Neural Hamiltonian Deformation Fields for Dynamic Scene Rendering

SIGGRAPH ASIA 2025 HONG KONG 香港
ASIA.SIGGRAPH.ORG/2025

Hai-Long Qin[1], Sixian Wang[2], Guo Lu[2], Jincheng Dai[1]

[1] Beijing University of Posts and Telecommunications (BUPT)
[2] Shanghai Jiao Tong University (SJTU)

Paper Page

Abstract

Representing dynamic scenes with realistic motion remains challenging as existing methods often produce physically implausible deformations. **We introduce NeHaD, a neural deformation field for dynamic Gaussian splatting governed by Hamiltonian mechanics.** Our key innovation replaces MLP-based deformation with Hamiltonian neural networks that model Gaussians evolving along energy-conserving trajectories in phase space, ensuring natural dynamics. We introduce Boltzmann equilibrium decomposition for energy-aware static/dynamic separation, and employ symplectic integration with rigidity constraints to handle real-world dissipation. Additionally, we extend NeHaD to adaptive streaming through scale-aware mipmapping. Extensive experiments demonstrate that NeHaD achieves physically plausible dynamic rendering with good quality-efficiency trade-offs, representing the **first** application of Hamiltonian mechanics to neural Gaussian deformation.

Motivation

Just as human cognition relies on physical intuition, dynamic scene rendering should respect physical laws to achieve realistic results. Current 4D Gaussian splatting methods produce physically implausible motions[1]. Our key insight: Gaussian covariances naturally exist on symplectic manifolds, making Hamiltonian mechanics ideal for deformation modeling. Using Hamiltonian neural networks, we learn energy-conserving trajectories from data, embedding physical knowledge into Gaussian deformation fields for stable, coherent dynamics.

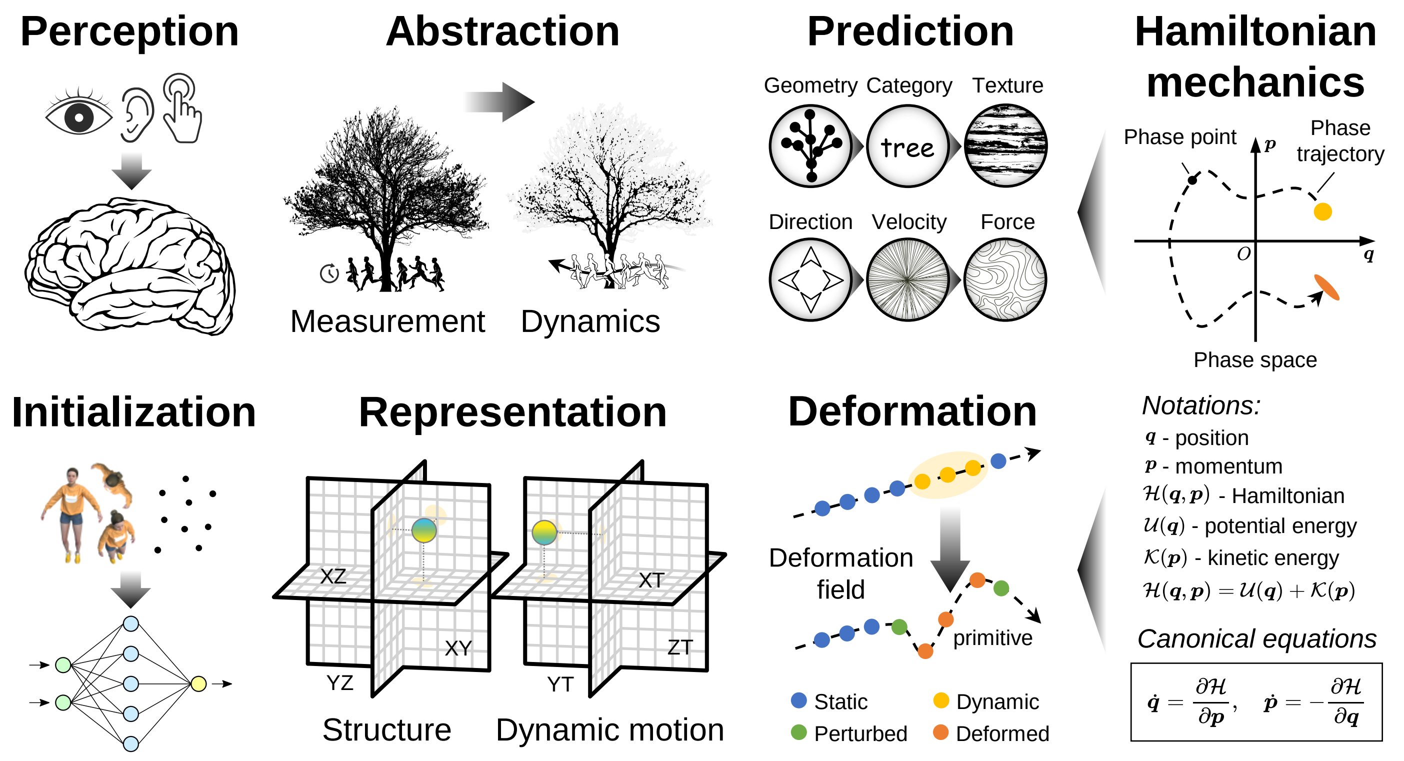

Figure 1. Both human cognition and scene rendering processes follow physical laws, with Hamiltonian mechanics offering mathematical frameworks aligned with physical intuition.

Methodology

Neural Hamiltonian Deformation Fields. Replace MLP-based deformation with Hamiltonian neural networks (HNN) that learn scalar potentials generating conservative and solenoidal vector fields[2].

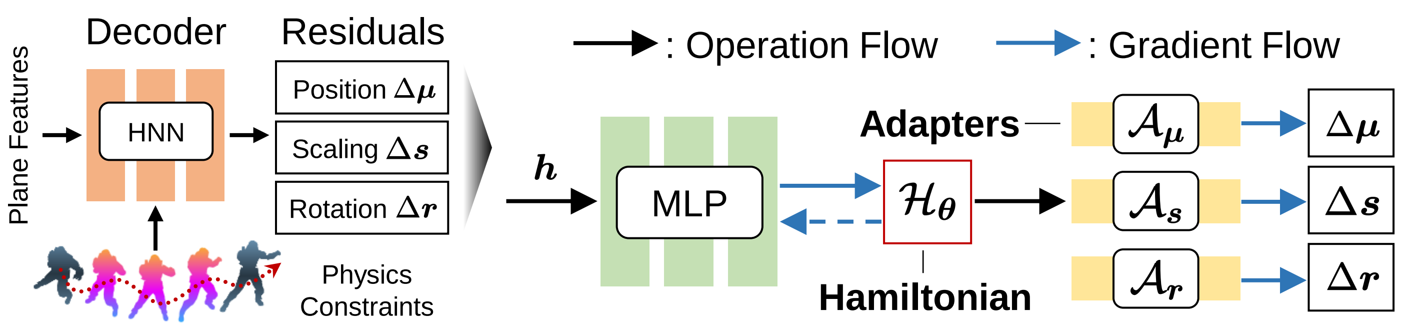

Figure 2. Through backpropagation of Hamiltonian gradients, the HNN optimizes vector fields and predicts Gaussian deformations via adapters.

Boltzmann Equilibrium Decomposition. *(continued)*

Adaptively separate static and dynamic Gaussians based on spatial-temporal energy deviation from equilibrium.

Physics-Informed Constraints.

- Second-order symplectic integration: Position Verlet integration maintains conservation.
- Local rigidity regularization: ARAP-motivated constraints clamp rotation angles.

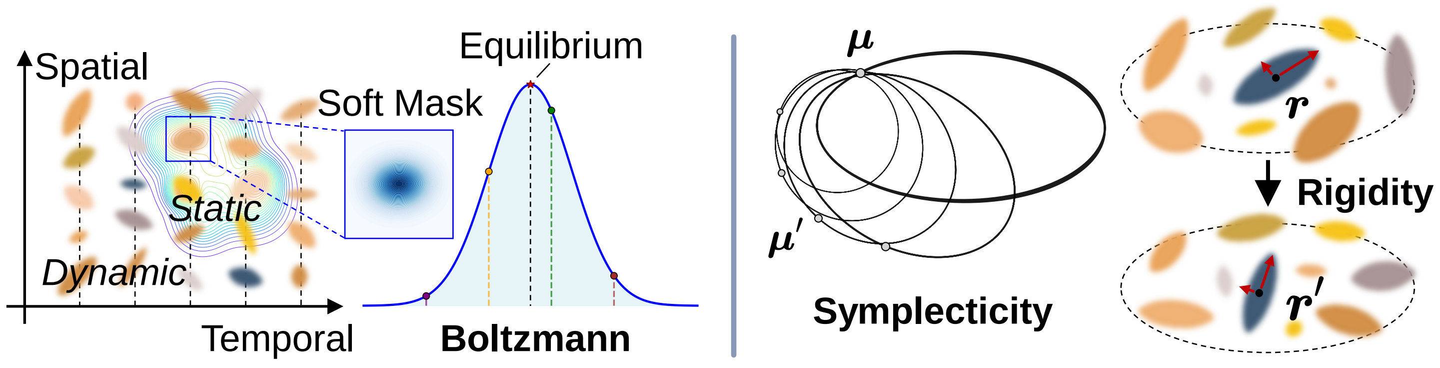

Figure 3. Left: decide which primitives should not be deformed with soft masks. Right: symplectic integration and rigidity regularization are used.

Experimental Results

Performance. State-of-the-art rendering quality across all benchmarks with physically plausible dynamics and real-time capability. More results available in the full text.

D-NeRF [Pumarola et al. 2021] (monocular, synthetic, 800×800)

Method	PSNR ↑	SSIM ↑	LPIPS ↓	Train Time ↓	FPS ↑
D-NeRF	29.68	0.947	0.058	48*hrs*	<1
TiNeuVox	32.74	0.972	0.051	28*min*	1.5
K-Planes	31.52	0.967	0.047	52*min*	0.97
HexPlane	31.04	0.97	0.04	**11.5*min***	2.5
4DGS	35.34	0.985	0.021	20*min*	82
SC-GS	40.26	0.992	0.009	30*min*	**164**
Ours	**40.91**	**0.995**	**0.008**	24*min*	62

HyperNeRF [Park et al. 2021b] (monocular, real-world, 536×960)

Method	PSNR ↑	MS-SSIM ↑	LPIPS ↓	Train Time ↓	FPS ↑
HyperNeRF	22.41	0.814	0.131	32*hrs*	<1
TiNeuVox	24.20	0.836	0.128	**30*min***	1
4DGS	25.24	0.845	0.116	34*min*	32
DeformGS	25.02	0.822	0.116	1.5*hrs*	13
SaRO-GS	25.38	0.850	0.110	1.2*hrs*	34
Grid4D	25.50	0.856	0.107	2.5*hrs*	**37**
Ours	**25.69**	**0.858**	**0.104**	45*min*	25

DyNeRF [Li et al. 2022] (multi-view, real-world, 1352×1014)

Method	PSNR ↑	D-SSIM ↓	LPIPS ↓	Train Time ↓	FPS ↑
DyNeRF	29.58	0.020	0.083	1344*hrs*	<1
HexPlane	31.70	0.014	0.075	12*hrs*	0.2
4DGS	31.17	0.016	0.049	42*min*	30
STG	32.05	0.014	0.044	10*hrs*	110
SaRO-GS	32.15	0.014	0.044	1.5*hrs*	32
Swift4D	32.23	0.014	0.043	**25*min***	**125**
Ours	**32.35**	**0.013**	**0.042**	50*min*	21

Table 1. Quantitative results.

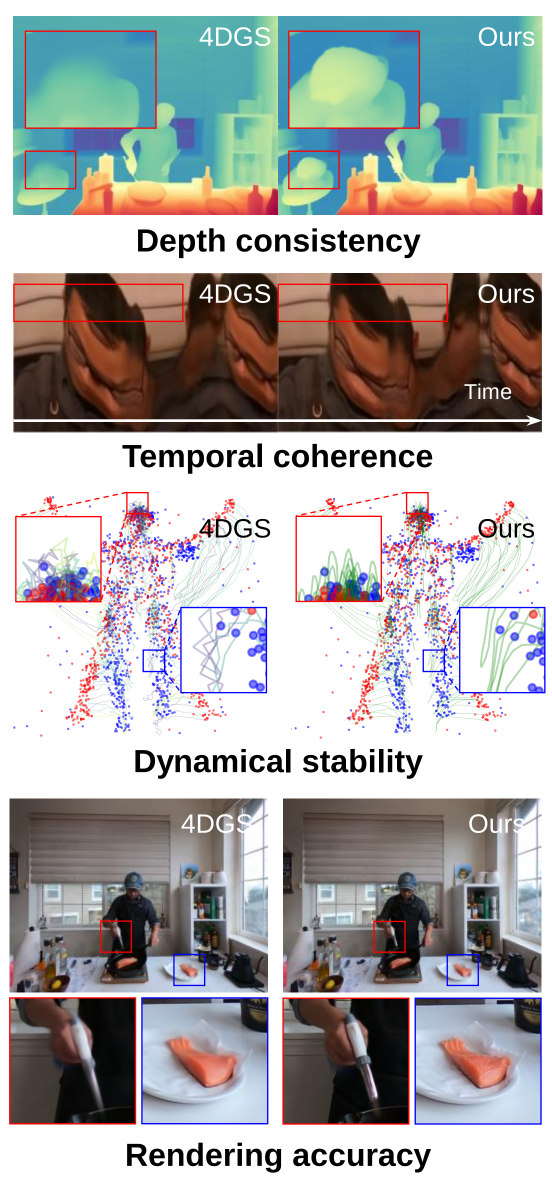

Figure 4. Qualitative results.

References.
[1] Wu G, et al. 4D Gaussian Splatting for Real-Time Dynamic Scene Rendering. CVPR 2024.
[2] Greydanus S, Dzamba M, Yosinski J. Hamiltonian Neural Networks. NeurIPS 2019.