

Spline Deformation Field

MINGYANG SONG, DisneyResearch|Studios, Switzerland and ETH Zürich, Switzerland

YANG ZHANG, DisneyResearch|Studios, Switzerland

MARKO MIHAJLOVIC, ETH Zürich, Switzerland

SIYU TANG, ETH Zürich, Switzerland

MARKUS GROSS, DisneyResearch|Studios, Switzerland and ETH Zürich, Switzerland

TUNÇ OZAN AYDIN, DisneyResearch|Studios, Switzerland




Fig. 1. We represent trajectories with splines, in which analytically derived velocity v enables effective spatial coherency preservation through our novel loss term \mathcal{L}_v . Furthermore, we propose a metric concepting on Moran's I to quantify it. Spline interpolation can generate fair interpolated and extended motions, compared to purely relying on the smoothness inductive bias of the implicit deformation field, such as ResFields [Mihajlovic et al. 2024].

Trajectory modeling of dense points usually employs implicit deformation fields, represented as neural networks that map coordinates to relate canonical spatial positions to temporal offsets. However, the inductive biases inherent in neural networks can hinder spatial coherence in ill-posed scenarios. Current methods focus either on enhancing encoding strategies for deformation fields, often resulting in opaque and less intuitive models, or adopt explicit techniques like linear blend skinning, which rely on heuristic-based node initialization. Additionally, the potential of implicit representations for interpolating sparse temporal signals remains under-explored. To address these challenges, we propose a spline-based trajectory representation, where the number of knots explicitly determines the degrees of freedom. This approach enables efficient analytical derivation of velocities, preserving spatial coherence and accelerations, while mitigating temporal fluctuations.

Authors' Contact Information: Mingyang Song, DisneyResearch|Studios, Zürich, Switzerland and ETH Zürich, Zürich, Switzerland, mingyang.song@inf.ethz.ch; Yang Zhang, DisneyResearch|Studios, Zürich, Switzerland, yang.zhang@disneyresearch.com; Marko Mihajlovic, ETH Zürich, Zürich, Switzerland, marko.mihajlovic@inf.ethz.ch; Siyu Tang, ETH Zürich, Zürich, Switzerland, siyu.tang@inf.ethz.ch; Markus Gross, DisneyResearch|Studios, Zürich, Switzerland and ETH Zürich, Zürich, Switzerland, gross@disneyresearch.com; Tunç Ozan Aydın, DisneyResearch|Studios, Zürich, Switzerland, tunc@disneyresearch.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGGRAPH Conference Papers '25, Vancouver, BC, Canada

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1540-2/25/08

<https://doi.org/10.1145/3721238.3730676>

To model knot characteristics in both spatial and temporal domains, we introduce a novel low-rank time-variant spatial encoding, replacing conventional coupled spatiotemporal techniques. Our method demonstrates superior performance in temporal interpolation for fitting continuous fields with sparse inputs. Furthermore, it achieves competitive dynamic scene reconstruction quality compared to state-of-the-art methods while enhancing motion coherence without relying on linear blend skinning or as-rigid-as-possible constraints.

CCS Concepts: • **Computing methodologies** → **Reconstruction; Motion processing; Neural networks.**

Additional Key Words and Phrases: Cubic Hermite Spline, Temporal Consistency, Moran's I .

ACM Reference Format:

Mingyang Song, Yang Zhang, Marko Mihajlovic, Siyu Tang, Markus Gross, and Tunç Ozan Aydın. 2025. Spline Deformation Field. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25)*, August 10–14, 2025, Vancouver, BC, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3721238.3730676>

1 Introduction & Related Work

The demand for precise and computationally efficient trajectory modeling has significantly increased in fields such as character animation, motion capture, and 3D reconstruction. Traditionally, motion modeling was achieved through explicit constraints or hand-crafted priors [Eisenberger et al. 2018; Huang et al. 2008; Innmann et al. 2020; Kilian et al. 2007; Myronenko and Song 2010; Sorkine and

Alexa 2007], ensuring spatial coherence and adherence to physical laws. However, such priors are not always practical, particularly when dealing with sparse or noisy data.

To address these limitations, neural network-based deformation modeling approaches have emerged, offering task-agnostic, differentiable, and continuous representations. Among these, coordinate neural networks (Coord.-NN) have gained prominence for their ability to encode continuous signals directly as network weights using query coordinates as input. This is typically achieved through a compact implicit function comprising a 3+1D encoding followed by a Multi-Layer Perceptron (MLP). Coord.-NNs generally produce reasonably smooth results regardless of the encoding scheme (e.g., sinusoidal, dense grid, hash grid), benefiting from an inherent smoothness inductive bias. This bias can be attributed to the continuity of MLP, the smoothness of activation functions, the local coherence of multi-resolution grids interpolation, or their interplay. However, while Coord.-NNs are widely used for modeling various types of temporal changes, it remains unclear which source of bias offers the best trade-off between performance and adaptability across different scenarios.

Recent Dynamic Point Fields (DPF) [Prokudin et al. 2023] leverages the intrinsic regularization power of SIREN [Sitzmann et al. 2020] to map unregistered point pairs. DOMA [Zhang et al. 2024] extends this further to multi-frame sequences by incorporating time conditioning. While this approach achieved improved temporal consistency compared with per-frame modeling, the coupled 4D signals introduced potential wavy fluctuations. ResFields [Mihajlovic et al. 2024] mitigated such biases by disentangling spatial and temporal signals, enhancing the capacity of MLPs for spatiotemporal signal modeling. However, its reliance on temporal discretization limited its interpolation performance. Constant canonical points and direct supervision are common characteristics across the aforementioned tasks. As a result, SIREN [Sitzmann et al. 2020], with its higher non-linearity and representational capacity, is typically preferred over sinusoidal positional encoding in such more-constrained scenarios.

In scenarios where the points to be deformed are also optimization targets, such as dynamic scene reconstruction with Gaussian Splatting (GS) [Kerbl et al. 2023], Coord.-NN’s inductive bias preserves spatial consistency fairly for free. In such cases, the canonical positions of 3D Gaussians are optimized alongside the deformation field. 4DGS [Wu et al. 2023] represented the deformation field using Hexplane [Cao and Johnson 2023; Fridovich-Keil et al. 2023], achieving strong photometric results on multi-view datasets. However, it suffered from floaters due to the limited regularization capacity of the total variation loss. Grid4D [Xu et al. 2024] addressed overlapping issues inherent to Hexplanes and introduced a perturbation loss, effectively reducing artifacts in grid-based methods. Similar to D-NeRF [Pumarola et al. 2021], Deformable3DGS (Def.GS) [Yang et al. 2023] utilized MLP to predict offsets, whose global continuity yielded better performance on synthetic datasets, even without extensive regularization. SC-GS [Huang et al. 2023] tackled the limited rigidity of Def.GS by introducing the concept of sparse-controlled GS, which relies on sampling sparse control nodes from the scene trained in the first stage. Moreover, the ARAP constraints used in SC-GS [Huang et al. 2023], which originate from meshes

with fixed vertex connectivity, may not perform effectively in large-scale scenes where densification causes significant changes in local neighborhoods. ReMatching [Oblak et al. 2024] adopted deformation priors to enhance spatial coherence in a single stage. While effective, the heuristics introduced by multi-stage training and prior choice limit generalization ability.

Inspired by previous works [Chugunov et al. 2023; Knodt 2022], we propose representing trajectories using splines, a classic tool from applied mathematics, that naturally model continuous motion over time. With the timeline discretized by knots, we concept on the design of ResFields [Mihajlovic et al. 2024], which decouples spatial and temporal signals, and extend this approach to multiple existing encoding schemes. Our experiments demonstrate that explicit spline interpolation outperforms deep feature interpolation, particularly in scenarios with sparse temporal signals. More importantly, the analytical derivation of kinematic properties from the interpolation function enables effective regularization of spatial coherence and reduces temporal jitter without requiring two-stage training or reliance on priors, thus reducing heuristics in the training process. Additionally, we introduce a novel measurement for spatial coherence that aligns well with other metric evaluations and visual reconstruction results. Finally, we analyze the limitations of our representation, offering insights that may guide future research in this area.

2 Method

2.1 Spline Interpolation

A trajectory from $t = 0$ to $t = 1$ is uniformly divided into $N - 1$ time intervals, resulting in N knots. Unlike global fitting methods that utilize a single polynomial, Fourier series, or their combination [Lin et al. 2024], the localized nature of splines minimizes potential temporal oscillations, particularly in static regions. To prevent under-fitting, the choice of N must correspond to the number of input temporal steps, ensuring a well-determined system. Specifically, for T training timestamps, the number of parameters introduced by N knots must satisfy the following equation:

$$K \cdot N = T, \quad (1)$$

where K is a factor determined by the order of the polynomial. Here, we focus on the Cubic Hermite Spline [Wikipedia 2024], which represents a trajectory using consecutive third-order polynomial segments. For a given query timestamp t_{query} and N knots, we first locate the time interval to which it belongs and identify the corresponding starting and ending knots, denoted as t_{start} and t_{end} , respectively. The interval is then normalized to a relative time $\bar{t} \in [0, 1]$. The interpolation function for each segment is defined as:

$$p(\bar{t}) = (2\bar{t}^3 - 3\bar{t}^2 + 1)p_0 + (\bar{t}^3 - 2\bar{t}^2 + \bar{t})m_0 + (-2\bar{t}^3 + 3\bar{t}^2)p_1 + (\bar{t}^3 - \bar{t}^2)m_1, \quad (2)$$

where p_0 and p_1 represent the properties at the starting and ending knots, while m_0 and m_1 are the corresponding starting and ending tangents. We treat the tangents as independent optimizable parameters. We follow Eq. 1 and set $K = 2$, which corresponds to choosing $N = T/2$, thereby guaranteeing a theoretically well-determined fit.

To preserve spatial coherence, we employ a Canonical-Deformation design, where the parameters p and m of points are predicted by

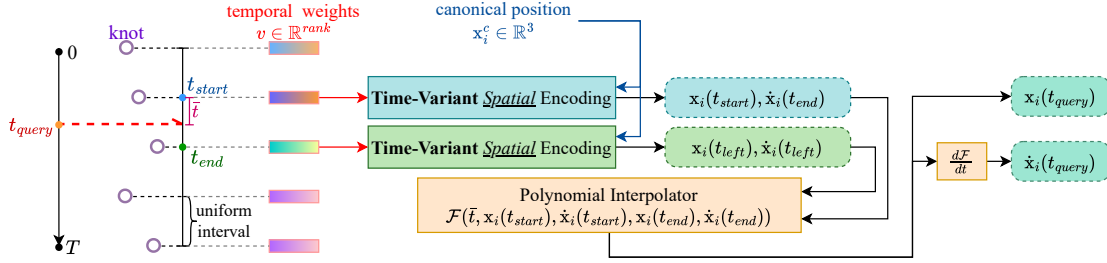


Fig. 2. The pipeline of our method. We explain how to query value at arbitrary timestamps through polynomial interpolation in Sec. 2.1, corresponding to the timelines in the left part and orange block. We describe our design of spatiotemporal conditioning in Sec. 2.2, which corresponds to the red (temporal signal) and blue (spatial signal) arrows.

a Coord.-NN using spatial coordinates in the canonical space as inputs. This process can be expressed as:

$$\begin{aligned} \mathbf{X}^c &= \{\mathbf{x}_i^c, \dot{\mathbf{x}}_i^c \in \mathbb{R}^3\}_{i=1, \dots, N_p}, \\ \mathbf{x}_i(t_{query}) &= \mathcal{F}(\bar{t}, \mathbf{x}_i(t_{start}), \dot{\mathbf{x}}_i(t_{start}), \mathbf{x}_i(t_{end}), \dot{\mathbf{x}}_i(t_{end})), \\ \Delta \mathbf{x}_i(t_{start}), \dot{\mathbf{x}}_i(t_{start}) &= \Phi_\theta(\mathbf{x}_i^c, t_{start}), \\ \mathbf{x}_i(t_{start}) &= \mathbf{x}_i^c + \Delta \mathbf{x}_i(t_{start}), \end{aligned} \quad (3)$$

where \mathbf{x}_i^c denotes the spatial coordinates of points in the canonical space, N_p is the total number of points, and $\mathcal{F}(\dots)$ represents the interpolation function described in Eq. 2. To enhance clarity, we substitute p_0 with $\mathbf{x}_i(t_{start})$ and m_0 with $\dot{\mathbf{x}}_i(t_{start})$. For brevity, we omit the derivation of t_{end} here and include it in the supp. material.

2.2 Time-variant Spatial Encoding (TVSE)

Previous works [Fridovich-Keil et al. 2023; Huang et al. 2023; Pumarola et al. 2021; Wu et al. 2023; Xu et al. 2024; Yang et al. 2023] treat four-dimensional inputs (three spatial coordinates and one timestamp) equally. Recently, ResFields [Mihajlovic et al. 2024] introduced a novel approach by injecting temporal information into the weights of pure MLP-based spatial encodings through time-variant residual and base weights, enhancing the performance of various previous methods. Its architecture-agnostic design naturally extends to other encoding schemes. The key conceptual difference can be expressed as:

$$\begin{aligned} \text{Previous work: } & \Phi_\theta(\gamma(x), \gamma(y), \gamma(z), \gamma(t)), \\ \text{Ours: } & \Phi_\theta(\gamma_{\varphi(t)}(x), \gamma_{\varphi(t)}(y), \gamma_{\varphi(t)}(z)), \end{aligned} \quad (4)$$

where $\gamma(\cdot)$ represents the encoding (e.g., Sinusoidal [Mildenhall et al. 2021] or grids [Cao and Johnson 2023; Müller et al. 2022]), $\gamma_{\varphi(t)}(\cdot)$ denotes time-variant encoding, and $\varphi(t)$ is the temporal signal injection function, which varies based on the specific design of $\gamma(\cdot)$.

Building on the concept of ResFields [Mihajlovic et al. 2024], we incorporate low-rank decomposition in the temporal domain to achieve both compactness and implicit regularization. This approach can be expressed as:

$$\phi(t) = \mathbf{b}_{base} + \sum_{r=1}^{rank} \mathbf{v}_t[r] \cdot \mathbf{B}_{res}[r], \quad (5)$$

where $[r]$ denotes element indexing, $\mathbf{v}_t \in \mathbb{R}^{rank}$ represents trainable weights associated with each timestep, $\mathbf{B}_{res} \in \mathbb{R}^{rank \times \dots}$ are

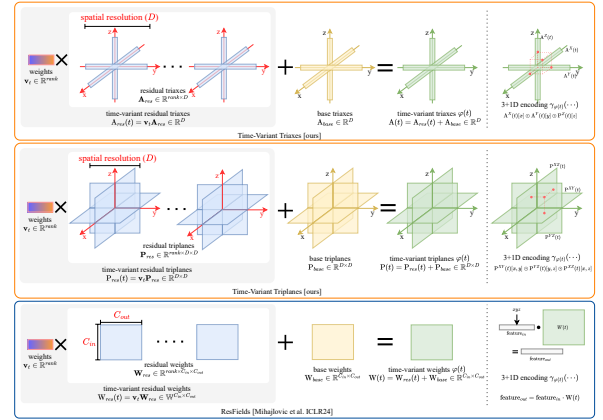


Fig. 3. An intuitive diagram of our TVSE. Rounded boxes with different edge colors mark grid-based and MLP and methods. $A'(\cdot)[\cdot]$ and $P''(\cdot)[\cdot, \cdot]$ denote sampling values at specific position from an axis and plane, respectively. Following [Fridovich-Keil et al. 2023], we aggregate features sampled from the three axes and planes using element-wise multiplication \odot .

the residual bases, and \mathbf{b}_{base} corresponds to the time-invariant encoding. These notations carry different meanings depending on the chosen encoding scheme. In this work, we primarily focus on cosine functions and multi-resolution dense grids. Fig. 3 provides an intuitive visualization of the design of $\gamma_{\varphi(t)}$. We refer to the supplementary material for detailed implementations.

2.3 Velocity and Acceleration Regularization

By taking derivative w.r.t time in Eq. 2, the velocity of a point on the spline can be determined by the following function:

$$\begin{aligned} v(\bar{t}) &= (6\bar{t}^2 - 6\bar{t})p_0 + (3\bar{t}^2 - 4\bar{t} + 1)m_0 + \\ &\quad (-6\bar{t}^2 + 6\bar{t})p_1 + (3\bar{t}^2 - 2\bar{t})m_1, \end{aligned} \quad (6)$$

where $v(\bar{t})$ has physical meaning, i.e., velocity. With the above closed-form velocity function, the velocity loss can be formulated as:

$$\mathcal{L}_v^i = \sum_{j \in \mathcal{N}_k(i)} w_{ij} \|v_i - v_j\|_2^2, \quad (7)$$

where i is the index of points, $\mathcal{N}_k(i)$ represents the k nearest neighbors of i , and j is the local index in neighborhood, w_{ij} is weight calculated from relative distance.

We propose one additional constraint to reduce high-frequency jitter and alleviate low-frequency oscillation. Specifically, taking derivation w.r.t time in Eq. 6, we can get the analytical acceleration:

$$a(\bar{t}) = (12\bar{t} - 6)p_0 + (6\bar{t} - 4)m_0 + (-12\bar{t} + 6)p_1 + (6\bar{t} - 2)m_1. \quad (8)$$

Then, we regularize the acceleration of points through the following loss function:

$$\mathcal{L}_{acc}^i = |a_i|. \quad (9)$$

Finally, the loss function is:

$$\mathcal{L} = \mathcal{L}_{recon} + \alpha \mathcal{L}_v + \beta \mathcal{L}_{acc}, \quad (10)$$

where \mathcal{L}_{recon} represents supervisions from ground truth to prediction, α and β are hyperparameters for regularization.

3 Results

3.1 Learning Continuous Deformation Field

As a proof of concept, we start with a more constrained problem. Given sparsely sampled point trajectories, the goal is to fit an implicit continuous deformation field that can infer the trajectories of unseen points, ensuring alignment with observed data. We focus on a more challenging yet practical scenario where available temporal signals are scarce to demonstrate the strength of explicit spline interpolation. Specifically, we choose nine long sequences (each exceeding 500 frames) from DeformingThings4D (humanoids) [Li et al. 2021] and train models with every 4th/6th timestep, leaving 75%/83.3% timesteps of timesteps for evaluation. Following DOMA [Zhang et al. 2024], 25% vertices from the starting mesh are sampled to calculate L1 distance for supervision during training. The determination of N follows Eq. 1. In addition to end point error (EPE), we propose a new metric for evaluating spatial coherence using Moran’s I [Mihajlovic et al. 2025; Moran 1950]. We provide its detailed implementation in the supplementary material.

We benchmark our method against state-of-the-art approaches in scene flow estimation, as detailed in Tab. 1. DOMA [Zhang et al. 2024] employs periodical activation layers, and the coupled spatial and temporal signals introduce significant temporal fluctuation artifacts, which also impairs spatial performance. This issue cannot be resolved simply by adjusting the model’s number of parameters. ResFields [Mihajlovic et al. 2024] disentangles time and position coordinates and results in clear improvements, motivating our design in Sec. 2.2. Its performance can be further refined by adjusting temporal capacity and rank or adapting as-isometric-as-possible (AIAP) regularization [Prokudin et al. 2023; Qian et al. 2024] (based on setting †). For example, $c=50\%$ indicates that the number of temporal weights, $v \in \mathbb{R}^{rank}$, is reduced to half the number of training timesteps, thereby decreasing the temporal DoF. However, motions generated by interpolating deep features may still exhibit distortions, as shown in Fig. 4. In contrast, our method relies on explicit spline interpolation and achieves better performance with $c=100\%$, reducing the risk of underfitting. Based on setting $c=100\%$, $r=60$ (marked as ‡), we conduct ablation studies on \mathcal{L}_v , \mathcal{L}_{acc} and AIAP, and the relative performance shows their effectiveness.

Table 1. Interpolation of scene flow on DeformingThings4D’s long sequences. EPE results are multiplied by $\times 10^4$ for better readability. The highlighted rows denote 1st, 2nd, and 3rd best models. Note that when $c=90\%$, our methods still produce plausible interpolations (shown in Fig. 10).

Method	Settings	x4		x6	
		EPE↓	M.’s I ↑	EPE↓	M.’s I ↑
DOMA [Zhang et al. 2024]	#p=0.03M	138.80	0.912	146.51	0.911
	#p=0.1M	92.59	0.876	105.28	0.877
	#p=1.6M	473.17	0.249	452.23	0.254
	#p=2.1M	927.26	0.115	906.67	0.118
ResFields [Mihajlovic et al. 2024]	r=30	52.09	0.899	80.74	0.905
	c=100% r=60	54.23	0.891	80.63	0.898
	r=90	55.01	0.882	84.19	0.885
	c=90%	52.19	0.894	81.60	0.898
	c=75% r=60	50.75	0.896	80.56	0.898
	c=50% †	45.79	0.896	74.35	0.894
	c=25%	48.60	0.894	74.92	0.889
	† w/ AIAP	45.72	0.927	73.00	0.923
SF-Siren-ResFields	r=30	42.98	0.917	68.73	0.927
	c=100% r=60‡	40.74	0.919	68.28	0.926
	r=90	41.67	0.917	69.38	0.926
	‡ w/o \mathcal{L}_v	42.15	0.916	68.40	0.924
	‡ w/o \mathcal{L}_{acc}	42.66	0.917	70.12	0.929
	‡ w/ AIAP	41.64	0.938	68.93	0.943
	c=90% r=60	45.97	0.921	80.22	0.927

This explicitness also ensures intuitive feedback when adjusting hyperparameters: reducing N in such large-motion dataset can lead to large misalignment that significantly penalizes reconstruction accuracy. The following section shows a practical case where our explicit constraints on N enhance reconstruction performance by leveraging motion priors.

3.2 Dynamic Scene Reconstruction with GS

We evaluate our method against canonical-deformation-based scene reconstruction approaches: Def.3D [Yang et al. 2023], 4DGS [Wu et al. 2023], SC-GS [Huang et al. 2023], and Grid4D [Xu et al. 2024]. Since positional encoding performs better in such an ill-constrained problem, we use it instead of Siren in this section, which we name SF-PE-ResFields. We also implement Eq. 5 for grid-based methods, which we name SF-Triaxes and SF-Triplanes, and refer to supplementary material for details.

3.2.1 Real-world datasets. NeRF-DS [Yan et al. 2023] features slow-moving objects with specular surfaces. In this dataset, only a single view is available per timestamp, and the camera positions are less accurate than those in synthetic datasets. To account for the motion characteristics in this dataset, we set $N = 30$ and $rank = 10$, showcasing the utility of explicit (N) regularization. As shown in Tab. 2, both grid-based and MLP-based methods benefit from our representation, achieving superior fidelity and perceptual quality than the baselines. Our method delivers competitive reconstruction quality compared to two-stage LBS-based SC-GS [Huang et al. 2023], with notably better temporal stability in static, textureless regions (white table, curtain in background). Additionally, we evaluate our method




Fig. 4. Qualitative comparison of $\times 6$ (left) and $\times 4$ (right) scene flow interpolation. Frames boxed in red are training timesteps, and the following ones are interpolated. Our method achieves fair interpolated motions without skinning. We provide more visual results in the extra pages.

Table 2. Quantitative results on NeRF-DS [Yan et al. 2023] dataset. Notably, traditional metrics (PSNR and SSIM) penalize sharper but misaligned results over blurry results. We introduce perturbation in Grid4D [Xu et al. 2024] to 4DGS [Wu et al. 2023] since it consistently improves grid-based methods, which we denote with *. SC-GS’s [Huang et al. 2023] metrics except for M’s I are adopted from the official script.

Method	As				Basin				Bell				Cup			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow
Def.3D [Yang et al. 2023]	26.01	87.61	12.33	0.685	19.51	77.97	13.43	0.735	25.23	83.94	12.07	0.676	24.53	88.11	11.47	0.719
SC-GS [Huang et al. 2023]	26.00	-	11.40	0.924	19.60	-	15.40	0.919	25.10	-	11.70	0.946	24.50	-	11.50	0.955
4DGS* [Wu et al. 2023]	25.60	86.27	13.36	0.780	18.90	76.97	15.43	0.775	24.60	82.66	14.05	0.758	24.42	88.23	11.55	0.787
4DGS [Wu et al. 2023]	25.69	86.43	13.60	0.845	18.88	75.92	17.59	0.723	24.46	82.30	13.87	0.732	24.44	88.04	12.02	0.739
SF-Triaxes	26.89	88.46	13.40	0.865	19.45	79.29	16.04	0.936	26.05	86.06	12.58	0.920	24.24	88.53	11.90	0.925
SF-Triplanes	26.34	88.75	12.43	0.891	19.61	79.85	14.67	0.937	25.68	85.27	13.01	0.928	24.40	88.60	11.33	0.932
SF-PE-ResFields	26.82	88.90	12.01	0.932	19.70	79.18	13.67	0.959	25.55	84.50	12.41	0.972	24.85	88.92	11.21	0.962

Method	Plate				Press				Sieve				Average			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	M’s I \uparrow
Def.3D [Yang et al. 2023]	20.33	80.26	19.12	0.702	25.40	85.61	13.66	0.720	25.24	86.72	10.85	0.703	23.75	84.32	13.28	0.706
SC-GS [Huang et al. 2023]	20.20	-	20.20	0.938	26.60	-	13.50	0.945	26.00	-	11.40	0.927	24.10	-	14.00	0.939
4DGS* [Wu et al. 2023]	20.02	79.81	20.19	0.869	25.61	85.03	14.41	0.744	25.76	87.47	11.51	0.801	23.56	83.78	14.36	0.788
4DGS [Wu et al. 2023]	19.32	78.37	21.29	0.846	25.04	84.35	15.82	0.682	24.94	87.01	11.83	0.765	23.25	83.20	15.15	0.762
SF-Triaxes	20.92	81.74	18.63	0.986	26.22	88.21	14.09	0.908	26.71	87.68	11.47	0.943	24.36	85.71	14.01	0.926
SF-Triplanes	21.17	82.57	18.14	0.982	26.49	88.29	13.10	0.917	26.57	87.78	11.24	0.927	24.32	85.87	13.42	0.931
SF-PE-ResFields	21.02	82.51	17.41	0.993	27.09	88.33	12.29	0.967	25.80	88.32	10.96	0.965	24.40	85.81	12.85	0.964

on another single-view-per-timestep dataset, Hyper-NeRF [Park et al. 2021]. It contains higher-resolution images but suffers from less accurate camera positions. We follow Eq. 1 to determine N and set $rank = 16$. As shown in Tab. 3, our method strikes a balance between fidelity and perceptual quality, with fewer abnormal floaters compared to Grid4D [Xu et al. 2024]. As pointed out in Hyper-NeRF’s original script, LPIPS (*Alex*) aligns better with visual results, as illustrated in Fig. 6. The imprecise camera poses in these two datasets lead to significant visual variations in sharpness across different views, a challenge most evident in grid-based methods without additional regularization. With strong explicit constraints, our variant SF-Triplanes reasonably preserve rigidity, as shown in the torch in Fig. 6. As expected, MLP-based methods maintain better spatial coherence in scenes with simple motions due to their inherent global continuity but frequently fail in extreme cases.

Neu3D [Li et al. 2022] is a multi-camera-per-timestamp dataset with significantly more accurate camera positions and features more

complex motions than the previously mentioned datasets. As shown in Tab. 4, the well-defined camera setup and large static regions yield comparable reconstruction metrics across all methods, while our approach achieves notable improvements in spatial coherence, enhancing the reconstruction of challenging motions, as illustrated in Fig. 7. The accurate multi-view camera setup reduces the performance gap of photometric metrics between MLP and grid-based methods. Notably, as evidenced by Moran’s I comparison between 4DGS [Wu et al. 2023] and Def.3D [Yang et al. 2023], grid-based methods demonstrate superior spatial coherence when supplemented with total variation [Fridovich-Keil et al. 2023; Wu et al. 2023] loss. This issue likely stems from the limited capacity and temporal fluctuation bias of a single MLP with coupled 4D inputs to model long-term signals, aligning with the experimental observations in Sec. 3.1. However, our approach effectively mitigates this limitation by leveraging decoupling, low-rank factorization, and velocity regularization.

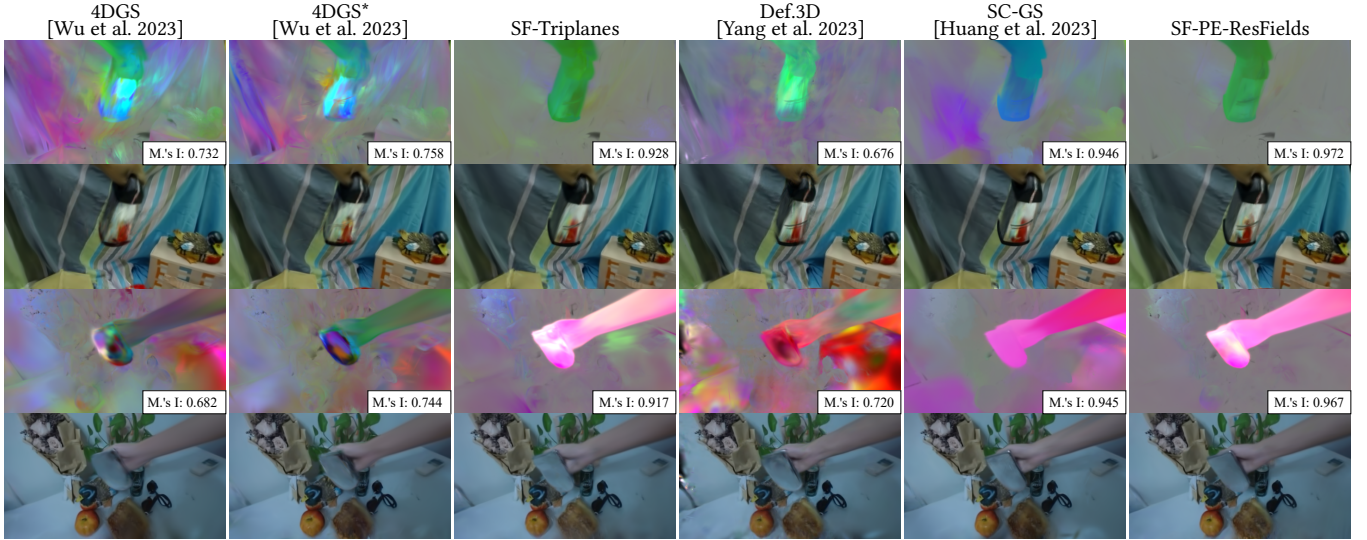


Fig. 5. Qualitative comparison of scene flows on NeRF-DS [Yan et al. 2023] dataset. With the derived regularizations in Sec. 2.3, our representation effectively alleviates the high frequency caused by imprecise camera positions. We normalize all frames’ motion vectors or velocities, so the gray color indicates static regions. We encourage readers to focus on the smoothness of motions.

Table 3. Quantitative results on Hyper-NeRF [Park et al. 2021] dataset. Scenes in which most canonical-deformation-field-based methods fail are excluded. Results with abnormal floaters or diverged ones are excluded from the comparison (denoted as ‘-’). Jittery caused by inaccurate cameras is alleviated by applying \mathcal{L}_{acc} in our method, resulting in higher fidelity performance (PSNR) albeit reduced sharpness in some scenes.

Method	Chicken (vrig)			3D Printer (vrig)			Peal Banana (vrig)			Chickchicken (interp)		
	PSNR↑	LPIPS↓	M’s $I \uparrow$	PSNR↑	LPIPS↓	M’s $I \uparrow$	PSNR↑	LPIPS↓	M’s $I \uparrow$	PSNR↑	LPIPS↓	M’s $I \uparrow$
4DGS [Wu et al. 2023]	28.381	17.147	0.758	22.032	23.669	0.648	28.286	15.860	0.804	26.680	20.228	0.703
Grid4D [Xu et al. 2024]	29.197	18.834	0.890	22.310	20.525	0.918	28.512	14.841	0.898	27.428	20.961	0.901
Def.3D [Yang et al. 2023]	21.816	18.742	0.943	18.814	27.319	0.897	-	-	-	25.040	19.507	0.890
SF-Triplanes w/o acc	24.393	18.604	0.898	21.786	19.802	0.902	28.334	14.363	0.947	26.037	17.878	0.954
SF-Triplanes w/ acc	28.829	16.889	0.877	22.169	22.448	0.866	28.083	17.031	0.894	26.154	19.280	0.945
SF-PE-ResFields w/o acc	23.204	18.465	0.960	20.219	18.696	0.975	26.164	17.289	0.975	26.278	15.788	0.985
SF-PE-ResFields w/ acc	28.751	16.305	0.949	21.934	24.114	0.946	27.579	17.175	0.960	26.543	15.788	0.983
Method	Cut Lemon (interp)			Hand (interp)			Slice Banana (interp)			Torch Chocolate (interp)		
	PSNR↑	LPIPS↓	M’s $I \uparrow$	PSNR↑	LPIPS↓	M’s $I \uparrow$	PSNR↑	LPIPS↓	M’s $I \uparrow$	PSNR↑	LPIPS↓	M’s $I \uparrow$
4DGS [Wu et al. 2023]	31.427	12.182	0.696	28.975	20.379	0.471	25.567	26.532	0.715	27.437	16.077	0.754
Grid4D [Xu et al. 2024]	32.101	11.862	0.746	31.338	12.656	0.694	25.869	25.887	0.801	28.256	10.842	0.779
Def.3D [Yang et al. 2023]	31.047	12.480	0.866	28.824	19.685	0.838	-	-	-	27.329	13.120	0.921
SF-Triplanes w/o acc	30.674	12.984	0.837	30.135	16.175	0.904	25.313	24.683	0.943	28.077	12.917	0.971
SF-Triplanes w/ acc	30.775	13.855	0.824	30.137	15.490	0.898	25.180	25.867	0.937	28.129	13.176	0.965
SF-PE-ResFields w/o acc	31.457	12.148	0.938	30.432	14.705	0.967	25.505	22.281	0.975	28.185	12.390	0.987
SF-PE-ResFields w/ acc	31.509	11.770	0.942	30.356	14.810	0.964	25.874	25.191	0.968	28.310	12.421	0.984

3.2.2 Synthetic dataset. We further evaluate our method on the synthetic, one-camera-per-timestep dataset, D-NeRF [Pumarola et al. 2021]. The quantitative results are included in the supplementary material. Our method slightly lags behind state-of-the-art methods regarding fidelity. This performance gap can be attributed to minor temporal misalignments, which we believe are due to biases in our representation. A detailed discussion of this issue is provided in the supplementary material, and we leave its resolution as future

work. Despite the slightly reduced fidelity performance, our representation consistently generates spatially coherent motions, and the misalignment is visually negligible, as shown in Fig. 12.

4 Discussions & Limitations

To further reason on our choice of $\gamma_{\varphi(t)}(\cdot)$ over $\gamma(\cdot)$, we illustrate the temporal weights, as shown in Fig. 11. The weights, aligned with the actual motions, demonstrated obvious symmetricity and periodicity. We expect this evidence to provide insights for the following works. We briefly discuss the potential applications based on our




Fig. 6. Qualitative comparison of rendered images and scene flows on Hyper-NeRF [Park et al. 2021] dataset. Since the whole scene is jittery due to inaccurate cameras, static regions are also colored, compared with the NeRF-DS[Yan et al. 2023] dataset. Additional results are included in Fig. 9.

Table 4. Quantitative results on Neu3D [Li et al. 2022]. 4DGS and our SF-Triplanes perform better than MLP-based methods in the long term with complex motion scenes. Grid4D [Xu et al. 2024] deactivated or reduced weights for smooth regularization on this dataset in their official implementation.

Method	Coffee Martini				Cook Spinach				Cut Roasted Beef			
	PSNR↑	SSIM↑	LPIPS↓	M.'s I↑	PSNR↑	SSIM↑	LPIPS↓	M.'s I↑	PSNR↑	SSIM↑	LPIPS↓	M.'s I↑
4DGS [Wu et al. 2023]	28.59	91.45	10.96	0.938	32.61	94.86	8.17	0.888	32.96	95.00	8.16	0.855
Def.3D [Yang et al. 2023]	28.81	91.64	11.07	0.915	32.06	94.83	8.46	0.842	32.84	94.97	8.43	0.809
Grid4D [Xu et al. 2024]	28.30	90.02	13.59	0.418	32.64	94.96	8.62	0.700	33.22	94.47	8.75	0.669
SF-Triplanes	28.88	91.67	10.86	0.965	32.02	94.79	8.15	0.931	33.11	95.06	8.02	0.908
SF-PE-Resfields	28.74	91.48	10.86	0.967	32.47	94.89	8.54	0.939	32.57	95.04	7.99	0.913

Method	Flame Salmon				Flame Steak				Sear Steak				Average			
	PSNR↑	SSIM↑	LPIPS↓	M.'s I↑	PSNR↑	SSIM↑	LPIPS↓	M.'s I↑	PSNR↑	SSIM↑	LPIPS↓	M.'s I↑	PSNR↑	SSIM↑	LPIPS↓	M.'s I↑
4DGS [Wu et al. 2023]	29.16	91.98	10.22	0.930	32.82	95.52	7.04	0.869	33.39	95.59	7.31	0.909	31.59	94.04	8.64	0.898
Def.3D [Yang et al. 2023]	28.28	91.79	10.56	0.924	32.76	95.49	7.23	0.825	33.00	95.49	7.67	0.760	31.29	94.04	8.90	0.846
Grid4D [Xu et al. 2024]	29.12	90.90	12.91	0.452	32.56	95.71	7.31	0.676	33.16	95.68	7.47	0.650	31.50	93.62	9.78	0.594
SF-Triplanes	29.21	91.90	10.13	0.961	33.21	95.65	7.05	0.914	33.15	95.43	7.31	0.901	31.60	94.08	8.59	0.930
SF-PE-Resfields	28.97	91.99	10.00	0.960	32.82	95.65	6.95	0.934	33.23	95.61	7.45	0.908	31.47	94.11	8.63	0.937




Fig. 7. Qualitative comparison of rendered images and scene flows on Neu3D [Li et al. 2022] dataset. Our velocity regularizations preserve the handle with reflection better than the baselines. The differences are more evident in rendered motion vectors when photometric reconstructions are mixed.

representation. Firstly, we qualitatively demonstrate the possibility of advecting points to verify the reasonableness of the derived velocity in Eq. 6. Specifically, we calculate the velocity of the last timestamp and propagate points with such constant speed. Secondly, we show the possibility of editing in the trajectory manifold. Furthermore, the established coherency enables us to propagate appearance modification in one frame to the whole sequence. The results are qualitatively shown in Fig. 8 and 13.

Explicit spline interpolation demonstrates superior performance in large-scale interpolation and smooth motion modeling, which

aligns with our intuition and is confirmed by our experimental findings. Since reconstruction errors are minimized through gradient descent and cannot be reduced to precisely zero, the misalignment issue caused by deviations from the ground truth is exacerbated by explicit spline interpolation. More specifically, we set the number of knots to half the number of training signals when no reliable prior on motion is available. This ensures a clean pipeline and a well-determined fit, avoiding the need to heuristically choose the number of knots or optimize interval lengths. To address the above issue, we apply a regularization term akin to penalized splines. This

inevitably introduces slight offsets from the ground-truth trajectories, though this discrepancy is imperceptible in the absence of residual maps. Stemming from it, whether other curves (e.g., hybrid circular and elliptical [Yuksel 2020]) can solve such unexpected bias is a worthy topic.

5 Conclusion

In this work, we demonstrate that splines, though a classical tool from applied mathematics, remain effective in modern settings, achieving strong results across several datasets with our implementation. In early experiments, we attempted to model deformations using independent splines without the support of Coord.-NNs, but achieved only modest results, even with strong regularization. Interestingly, our findings reveal that splines and Coord.-NNs complement each other well. Coord.-NNs provide coherent guidance for splines through their smoothness inductive bias, while the explicit regularization derived from spline interpolation reinforces this smoothness, particularly in challenging or ill-constrained scenarios. Moreover, our improved performances also come from the decoupling of spatial and temporal signals, which is a consideration often overlooked. Finally, we provide a principled and robust evaluation on spatial coherency concept on Moran's I , which offers more informative insights in cases where photometric scores are ambiguous or inconsistent.

Acknowledgments

We gratefully acknowledge the creators of the open-source datasets used in this work. The Deer, Astra SambaDancing, and Pumpknhulk SwingDancing scenes shown in Fig. 1, 4, and 10 were created by Yang Li et al. [Li et al. 2021]. The Flame Steak and Cook Spinach scenes in Fig. 1 and 7 were created by Zhaoyang Lv et al. [Li et al. 2022]. The Bell and Press scenes in Fig. 5 were created by Zhiwen Yan et al. [Yan et al. 2023]. The Torch Chocolate, Slice Banana, 3D Printer, and ChickChicken scenes in Fig. 6 and 9 were created by Keunhong Park et al. [Park et al. 2021]. The Standup, Mutant, T-Rex, and Lego scenes in Fig. 8, 11, 12, and 13 were created by Albert Pumarola et al. [Pumarola et al. 2021].

References

- Ang Cao and Justin Johnson. 2023. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 130–141.
- Ilya Chugunov, David Shustin, Ruyu Yan, Chenyang Lei, and Felix Heide. 2023. Neural Spline Fields for Burst Image Fusion and Layer Separation. *arXiv preprint arXiv:2312.14235* (2023).
- Marvin Eisenberger, Zorah Löhner, and Daniel Cremers. 2018. Divergence-free shape interpolation and correspondence. *arXiv preprint arXiv:1806.10417* (2018).
- Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. 2023. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12479–12488.
- Qi-Xing Huang, Bart Adams, Martin Wicke, and Leonidas J Guibas. 2008. Non-rigid registration under isometric deformations. In *Computer Graphics Forum*, Vol. 27. Wiley Online Library, 1449–1457.
- Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. 2023. SC-GS: Sparse-Controlled Gaussian Splatting for Editable Dynamic Scenes. *arXiv preprint arXiv:2312.14937* (2023).
- Matthias Innmann, Kihwan Kim, Jinwei Gu, Matthias Nießner, Charles Loop, Marc Stamminger, and Jan Kautz. 2020. Nrmvs: Non-rigid multi-view stereo. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 2754–2763.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* 42, 4 (2023), 1–14.
- Martin Kilian, Niloy J Mitra, and Helmut Pottmann. 2007. Geometric modeling in shape space. In *ACM SIGGRAPH 2007 papers*. 64–68.
- Julian Knodt. 2022. Continuous Dynamic-NeRF: Spline-NeRF. *arXiv preprint arXiv:2203.13800* (2022).
- Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. 2022. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5521–5531.
- Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 2021. 4dcomplete: Non-rigid motion estimation beyond the observable surface. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12706–12716.
- Youtian Lin, ZuoZhuo Dai, Siyu Zhu, and Yao Yao. 2024. Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21136–21145.
- Marko Mihajlovic, Sergey Prokudin, Marc Pollefeys, and Siyu Tang. 2024. ResFields: Residual Neural Fields for Spatiotemporal Signals. In *International Conference on Learning Representations (ICLR)*.
- Marko Mihajlovic, Sergey Prokudin, Siyu Tang, Robert Maier, Federica Bogo, Tony Tung, and Edmond Boyer. 2025. Splatfields: Neural gaussian splats for sparse 3d and 4d reconstruction. In *European Conference on Computer Vision*. Springer, 313–332.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- Patrick AP Moran. 1950. Notes on continuous stochastic phenomena. *Biometrika* 37, 1/2 (1950), 17–23.
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)* 41, 4 (2022), 1–15.
- Andriy Myronenko and Xubo Song. 2010. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence* 32, 12 (2010), 2262–2275.
- Sara Oblak, Despoina Paschalidou, Sanja Fidler, and Matan Atzmon. 2024. ReMatching Dynamic Reconstruction Flow. *arXiv preprint arXiv:2411.00705* (2024).
- Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. 2021. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228* (2021).
- Sergey Prokudin, Qianli Ma, Maxime Raafat, Julien Valentin, and Siyu Tang. 2023. Dynamic Point Fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 7964–7976.
- Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. 2021. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10318–10327.
- Zhiyin Qian, Shaofei Wang, Marko Mihajlovic, Andreas Geiger, and Siyu Tang. 2024. 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5020–5030.
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. 2020. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems* 33 (2020), 7462–7473.
- Olga Sorkine and Marc Alexa. 2007. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, Vol. 4. Citeseer, 109–116.
- Wikipedia. 2024. Cubic Hermite spline — Wikipedia, The Free Encyclopedia. <http://en.wikipedia.org/w/index.php?title=Cubic%20Hermite%20spline&oldid=1219643794>. [Online; accessed 05-December-2024].
- Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 2023. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528* (2023).
- Jiawei Xu, Zexin Fan, Jian Yang, and Jin Xie. 2024. Grid4D: 4D Decomposed Hash Encoding for High-fidelity Dynamic Gaussian Splatting. *arXiv preprint arXiv:2410.20815* (2024).
- Zhiwen Yan, Chen Li, and Gim Hee Lee. 2023. Nerf-ds: Neural radiance fields for dynamic specular objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8285–8295.
- Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. 2023. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101* (2023).
- Cem Yuksel. 2020. A class of C 2 interpolating splines. *ACM Transactions on Graphics (TOG)* 39, 5 (2020), 1–14.
- Yan Zhang, Sergey Prokudin, Marko Mihajlovic, Qianli Ma, and Siyu Tang. 2024. Degrees of Freedom Matter: Inferring Dynamics from Point Trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2018–2028.



Fig. 8. Examples of motion advection and editing. **Left:** We propagate points in the last frame with the derived velocity. The plausible visual results showcase the effectiveness of \mathcal{L}_v . **Right:** We adjust motions by only editing keyframes. In this work, we can only neatly achieve trifling modifications since the definition of ‘motion editing’ is still ambiguous, and LBS-based deformation is tedious for users. We leave a clearer description of such a task and a streamlined editing pipeline as future work (e.g., cage-based deformation).

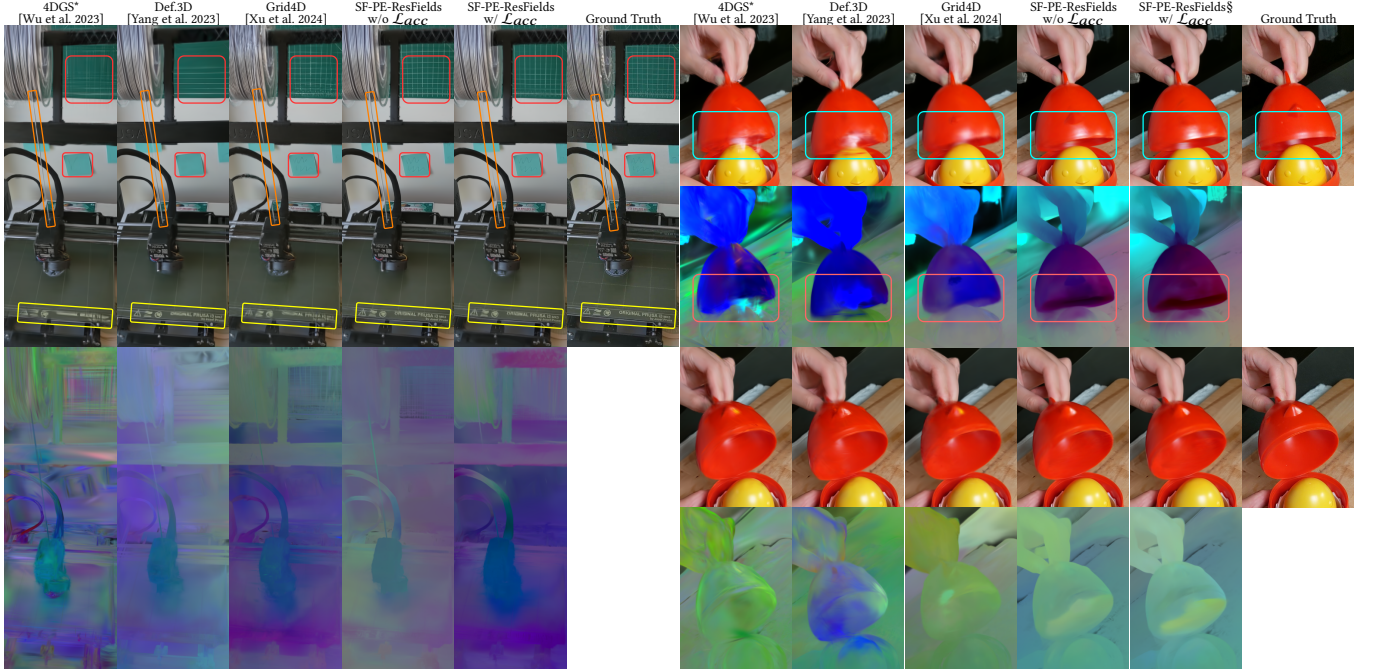


Fig. 9. Additional qualitative comparison of rendered images and scene flows on Hyper-NeRF [Park et al. 2021] dataset. Bounding boxes highlight major differences. We provide artifacts and failure cases of our method on this dataset in the supplementary material.




Fig. 11. Visualization of $\mathbf{v} \in \mathbb{R}^{rank}$ in Eq. 4. We concatenate the temporal weights of SF-Triplanes along the temporal dimension, i.e., each **column** in the heat map is a vector representing a deformed state. Not surprisingly, scene *lego* can be reconstructed with $rank = 1$.




Fig. 10. We present additional scene flow interpolation results, a task with significant practical value yet remains underexplored. Our model demonstrates strong performance and can serve as a baseline for future research.

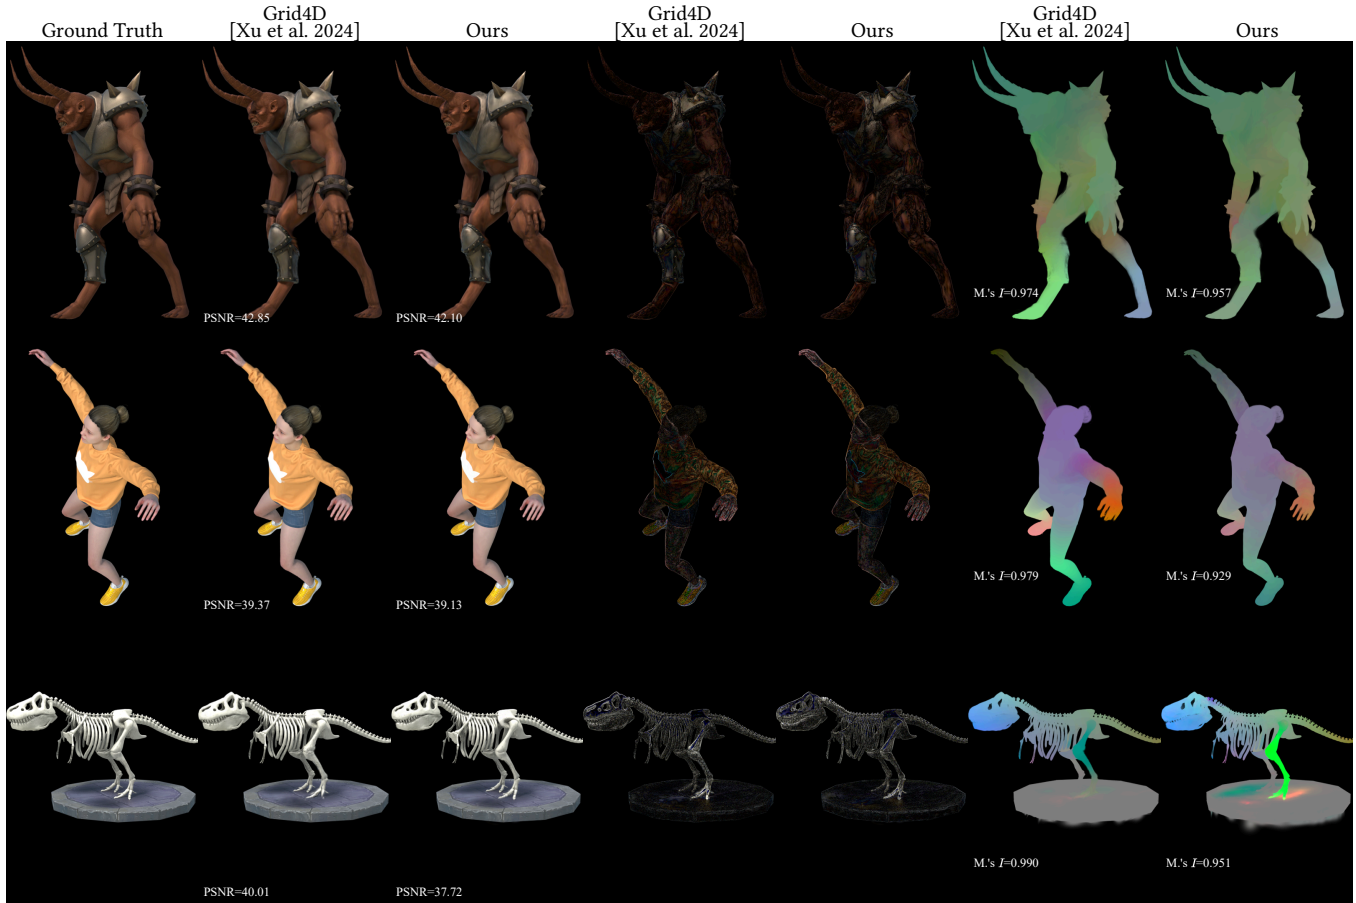


Fig. 12. We present the comparison between our method and Grid4D [Xu et al. 2024], the SOTA method on synthetic D-NeRF [Pumarola et al. 2021] dataset. While our results lag slightly behind on this dataset, our method is capable of reasonably reconstructing challenging regions. Our evaluations of spatial coherency by using Moran’s I align with these observations, verifying informativeness of this metric. As highlighted in the main paper, our explicit regularization demonstrates superior performance in real-world scenarios. We provide artifacts of our method on this dataset in the supplementary material.



Fig. 13. The canonical space of our variant, SF-PE-ResFields, achieves improved consistency with dynamic observations. As a result, we can perform style transfer on the avatar’s texture in a straightforward manner. Specifically, by transferring the style in the static canonical space, the new texture can be seamlessly propagated throughout the entire sequence. While new points are densified without constraint during style transfer, the continuity of the deformation field and regularization on the rendered depth map allow us to achieve these appealing results without the need for additional registration (e.g., Chamfer distance).