distortion that some pixels significantly deviate from the ground-truth outputs, while our method does not suffer from such artifacts. We also note that the variance of NVP is relatively small compared with other baselines, which shows the robustness of videos with diverse scenes and motion. See Appendix C and D for video-wise results and discussion on decoding time, respectively.

4.2 Applications

In this section, we provide several applications of our method, NVP, as video CNRs. For better, playable illustrations and qualitative results, please refer to our project page.

Video inpainting. Intriguingly, our method has the capability of video inpainting, *i.e.*, the desired moving object in the video can be naturally removed by capturing shared video contents with learnable keyframes. Figure 4 visualizes the illustration of inpainting results from NVP; as shown in this figure, one can see the car is removed where such parts are filled with a natural background.

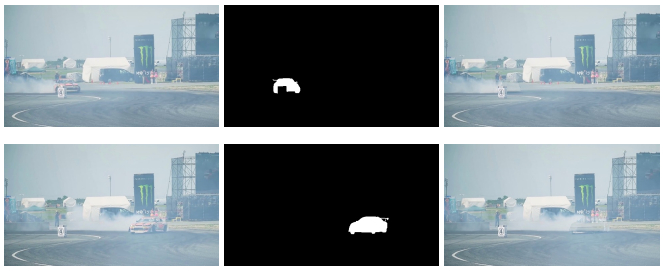

Figure 4: Video inpainting result of NVP on the drift-chicane video in DAVIS 2017 [35] to remove the masked car.

Table 2: Quantitative interpolation result of different methods on Big Buck Bunny measured with PSNR, LPIPS, and SSIM metrics. Bold indicates the best result.

Metric	NeRV [4]	NVP (ours)
PSNR (↑)	23.05	**33.76**
LPIPS (↓)	0.480	**0.311**
SSIM (↑)	0.690	**0.960**

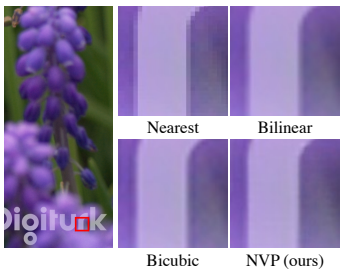

Figure 5: Super-resolution result (×8) of NVP (HoneyBee).

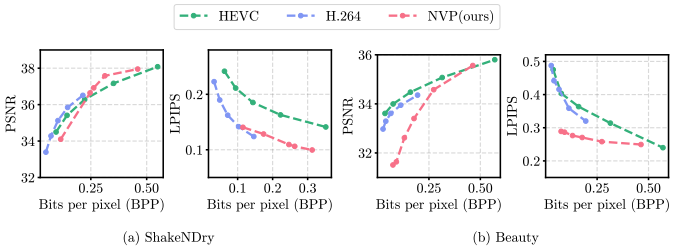

Figure 6: PSNR and LPIPS values of NVP and well-known video codecs over different BPP values computed on (a) the ShakeNDry video and (b) the Beauty video in UVG-HD.

Video frame interpolation. Since video CNRs approximate videos as temporally continuous signals, they should interpolate among two different frames at an arbitrary time even if such a frame does not exist in the training dataset. To validate the interpolation capability of NVP, we provide the quantitative results on Big Buck Bunny sequences; our method shows better interpolation results measured with various metrics.

Video super-resolution. We remark that NVP encodes a given video as a *spatio-temporally* continuous signal. Thus, our method can interpolate the frames across spatial directions, *i.e.*, frame-wise super-resolution. Figure 5 exhibits how well NVP smoothly interpolates in video frames across spatial directions while preserving the sharp edges, compared with naïve upsampling methods.

Video compression. Recall that one of the major advantages of CNRs is their succinct encoding to represent a given signal; one may consider utilizing CNRs for video compression [4, 10, 11]. To verify the potential of our method on video compression, we compare the quality of compressed videos from NVP with the ones from the current state-of-the-art video codecs. As shown in Figure 6, compressed videos from NVP show the comparable reconstruction quality (measured with PSNR metrics) while outperforming perceptual similarity (measured with LPIPS metrics).

4.3 Ablation studies

Effect of architecture components. To verify the effectiveness of each component, we train our model with the Jockey video by removing each component while maintaining the total number of parameters, then measure PSNR metrics from these models. Table 3 summarizes the effect of three different architecture components. Without any of the components consisting of our positional features, the reconstruction quality gets dramatically worse, which validates how NVP succinctly encodes a given video as latent codes. We also note that the Jockey video contains a dynamic motion of a running horse; the modulation indeed improves the encoding of videos with complex motions.

Analysis of non-temporal keyframes. Recall that we additionally design keyframes across spatial directions, unlike conventional approaches that designate the keyframe over only the temporal direction. To validate the effect of such keyframes, we compare NVP with (a) NVP without spatial keyframes and (b) NVP without all of the keyframes in Figure 8, where the number of parameters is equally set for a fair comparison. While utilizing latent learnable keyframes only over temporal direction is already fairly effective for high-quality encoding, one can observe the consideration of keyframes across other directions provides a further improvement.

Table 3: Ablation study of each component of NVP: learnable keyframes, sparse feature, and modulation.

Keyframes	Sparse feat.	Module.	# Params.	PSNR (↑)
✗	✓	✓	136M	31.21±2.80
✓	✗	✓	138M	32.44±4.48
✓	✓	✗	147M	38.04±2.27
✓	✓	✓	136M	**38.89±2.11**

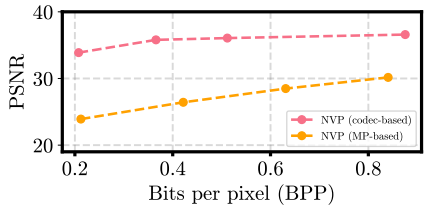

Figure 7: Rate-distortion plot of different compression strategies on the ReadySetGo video.

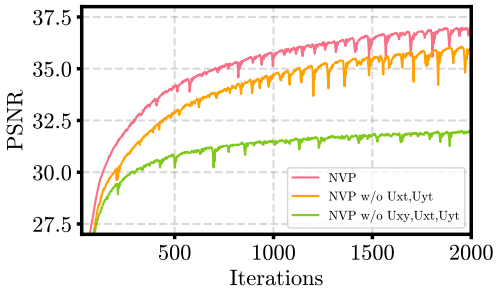

Figure 8: Convergence plot of NVP under different keyframe choices on the Jockey video.

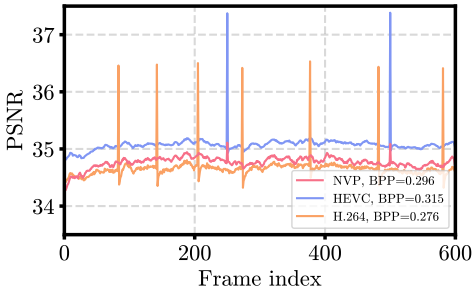

Figure 9: Frame-wise encoding quality of NVP and existing video codecs on Beauty.

Compression procedure. To validate the effectiveness of our compression scheme, we compare the encoding quality of (a) NVP with our compression scheme and (b) NVP compressed through magnitude-based pruning. Figure 7 shows the proposed compression pipeline outperforms conventional magnitude-based pruning under various BPP. We also remark that our compression method does not require re-training and is thus much more time-efficient.

Consistent frame-wise encoding quality. Existing keyframe-based compression approaches often suffer from inconsistent encoding quality: the frame-wise quality of the compressed video highly depends on whether it is the designated keyframe or not. In contrast, NVP *learns* the keyframes and does not have this problem. Figure 9 validates the result: our method exhibits consistent encoding performance while conventional popular video codecs show several peaks that the reconstruction quality (measured with PSNR metrics) highly deviates from others.

5 Discussion and conclusion

We proposed NVP, a new coordinate-based neural representation (CNR) to encode videos as succinct latent codes. Our main idea is to decompose a video into "image-like" and "video-like" structures to learn a coordinate-to-latent mapping efficiently. Extensive experiments have verified the effectiveness of NVP on all of parameter-/compute-efficiency and the encoding quality. We hope our method will facilitate various future research directions in the CNR area.

Limitations and future works. Each video contains different scenes and motions so that it can be either static or dynamic, yet we utilize the same hyperparameters and architectures for encoding any video. Although such a video-agnostic design is fairly effective and outperforms prior works, we believe the video-wise consideration of the architecture and hyperparameter can remarkably boost the performance further. Moreover, we have shown the potential of utilizing powerful image and video codes for compressing latent codes in NVP; extending such codecs to be specialized for the compression of latent codes should be an interesting direction.

Negative social impacts. A side effect of CNRs is their potential unexpected behavior on encoding; they may cause undesirable artifacts in representing the given signal but are challenging to predict due to the under-explored behavior of training CNRs. Furthermore, in the case of representing videos, the encoded videos may suffer from severe distortions and conceivably cause ethical problems. In this respect, such behaviors should be extensively and carefully investigated and mitigated to exploit CNRs as the standard for encoding videos in real-world situations.

References

[1] P. Andersson, J. Nilsson, T. Akenine-Möller, M. Oskarsson, K. Åström, and M. D. Fairchild. FLIP: A Difference Evaluator for Alternating Images. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 3(2):15:1–15:23, 2020.

[2] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. In *IEEE International Conference on Computer Vision*, pages 5855–5864, 2021.

[3] R. Chabra, J. E. Lenssen, E. Ilg, T. Schmidt, J. Straub, S. Lovegrove, and R. Newcombe. Deep local shapes: Learning local SDF priors for detailed 3d reconstruction. In *ECCV*, pages 608–625. Springer, 2020.

[4] H. Chen, B. He, H. Wang, Y. Ren, S. N. Lim, and A. Shrivastava. NeRV: Neural representations for videos. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

[5] Y. Chen, S. Liu, and X. Wang. Learning continuous image representation with local implicit image function. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[6] Z. Chen, Y. Chen, J. Liu, X. Xu, V. Goel, Z. Wang, H. Shi, and X. Wang. VideoINR: Learning video implicit neural representation for continuous space-time super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

[7] J. Chibane, T. Alldieck, and G. Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *CVPR*, pages 6970–6981, 2020.

[8] S.-F. Chng, S. Ramasinghe, J. Sherrah, and S. Lucey. GARF: Gaussian activated radiance fields for high fidelity reconstruction and pose estimation. *arXiv e-prints*, pages arXiv–2204, 2022.

[9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.

[10] E. Dupont, A. Golinski, M. Alizadeh, Y. W. Teh, and A. Doucet. COIN: COmpression with Implicit Neural representations. In *Neural Compression: From Information Theory to Applications – Workshop @ ICLR 2021*, 2021.

[11] E. Dupont, H. Loya, M. Alizadeh, A. Goliński, Y. W. Teh, and A. Doucet. COIN++: Data agnostic neural compression, 2022.

[12] R. Fathony, A. K. Sahu, D. Willmott, and J. Z. Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2021.

[13] J. Gu, L. Liu, P. Wang, and C. Theobalt. StyleNeRF: A style-based 3d aware generator for high-resolution image synthesis. In *International Conference on Learning Representations*, 2022.

[14] P. Hedman, P. P. Srinivasan, B. Mildenhall, J. T. Barron, and P. Debevec. Baking neural radiance fields for real-time view synthesis. In *IEEE International Conference on Computer Vision*, pages 5875–5884, 2021.

[15] A. Hore and D. Ziou. Image quality metrics: PSNR vs. SSIM. In *2010 20th international conference on pattern recognition*, pages 2366–2369. IEEE, 2010.

[16] Z. Huang, S. Bai, and J. Z. Kolter. Implicit^2: Implicit layers for implicit representations. In *Advances in Neural Information Processing Systems*, 2021.

[17] A. Jain, M. Tancik, and P. Abbeel. Putting NeRF on a diet: Semantically consistent few-shot view synthesis. In *IEEE International Conference on Computer Vision*, 2021.

[18] C. Jiang, A. Sud, A. Makadia, J. Huang, M. Nießner, T. Funkhouser, et al. Local implicit grid representations for 3D scenes. In *CVPR*, pages 6001–6010, 2020.

[19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.

[20] J. Lee, J. Tack, N. Lee, and J. Shin. Meta-learning sparse implicit neural representations. In *Advances in Neural Information Processing Systems*, 2021.

[21] T. Li, M. Slavcheva, M. Zollhoefer, S. Green, C. Lassner, C. Kim, T. Schmidt, S. Lovegrove, M. Goesele, and Z. Lv. Neural 3D video synthesis from multi-view video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

[22] Z. Li, S. Niklaus, N. Snavely, and O. Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[23] C. H. Lin, H.-Y. Lee, Y.-C. Cheng, S. Tulyakov, and M.-H. Yang. InfinityGAN: Towards infinite-pixel image synthesis. In *International Conference on Learning Representations*, 2022.

[24] L. Liu, J. Gu, K. Zaw Lin, T.-S. Chua, and C. Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems*, volume 33, pages 15651–15663, 2020.

[25] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

[26] J. N. Martel, D. B. Lindell, C. Z. Lin, E. R. Chan, M. Monteiro, and G. Wetzstein. ACORN: Adaptive coordinate networks for neural representation. In *ACM Trans. Graph. (SIGGRAPH)*, 2021.

[27] R. Martin-Brualla, N. Radwan, M. S. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth. NeRF in the wild: Neural radiance fields for unconstrained photo collections. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[28] I. Mehta, M. Gharbi, C. Barnes, E. Shechtman, R. Ramamoorthi, and M. Chandraker. Modulated periodic activations for generalizable local functional representations. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 14214–14223, 2021.

[29] A. Mercat, M. Viitanen, and J. Vanne. UVG dataset: 50/120fps 4K sequences for video codec analysis and development. In *Proceedings of the 11th ACM Multimedia Systems Conference*, MMSys '20, page 297–302, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450368452. doi: 10.1145/3339825.3394937.

[30] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, 2020.

[31] T. Müller, A. Evans, C. Schied, and A. Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. doi: 10.1145/3528223.3530127.

[32] K. Park, U. Sinha, J. T. Barron, S. Bouaziz, D. B. Goldman, S. M. Seitz, and R. Martin-Brualla. Deformable neural radiance fields. In *IEEE International Conference on Computer Vision*, 2021.

[33] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

[34] S. Peng, M. Niemeyer, L. Mescheder, M. Pollefeys, and A. Geiger. Convolutional occupancy networks. In *ECCV*, pages 523–540. Springer, 2020.

[35] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 DAVIS challenge on video object segmentation. *arXiv:1704.00675*, 2017.

[36] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[37] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems*, 2020.

[38] I. Skorokhodov, S. Ignatyev, and M. Elhoseiny. Adversarial generation of continuous images. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[39] I. Skorokhodov, S. Tulyakov, and M. Elhoseiny. StyleGAN-V: A continuous video generator with the price, image quality and perks of StyleGAN2. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

[40] G. J. Sullivan, J.-R. Ohm, W.-J. Han, and T. Wiegand. Overview of the high efficiency video coding (HEVC) standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(12):1649–1668, 2012. doi: 10.1109/TCSVT.2012.2221191.

[41] C. Sun, M. Sun, and H.-T. Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

[42] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezahrai, A. Jacobson, M. McGuire, and S. Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[43] M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, 2020.

[44] M. Tancik, V. Casser, X. Yan, S. Pradhan, B. Mildenhall, P. P. Srinivasan, J. T. Barron, and H. Kretzschmar. Block-NeRF: Scalable large scene neural view synthesis. *arXiv preprint arXiv:2202.05263*, 2022.

[45] S. Tomar. Converting video formats with ffmpeg. *Linux journal*, 2006(146):10, 2006.

[46] G. K. Wallace. The JPEG still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992.

[47] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

[48] T. Wiegand, G. Sullivan, G. Bjontegaard, and A. Luthra. Overview of the H.264/AVC video coding standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(7):560–576, 2003. doi: 10.1109/TCSVT.2003.815165.

[49] W. Xian, J.-B. Huang, J. Kopf, and C. Kim. Space-time neural irradiance fields for free-viewpoint video. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[50] B. Xu, N. Wang, T. Chen, and M. Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.

[51] D. Xu, Y. Jiang, P. Wang, Z. Fan, H. Shi, and Z. Wang. SinNeRF: Training neural radiance fields on complex scenes from a single image. *arXiv preprint arXiv:2204.00928*, 2022.

[52] S. Yu, J. Tack, S. Mo, H. Kim, J. Kim, J.-W. Ha, and J. Shin. Generating videos with dynamics-aware implicit generative adversarial networks. In *International Conference on Learning Representations*, 2022.

[53] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[54] Y. Zhang, T. van Rozendaal, J. Brehmer, M. Nagel, and T. Cohen. Implicit neural video compression. In *ICLR Workshop on Deep Generative Models for Highly Structured Data*, 2022.

[55] P. Zhou, L. Xie, B. Ni, and Q. Tian. CIPS-3D: A 3D-Aware Generator of GANs Based on Conditionally-Independent Pixel Synthesis. arXiv preprint 2110.09788, 2021.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section 4
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 5
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 5.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section 4 and supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4 and supplementary material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 4.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 4.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Section 4.
   (b) Did you mention the license of the assets? [Yes] See Section 4.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] See the supplementary material.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] See the supplementary material.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]