

Neu-PiG: Neural Preconditioned Grids for Fast Dynamic Surface Reconstruction on Long Sequences

Supplementary Material

In this supplementary document, we provide additional experiments, analyses, and visualizations that complement the results presented in the main paper. Specifically, we extend the evaluation of Neu-PiG with further quantitative and qualitative results to validate the impact of grid design choices, preconditioning strength, and input conditions.

In Figs. 9 and 10, we present extended visual results on a broader range of datasets, demonstrating that Neu-PiG maintains high reconstruction quality and temporal consistency across diverse motion types, outperforming existing state-of-the-art approaches in both accuracy and stability.

5.1. Latent Grid Design

We analyze the design choices of our preconditioned multi-resolution grid, focusing on the impact of the number of levels and the smoothness introduced by preconditioning.

Grid Levels. In Tab. 6, we evaluate the influence of the number of grid levels, described in Sec. 3.2, on reconstruction quality using the AMA dataset. We keep our default configuration and vary only the total number of levels L . The results show that even with a small number of levels, and thus reduced latent capacity, Neu-PiG achieves surprisingly strong reconstructions. However, excessively increasing the grid resolution leads to performance degradation due to overfitting, as shown in Fig. 6.

Complementing this quantitative analysis, Fig. 7 provides a qualitative view of the learned multi-resolution decomposition. Coarser levels capture low-frequency global motion, whereas finer levels recover localized high-frequency deformations, visualized using a local rigidity proxy based on Kabsch alignment.

Smoothness Weights. We investigate the effect of the smoothness parameter λ in our grid preconditioning, introduced in Sec. 3.5, on reconstruction performance. To access its influence, we vary the base smoothness value λ^1 to both smaller and larger magnitudes, while using the same increase per level as in the default configuration. The results in Tab. 7 on the AMA dataset indicate that Neu-PiG is largely robust to moderate changes in smoothness, while extreme values, either excessively high or entirely absent, lead to a clear degradation in reconstruction quality. We further visualize the learned latent spaces in Fig. 8 by performing a PCA analysis and plotting the first principal component.

Table 6. Ablation study on the number of grid levels in \mathcal{G}_p , evaluated on the AMA dataset. Increasing the number of levels improves reconstruction quality up to $L = 8$, after which performance slightly degrades due to overfitting.

	CD [$\times 10^{-5}$] ↓	NC ↑	F-0.5% ↑	Corr. ↓
$L = 1$	1.96	0.850	0.810	0.063
$L = 2$	1.76	0.859	0.841	0.054
$L = 4$	0.54	0.945	0.980	0.023
$L = 6$	0.46	0.949	0.987	0.018
Ours ($L = 8$)	0.44	0.951	0.989	0.017
$L = 10$	0.45	0.951	0.989	0.021
$L = 12$	0.46	0.948	0.988	0.024




Figure 6. Qualitative analysis of the effect of the number of grid levels during optimization on reconstruction quality. Neu-PiG achieves accurate and temporally stable reconstructions even with few grid levels, while excessively increasing resolution leads to mild overfitting and loss of fine structural detail.




Figure 7. After optimization, we progressively restrict Neu-PiG to l_{\max} grid levels and visualize the deformations, colored by a rigidity proxy. It shows that coarse levels capture smooth, global motion, whereas finer levels contribute localized, high-frequency detail, supporting the role of the multi-resolution latent grid.

Table 7. Ablation study on the smoothness weight λ of the base grid level ($l = 1$), evaluated on the AMA dataset. Neu-PiG remains robust under moderate variations of λ , while extreme values reduce reconstruction quality.

	CD [$\times 10^{-5}$] ↓	NC ↑	F-0.5% ↑	Corr. ↓
$\lambda^1 = 0$	0.58	0.928	0.977	0.022
$\lambda^1 = 0.08$	0.45	0.950	0.988	0.020
Ours ($\lambda^1 = 0.4$)	0.44	0.951	0.989	0.017
$\lambda^1 = 2$	0.46	0.948	0.986	0.021
$\lambda^1 = 10$	0.50	0.946	0.983	0.025

5.2. Input Conditions

We additionally study how different input conditions, such as varying number of input points and noise levels, affect the reconstruction performance of Neu-PiG.

Noise. We evaluate the robustness of Neu-PiG to noisy input data by adding Gaussian noise with varying strengths to the input points based on the size of the bounding box diagonal. As shown in Tab. 8, the method remains stable for noise levels up to 0.25–0.5%, after which the performance degrades as noise begins to be reconstructed as part of the surface geometry.

Point Cloud Resolution. We evaluate the reconstruction performance across varying input point cloud resolutions, focusing on the generalization behavior in comparison to DynoSurf, which similarly employs MLP-based transformation modeling. As shown in Tab. 9, DynoSurf exhibits a noticeable drop in performance with increasing target resolution, indicating limited scalability. In contrast, our method demonstrates consistent improvements with higher input densities, closely matching the behavior of the direct optimization approach PDG, and achieving superior reconstruction quality at all tested resolutions. The difference compared to DynoSurf is particularly noteworthy: although both methods employ MLPs to model per-point transformations, DynoSurf fails to benefit from denser point clouds. This robustness stems from our hierarchical latent representation, which efficiently captures both coarse and fine geometric structures.

5.3. Neural Model

We further analyze the design of the neural components in Neu-PiG and evaluate how different architectural choices of the decoder influence reconstruction quality and temporal stability.

Time Deformation Model. We investigate how modifications to the temporal deformation network influence reconstruction performance. This component models temporal




Figure 8. Visualization of the learned latent space for different smoothness weights. We perform a PCA analysis on the latent vectors and display the first principal component as a spatial field.

Table 8. Ablation study on robustness to input noise with varying strengths based on the bounding box diagonal. Neu-PiG maintains stable reconstruction quality up to 0.5% noise, after which performance degrades as noise is captured in the surface geometry.

Noise	CD [$\times 10^{-5}$] ↓	NC ↑	F-0.5% ↑	Corr. ↓
0.25%	0.69	0.925	0.967	0.021
0.5%	1.34	0.866	0.853	0.026
1%	4.45	0.681	0.540	0.049
2%	21.71	0.517	0.247	0.086

shape evolution conditioned on the latent embeddings introduced in Sec. 3.2. We systematically vary the dimensionality of the positional and normal latent vectors z_p and z_n , as well as the frequency range of the time encoding γ . In addition, we assess the impact of network capacity by adjusting the number of hidden units per layer. As shown in Tab. 10, Neu-PiG remains stable across most configurations. However, excessive reduction in model capacity leads to lower reconstruction quality, while increasing the dimensionality of z_p or γ causes the network to overfit more easily.

Rotation Modeling. Our deformation model predicts parameters that are mapped to 3D rotations and deformations, as described in Sec. 3.4. Several representations exist for mapping parameters to rotations, including quaternions, the Cayley transform, and the exponential map, and their choice can affect optimization stability, convergence behavior, and final reconstruction accuracy. To study this influence, we run our reconstruction using each representation under identical settings and compare performance over different amounts of maximal optimization epochs. The resulting metrics are summarized in Tab. 11. The results show that while all parameterizations achieve comparable accuracy, quaternions converge slightly faster and yield marginally better performance during early training.

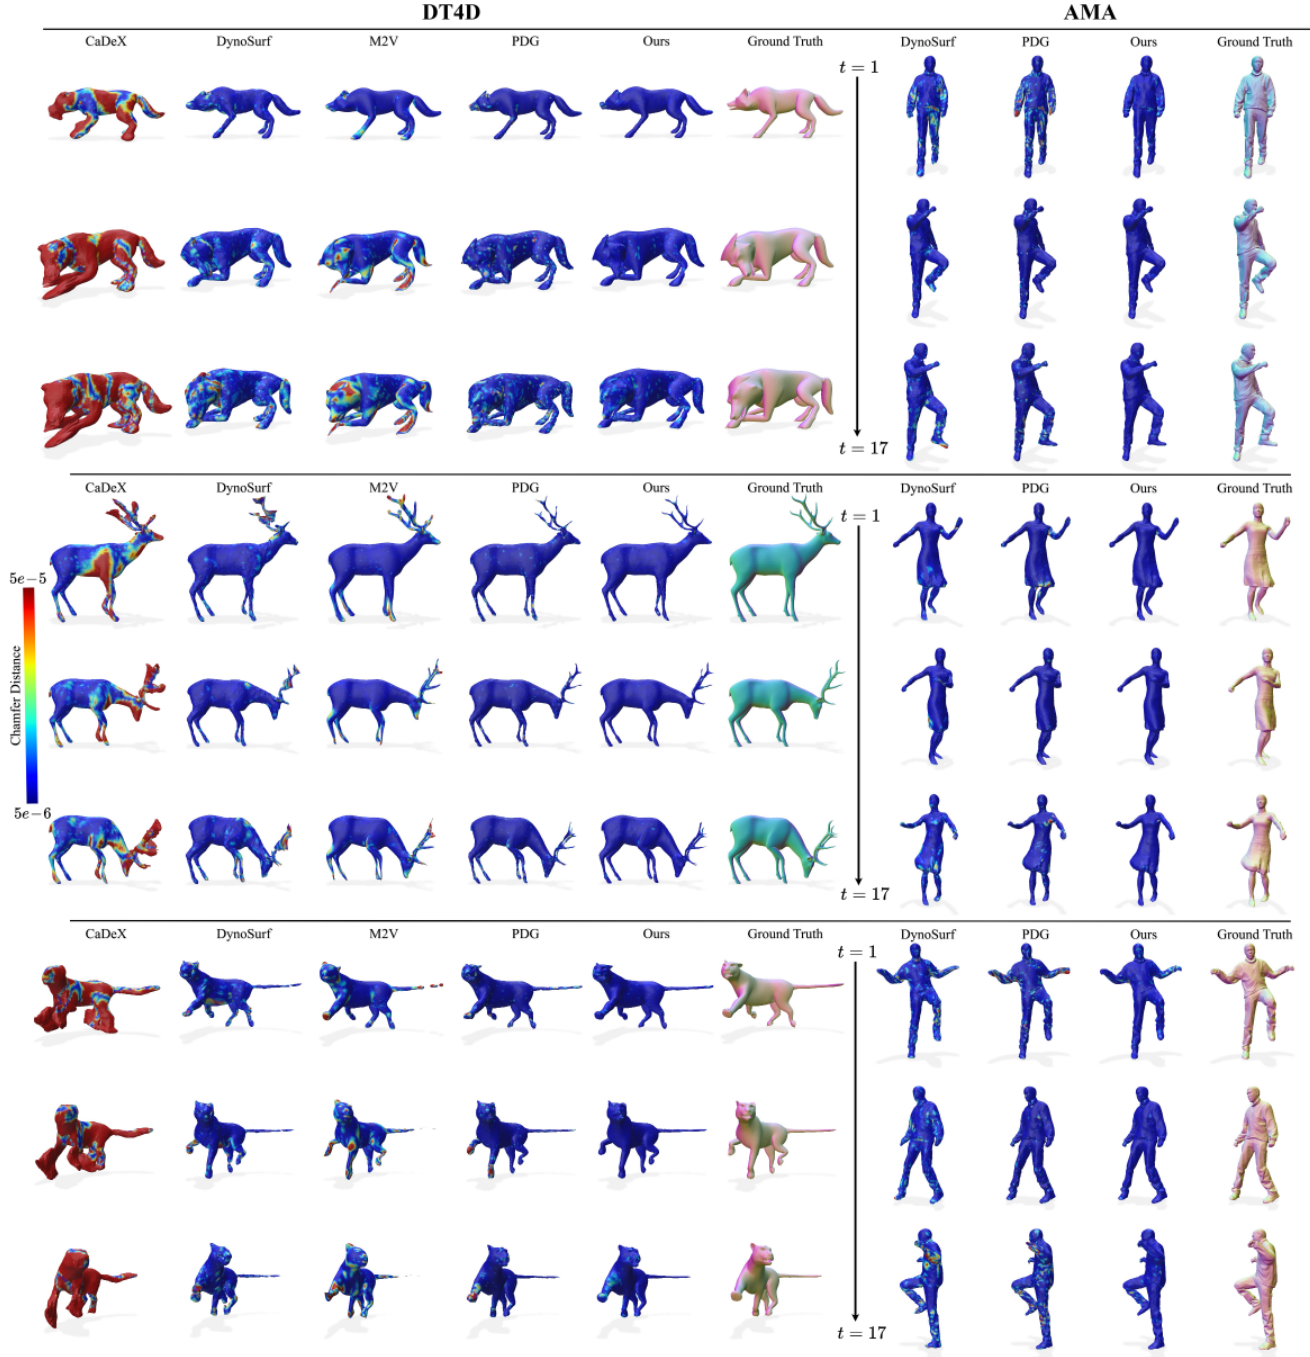


Figure 9. Extended qualitative results across diverse human and animal sequences. Neu-PiG consistently reconstructs temporally stable and detailed surfaces, even under large non-rigid deformations and challenging motion dynamics.

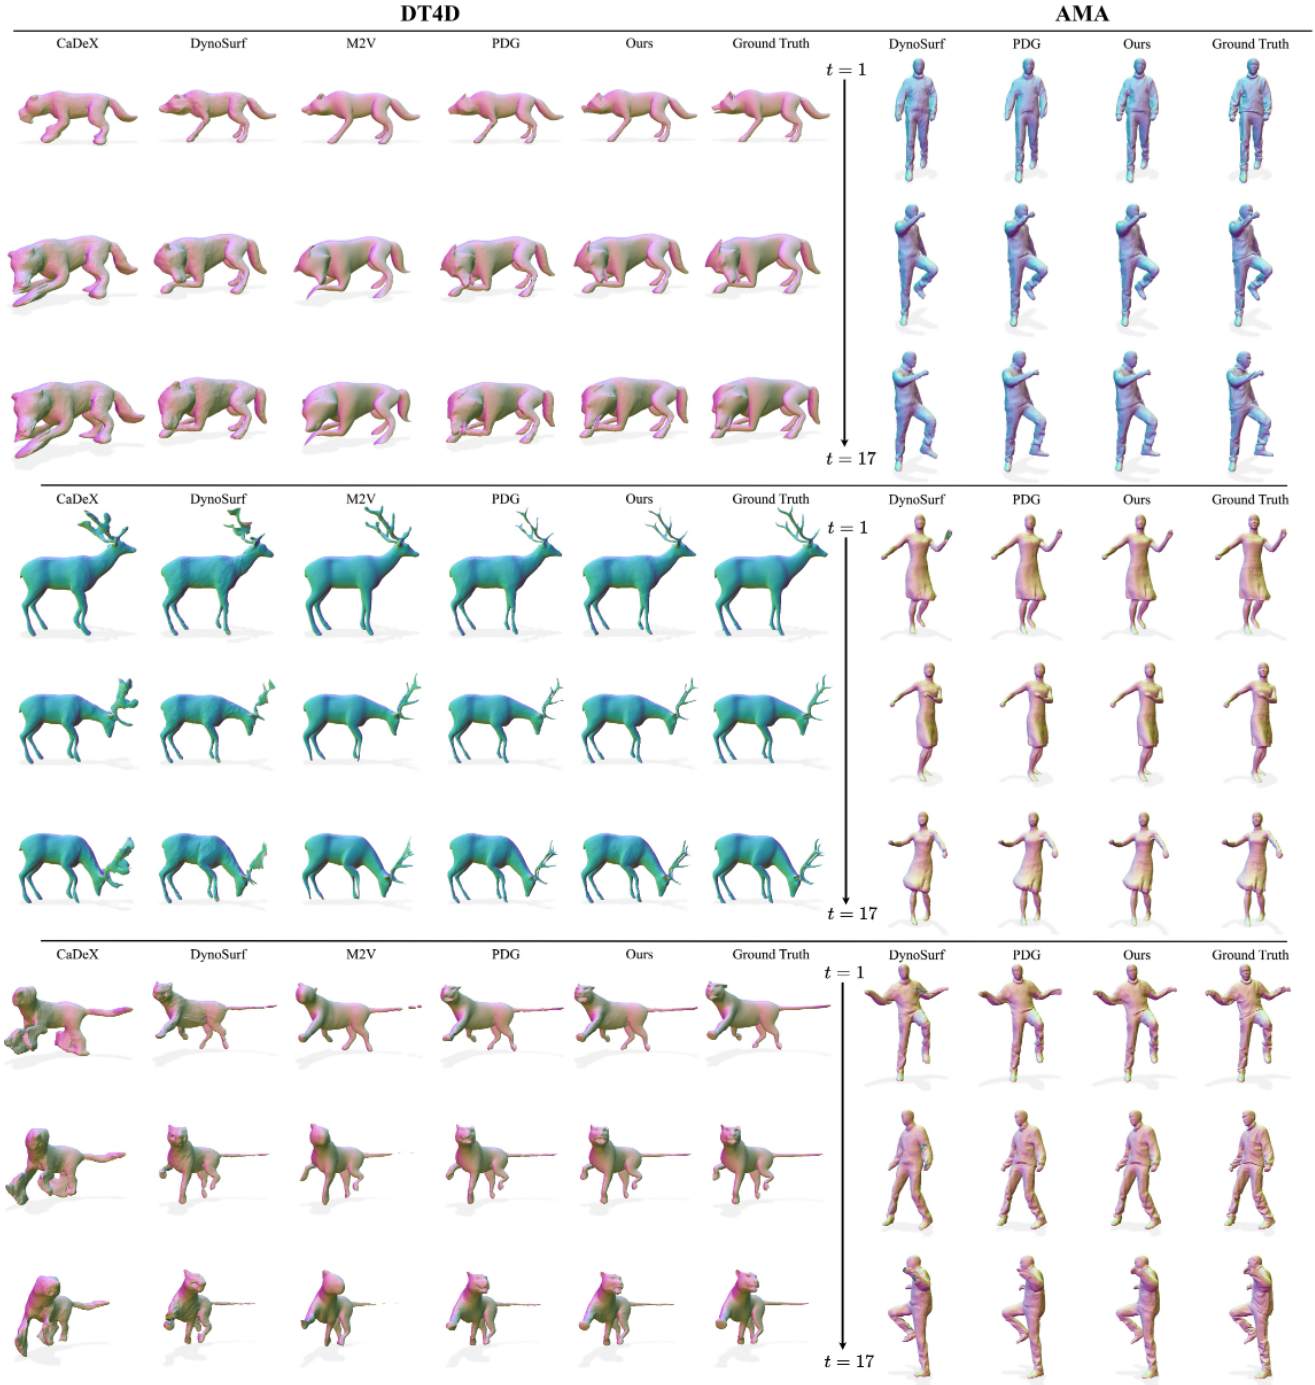


Figure 10. Extended visual results across diverse human and animal sequences. The surface normals of the deformed meshes are shown in comparison to ground truth surface normals.

Table 9. Reconstruction performance at varying input point cloud resolutions. Neu-PiG consistently improves with higher input densities, closely matching the behavior of PDG while maintaining superior quality across all settings. In contrast, DynoSurf shows limited scalability and degrading performance at higher resolutions.

$ \mathcal{P}_t $		CD $[\times 10^{-5}] \downarrow$	NC \uparrow	F-0.5% \uparrow	Corr. \downarrow
2500	DynoSurf	1.56	0.896	0.862	0.041
	PDG	0.79	0.922	0.948	0.037
	Ours	0.75	0.930	0.953	0.022
5000	DynoSurf	1.01	0.918	0.921	0.044
	PDG	0.47	0.939	0.985	0.030
	Ours	0.44	0.950	0.988	0.018
10000	DynoSurf	1.28	0.906	0.897	0.037
	PDG	0.40	0.950	0.993	0.027
	Ours	0.37	0.960	0.995	0.015
20000	DynoSurf	1.45	0.902	0.887	0.040
	PDG	0.37	0.959	0.996	0.023
	Ours	0.34	0.968	0.997	0.014

Table 10. Ablation study on the temporal deformation network architecture, evaluated on the AMA dataset with 40 time steps per sequence. We vary latent dimensionality, number of time-encoding frequencies, and network capacity to assess their effect on reconstruction quality. Neu-PiG remains stable under most configurations, though extreme reductions in capacity or excessive latent dimensionality lead to degraded performance or overfitting. Our default configuration is shown in the last row.

$ \psi $	$ \mathbf{z}_p $	$ \mathbf{z}_n $	$ \gamma $	CD $[\times 10^{-5}] \downarrow$	NC \uparrow	F-0.5% \uparrow	Corr. \downarrow
512	30	2	2	0.63	0.945	0.970	0.027
512	30	2	32	0.64	0.943	0.973	0.040
512	30	1	8	0.55	0.946	0.979	0.026
512	30	8	8	0.57	0.947	0.978	0.026
512	8	2	8	0.61	0.945	0.975	0.023
512	120	2	8	0.59	0.947	0.978	0.032
128	30	2	8	0.76	0.938	0.959	0.045
2048	30	2	8	0.54	0.948	0.980	0.021
512	30	2	8	0.53	0.947	0.981	0.019

Table 11. Comparison of different rotation parameterizations predicted by the deformation model. We evaluate quaternions, the Cayley transform, and the exponential map under identical optimization conditions. All representations achieve similar accuracy, while using quaternions leads to slightly faster convergence and improves performance with fewer epochs.

Epochs		CD $[\times 10^{-5}] \downarrow$	NC \uparrow	F-0.5% \uparrow	Corr. \downarrow
75	Cayley	1.96	0.921	0.866	0.047
	Exponential	1.62	0.924	0.886	0.043
	Quaternions	1.59	0.923	0.884	0.044
125	Cayley	0.72	0.940	0.955	0.029
	Exponential	0.68	0.941	0.961	0.027
	Quaternions	0.65	0.942	0.965	0.028
250	Cayley	0.52	0.947	0.978	0.021
	Exponential	0.54	0.948	0.976	0.021
	Quaternions	0.50	0.948	0.982	0.020
500	Cayley	0.47	0.950	0.985	0.019
	Exponential	0.47	0.950	0.985	0.018
	Quaternions	0.46	0.950	0.986	0.018
1000	Cayley	0.46	0.950	0.987	0.019
	Exponential	0.45	0.951	0.988	0.017
	Quaternions	0.44	0.951	0.988	0.017