

# 000 001 002 003 004 005 006 007 008 009 010 ODE-GS: LATENT ODES FOR DYNAMIC SCENE EXTRAPOLATION WITH 3D GAUSSIAN SPLATTING

005      **Anonymous authors**

006      Paper under double-blind review

## 009      ABSTRACT

011      We introduce ODE-GS, a novel approach that integrates 3D Gaussian Splatting with  
012      latent neural ordinary differential equations (ODEs) to enable future extrapolation  
013      of dynamic 3D scenes. Unlike existing dynamic scene reconstruction methods,  
014      which rely on time-conditioned deformation networks and are limited to interpola-  
015      tion within a fixed time window, ODE-GS eliminates timestamp dependency by  
016      modeling Gaussian parameter trajectories as continuous-time latent dynamics. Our  
017      approach first learns an interpolation model to generate accurate Gaussian trajec-  
018      tories within the observed window, then trains a Transformer encoder to aggregate  
019      past trajectories into a latent state evolved via a neural ODE. Finally, numerical  
020      integration produces smooth, physically plausible future Gaussian trajectories,  
021      enabling rendering at arbitrary future timestamps. On the D-NeRF, NVFi, and  
022      HyperNeRF benchmarks, ODE-GS achieves state-of-the-art extrapolation perfor-  
023      mance, improving metrics by 19.8% compared to leading baselines, demonstrating  
024      its ability to accurately represent and predict 3D scene dynamics.

## 025      1 INTRODUCTION

027      Recently, 3D Gaussian Splatting (3DGS) methods have emerged as an effective approach for dynamic  
028      scene reconstruction. By training on images taken from a time-dependent 3-dimensional (3D)  
029      scene, such methods enable photorealistic novel view synthesis (NVS) for any time within the  
030      observed window. However, the task of *prediction*—extrapolating future scene dynamics from past  
031      observations—remains largely underexplored. Performing such a task is a well-studied capability  
032      that humans possess (Rao & Ballard, 1999; Mrótek & Soechting, 2007; Battaglia et al., 2013; Khoei  
033      et al., 2017), while mirroring this capability in intelligent systems is of great interest for applications  
034      such as self-driving, robotics, and augmented reality. Our focus is then to bridge this gap and enable  
035      the ability to forecast future 3D states in the context of dynamic scene reconstruction, which we will  
036      refer to as **dynamic scene extrapolation**.

037      Performing dynamic scene extrapolation is fundamentally more challenging than reconstruction.  
038      Existing methods are primarily designed for temporal *interpolation* within the observed window,  
039      where arbitrary view points at given time can be reconstructed by conditioning models on timestamps.  
040      In contrast, temporal *extrapolation* is inherently under-constrained, as there exist infinitely many pos-  
041      sible future dynamics given our past observations. Therefore, popular dynamic scene reconstruc-  
042      tion methods, such as TiNeuVox (Fang et al., 2022), Deformable 3D Gaussians (Yang et al., 2024), and  
043      4D Gaussian Splatting (Wu et al., 2024), excel at “filling the gaps” between observed timestamps,  
044      but degrade when extended beyond them. In such cases, future timestamps fall outside the training  
045      distribution, leading to out-of-distribution (OOD) failures.

046      As with many under-constrained problems, our task then becomes estimating the *most likely* future  
047      dynamics given existing observations. Incorporating physical assumptions can substantially narrow  
048      this solution space. An example of such an assumption is the spatio-temporal smoothness of motion,  
049      which provides strong constraints on the evolution of scene dynamics. Differential equations have  
050      long served as principled tools for describing the evolution of physical systems (Chen et al., 2018),  
051      and **ordinary differential equations (ODEs)** in particular offer a natural formalism for representing  
052      continuous and physically plausible motion trajectories.

053      In light of these considerations, we propose to represent the temporal evolution of a dynamic scene in  
a continuous-time latent space, where the evolution is governed by an ODE. In this way, we impose a

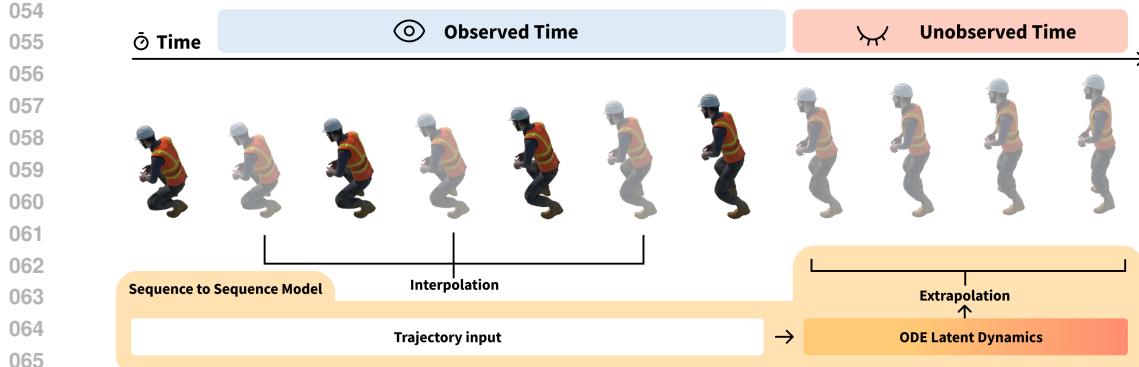


Figure 1: Unlike existing methods that focus on interpolation, i.e., reconstructing novel scene views at unseen timestamps *within the observed time window*, we focus on extrapolation, i.e., extending scene dynamics beyond the observed times, by first training a representation of the observed scene and then using a sequence-to-sequence model to reconstruct *future* novel views via latent ODE dynamics.

physical assumption of smooth motion into the predictive model (Chen et al., 2018; Rubanova et al., 2019), while retaining 3DGS’s capability of high-fidelity rendering. Unlike prior approaches that condition directly on explicit timestamps—leading to OOD failures—we reformulate dynamic scene extrapolation as a sequence-to-sequence forecasting problem, which naturally aligns with 3DGS’s explicit representation of the 3D scene. Specifically, we first encode a temporal sequence of Gaussian parameters using a Transformer (Vaswani et al., 2017) into a latent state that embeds past motion. We then model the temporal evolution of this latent state via a neural ODE, where a neural network parameterizes its velocity field. Extrapolation is thus realized by numerically integrating the latent dynamics beyond the observed window, yielding future latent states. These states are subsequently decoded back into Gaussian parameters, which can be rendered into novel future views. To further constrain extrapolation, we incorporate additional physical priors as lightweight regularizers during training. These include penalties that encourage smoothness in both the latent trajectories and the decoded Gaussian parameters, reinforcing the assumption of continuous motion.

The resulting method, **ODE-GS** (Ordinary Differential Equation-based Gaussian Splatting), decouples scene reconstruction from temporal forecasting. We first optimize a Gaussian interpolation model within the observed window, by training a set of canonical Gaussian parameters and a deformation network. This interpolation model is then frozen and used as a data generator, producing temporal trajectories of 3D Gaussian parameters. We train the Transformer-Latent ODE on these trajectories, but crucially, the model is conditioned only on a partial prefix of each sequence and learns to extrapolate the remainder. This training setup equips ODE-GS to extrapolate beyond the observed time horizon at inference, by integrating forward in time via our latent ODE and decoding back into future Gaussian parameters.

Our contributions are:

- We propose **ODE-GS**, which integrates 3D Gaussian Splatting with a Transformer-based latent ODE. We model scene dynamics as continuous trajectories in latent space to enforce smoothness priors, enabling stable extrapolation beyond the observed window.
- We validate our modeling strategy that decouples dynamic scene reconstruction from temporal forecasting. By first optimizing an interpolation model for reconstruction and then training the latent ODE forecaster on generated trajectories, ODE-GS avoids direct reliance on timestamp conditioning and mitigates out-of-distribution failures.
- We incorporate inductive biases as regularizers that encourage smoothness in both latent and decoded trajectories, improving stability and rendering quality at extrapolated timestamps.
- ODE-GS achieves state-of-the-art performance in dynamic scene extrapolation, improving over the best existing extrapolation method by an average of 21.4% PSNR, 7.4% SSIM, and 30.5% LPIPS across synthetic (D-NeRF, NVFi) and real-world (HyperNeRF) datasets.

108 

## 2 RELATED WORK

110 **Novel View Synthesis (NVS).** Recent advances in NVS have explored a diverse range of approaches,  
 111 including explicit mesh-based representations (Broxton et al., 2020; Dou et al., 2016; Newcombe  
 112 et al., 2015; Orts-Escalano et al., 2016), and implicit neural volume representations (Lombardi et al.,  
 113 2019). Among these, Neural Radiance Fields (NeRF) (Mildenhall et al., 2021) have emerged as a  
 114 dominant paradigm due to their leading performance in NVS. The foundational success of NeRF  
 115 has led to numerous extensions for dynamic scene reconstruction (Attal et al., 2023; Du et al., 2021;  
 116 Ost et al., 2021; Park et al., 2021a; Pumarola et al., 2021), enabling applications such as monocular  
 117 video-based scene reconstruction (Gao et al., 2021; Li et al., 2021; Tretschk et al., 2021), editable  
 118 scene representations (Kania et al., 2022; Park et al., 2021a), and human-centered reconstructions  
 119 (Peng et al., 2021b;a). Notably, NVFi (Li et al., 2023) investigates the problem of dynamic scene  
 120 extrapolation by adding geometric priors. However, its dependence on explicit timestamps during  
 121 training induces out-of-distribution errors when extrapolating.

122 **3DGS and Dynamic Scene Modeling.** 3DGS (Kerbl et al., 2023), first used for static NVS,  
 123 has become a popular choice for representing dynamic scenes due to its speed and explicit nature  
 124 (Huang et al., 2024; Li et al., 2024). Deformable 3D Gaussians (Yang et al., 2024) learn Gaussians  
 125 in a canonical space with a deformation network. By training both the canonical Gaussians and  
 126 the deformation network simultaneously, they enable continuous-time rendering by transforming  
 127 canonical Gaussians into arbitrary time within the training window. Concurrently, 4D Gaussian  
 128 Splatting (Wu et al., 2024) took the same canonical-deform strategy and additionally integrated 4D  
 129 neural voxels inspired by HexPlane (Cao & Johnson, 2023). Additional efforts, such as explicit  
 130 time-variant Gaussian features (Luiten et al., 2024), achieved interactive frame rates and enabled  
 131 flexible editing. GaussianVideo (Bond et al., 2025) uses neural ODEs, but for learning smooth  
 132 camera trajectories rather than scene motion. However, most 3DGS-based approaches rely on a  
 133 time-conditioned deformation field, making these models excel in interpolation tasks but unable  
 134 to extrapolate into unseen time in the future. GaussianPrediction (Zhao et al., 2024) has recently  
 135 explored this issue by combining a superpoint strategy with Graph Convolution Networks (GCN)  
 136 that directly conditions on past motion instead of time, but is only capable of sampling at discrete  
 137 steps when extrapolating.

138 **Neural Ordinary Differential Equations.** Neural Ordinary Differential Equations (Neural ODEs)  
 139 (Chen et al., 2018) introduced a novel approach for continuous-depth neural networks. Instead  
 140 of defining discrete layers, a Neural ODE specifies the continuous dynamics of a hidden state  
 141 using a neural network that parameterizes its derivative. The network’s output is then determined  
 142 by a numerical ODE solver that integrates these learned dynamics over a specified interval. Key  
 143 advantages include memory-efficient training via the adjoint sensitivity method, inherent handling of  
 144 irregularly-sampled data, and adaptive computation.

145 Relevant to dynamic modeling, Latent ODEs (Rubanova et al., 2019) typically encode an input  
 146 sequence into an initial latent representation whose continuous-time evolution is then governed by a  
 147 Neural ODE. A decoder can subsequently map these evolving latent states back to the observation  
 148 space at arbitrary times. This is effective for modeling continuous trajectories and is often combined  
 149 with Variational Autoencoders (VAEs), as in ODE2VAE (Yildiz et al., 2019), to learn distributions  
 150 over latent paths and capture uncertainty.

151 The synergy between recurrent methods and Neural ODEs has also been explored. For instance,  
 152 GRU-ODE (De Brouwer et al., 2019) adapts GRU-like gating mechanisms to continuously evolving  
 153 states, while ODE-RNN (Rubanova et al., 2019) interleaves discrete RNN updates at observation  
 154 points with continuous ODE-based evolution.

155 

## 3 METHODOLOGY

157 In the following sections, we formalize our problem setting and detail each component of ODE-GS:  
 158 first, the interpolation model that generates dense Gaussian trajectories within the observed window  
 159 (Sec. 3.1); second, the Transformer-based latent ODE architecture for extrapolation (Sec. 3.2);  
 160 third, a dynamic sampling strategy to train the extrapolation model on different forecasting horizons  
 161 (Sec. 3.3); and finally, the training and regularization objectives that ensure physically plausible and  
 162 stable extrapolated dynamics (Sec. 3.4). Implementation details are discussed in Sec. B.1.

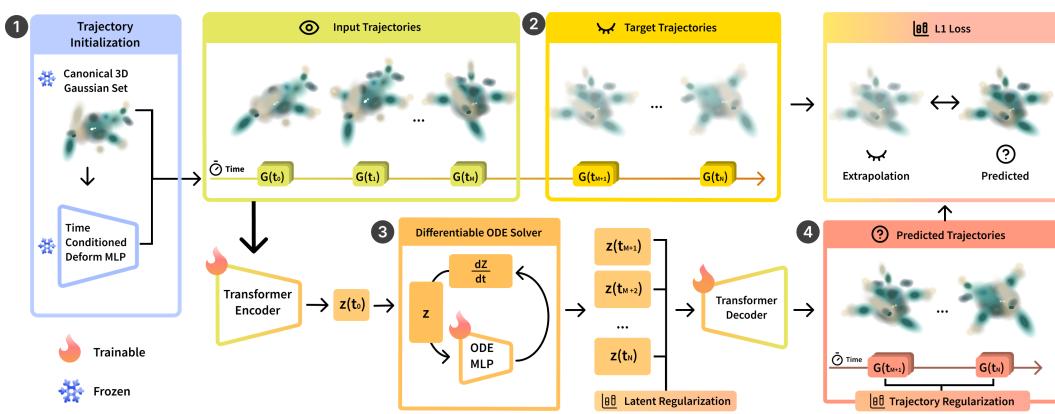


Figure 2: 1: We initialize temporal trajectories of 3D Gaussian parameters using the frozen interpolation model, which consists of the canonical 3D Gaussian set and a time-conditioned deformation MLP. These trajectories lie entirely within the observed temporal window. 2: Through our dynamic sampling strategy, each Gaussian trajectory is sampled into multiple observed prefix (input) and a held-out suffix (target) trajectories, providing training pairs for the Transformer latent ODE. 3: Latent-ODE training encodes the observed prefix with a Transformer, infers a latent initial state, and evolves it forward with a neural ODE. 4: A decoder maps the latent path back to Gaussian parameters, which are supervised against the ground-truth suffixes via an L1 loss and smoothness regularizers.

**Notation and Formalization.** Let  $\{I_i\}$  be a set of calibrated RGB images, with corresponding camera poses  $\{V_i\}$  capturing a dynamic 3D scene at timestamps  $\{t_i\}$ ,

$$\mathcal{D} = \{(I_i, V_i, t_i)\}_{i=1}^N, \quad I_i : \mathbb{R}^{3 \times H \times W}, \quad V_i \in SE(3), \quad t_i \in \mathbb{R}, \quad (1)$$

we aim to learn a continuous-time rendering operator  $\mathcal{F} : \mathbb{R} \times SE(3) \rightarrow \mathbb{R}^{3 \times H \times W}$  that generates RGB images for any time  $t$  and camera pose  $V$ . The operator decomposes as:

$$\mathcal{F}(t, V) = \mathcal{R}(\mathcal{G}(t), V), \quad \mathcal{G}(t) = \begin{cases} \bar{\mathcal{G}} + \mathcal{D}_\omega(t, \bar{\mathcal{G}}) & \text{if } t_{\min} \leq t \leq t_{\max}, \\ \mathcal{E}_\phi(\gamma, t) & \text{if } t > t_{\max}, \end{cases} \quad (2)$$

where  $\bar{\mathcal{G}} = \{\bar{G}_k\}_{k=1}^M$  is a learnable canonical set of 3D Gaussians,  $\mathcal{D}_\omega$  is an interpolation deformation function,  $\mathcal{E}_\phi$  is a sequence-to-sequence extrapolation model with Transformer Latent ODE architecture, and  $\mathcal{R}$  is the differentiable rasterizer from (Kerbl et al., 2023).  $\gamma$  is a sequence of past Gaussian parameters that serve as input, which is detailed in Sec. B.

Each canonical Gaussian  $\bar{G}_k = (\mu_k, q_k, s_k, c_k, \alpha_k)$  comprises position  $\mu_k \in \mathbb{R}^3$ , quaternion  $q_k \in \mathbb{R}^4$ , scales  $s_k \in \mathbb{R}^3$ , opacity  $\alpha_k \in \mathbb{R}$ , and spherical harmonics (SH) coefficients  $c_k \in \mathbb{R}^d$ , where  $d$  is the number of SH functions used, usually set to 3. For each  $k$ , we keep  $\alpha_k$  and  $c_k$  consistent across time, so that only  $\mu_k$ ,  $q_k$ , and  $s_k$  are time-dependent. For simplicity, from now on,  $G(t)$  refers to only these three parameters. We may then derive the rotation matrix  $R_k \in \mathbb{R}^{3 \times 3}$  from  $q_k$ , the scaling diagonal matrix  $S_k \in \mathbb{R}^{3 \times 3}$  from  $s_k$ , and the covariance matrix for each Gaussian by  $\Sigma_k = R_k S_k S_k^\top R_k^\top$ .

The differentiable rasterizer (Kerbl et al., 2023)  $\mathcal{R}$  renders images by projecting each 3D Gaussian onto the image plane, computing per-pixel alpha compositing with front-to-back blending:

$$C(p) = \sum_{k \in G(p)} c_k \alpha_k \prod_{j=1}^{k-1} (1 - \alpha_j), \quad (3)$$

where  $G(p)$  are Gaussians affecting pixel  $p$ ,  $c_k$  is the color, and  $\alpha_k$  is the opacity after 2D projection. We first learn the interpolation model  $\bar{\mathcal{G}}$  and  $\mathcal{D}_\omega$ , then freeze them and train  $\mathcal{E}_\phi$  for extrapolation.

### 3.1 INTERPOLATION MODEL FOR TRAJECTORY GENERATION

To obtain trajectories of Gaussian parameters within the observed window, we adopt a canonical-plus-deformation strategy (Yang et al., 2024; Wu et al., 2024) that has become standard in dynamic scene

216 modeling. A canonical set of 3D Gaussians,  $\bar{\mathcal{G}}$ , represents the static reference configuration of the  
 217 scene, while a lightweight, time-conditioned deformation Multi-Layer-Perceptron (MLP)  $\mathcal{D}_\omega$  predicts  
 218 offsets for position, rotation, and scale at each timestamp  $t$ . This enables continuous interpolation of  
 219 Gaussian states across time. The interpolation model is trained using a photometric reconstruction  
 220 objective:

$$221 \quad \mathcal{L}_{\text{render}} = (1 - \lambda) \cdot \|\hat{I}_i - I_i\|_1 + \lambda \cdot (1 - \text{SSIM}(\hat{I}_i, I_i)), \quad (4)$$

222 Where  $\hat{I}_i$  is rendered using the differentiable rasterizer:  $\hat{I}_i = \mathcal{R}(\bar{\mathcal{G}} + \mathcal{D}_\omega(t, \bar{\mathcal{G}}), V_i)$ . After training,  
 223 we freeze both the canonical Gaussians  $\bar{\mathcal{G}}$  and the deformation network  $\mathcal{D}_\omega$ , so that given any time  $t$   
 224 within the observed time window,  $G_k(t)$  may be generated for all  $k$ . SSIM refers to the Structural  
 225 Similarity Index Measure.

### 227 3.2 LATENT ODE MODEL FOR EXTRAPOLATION

229 Viewing temporal prediction as a mapping between an observed sequence and a future sequence,  
 230 we model scene dynamics as a sequence-to-sequence problem. Our extrapolation module  $\mathcal{E}_\phi$  is a  
 231 Transformer Latent ODE that predicts future dynamics from past Gaussian trajectories.

232 Given an input sequence

$$233 \quad \gamma_k = \{G_k(t_j)\}_{j=1}^{T_c},$$

235 uniformly sampled from a context window of length  $T_c$  for Gaussian  $k$ , we embed each step and add  
 236 sinusoidal positional encodings to preserve temporal order. The resulting sequence is processed by a  
 237 Transformer encoder

$$238 \quad \mathcal{F}_\phi : \mathbb{R}^{T_c \times 10} \rightarrow \mathbb{R}^d,$$

239 yielding a latent representation  $z(t_0) \in \mathbb{R}^d$  that summarizes past dynamics. This latent state initializes  
 240 a neural ODE, parameterized by an MLP:

$$241 \quad \dot{z} = \frac{dz}{dt} = f_\theta(z(t)).$$

243 Numerical integration produces a continuous latent trajectory  $z(t)$  for any  $t > t_{\max}$ . A decoder then  
 244 maps the evolved latent states back to Gaussian parameters:

$$245 \quad \delta_\psi : \mathbb{R}^d \rightarrow \mathbb{R}^{10}, \hat{G}_k(t) = \delta_\psi(z(t)).$$

247 This combination of sequence encoding and continuous latent evolution allows  $\mathcal{E}_\phi$  to generate smooth  
 248 Gaussian trajectories without explicit timestamp embeddings, enabling extrapolation to arbitrary  
 249 future horizons where unseen timestamps are no longer out-of-distribution.

### 252 3.3 DYNAMIC TRAJECTORY SAMPLING

253 To effectively train the extrapolation module, it is essential to expose the model to prediction tasks  
 254 spanning a wide range of forecasting horizons. Therefore, unlike common approaches where sampled  
 255 trajectories always occupy the same time span uniformly (Li et al., 2023), we design a dynamic  
 256 trajectory sampling strategy. The pre-trained interpolation model provides continuous trajectories  
 257 of Gaussian parameters, from which we extract an observed prefix and a future suffix. The prefix  
 258 is sampled at fixed intervals to ensure consistent input dimensionality, while the suffix varies in  
 259 temporal span depending on the selected starting time.

260 Our training dataset is constructed as the union over all possible prefix–suffix splits of Gaussian  
 261 trajectories across all Gaussians and starting times. This design compels the extrapolation model to  
 262 learn from both short-term and long-term forecasting instances within a unified training procedure,  
 263 thereby encouraging robust generalization beyond the observed window, as described in Sec. B. At  
 264 inference, the model conditions on the final observed prefix and extrapolates Gaussian trajectories  
 265 forward in time, which are subsequently rendered into novel frames. At test time, we take the *final*  
 266 *context segment* of length  $T_c$  from the observed window, generated by the interpolation model  $\mathcal{D}_\omega$ .  
 267 This sequence is encoded and evolved forward by the Latent ODE to produce extrapolated Gaussians  
 268  $\hat{G}_k(t)$  for  $t > t_{\max}$ . The full scene  $\hat{\mathcal{G}}(t) = \{\hat{G}_k(t)\}_{k=1}^M$  is then rendered with the differentiable  
 269 rasterizer:

$$\hat{I}(t, V) = \mathcal{R}(\hat{\mathcal{G}}(t), V).$$

270 3.4 TRAINING OBJECTIVE  
271

272 **Extrapolation Loss.** Given a training pair consisting of a context trajectory  $\gamma_c$  and its corre-  
273 sponding target trajectory  $\gamma_e = \{G_k(t_j)\}_{j=1}^{N_e}$ , the extrapolation module  $\mathcal{E}_\phi$  produces predictions  
274  $\hat{\gamma}_e = \{\hat{G}_k(t_j)\}_{j=1}^{N_e}$ . The extrapolation loss supervises these predictions by minimizing the mean  
275 absolute error (L1) between target Gaussian parameters generated by the interpolation model and the  
276 predicted parameters by the extrapolation model across the target temporal window:

$$277 \quad 278 \quad 279 \quad 280 \quad \mathcal{L}_e = \frac{1}{N_e} \sum_{j=1}^{N_e} \|\hat{G}_k(t_j) - G_k(t_j)\|_1. \quad (5)$$

281 This objective ensures that the predicted Gaussian trajectories align closely with true future dynamics  
282 before additional regularization terms are applied.

283 To ensure physical plausibility and prevent overfitting to the observed trajectory, we also introduce  
284 two complementary regularizations that promote smoothness in latent dynamics and 3D trajectories.

285 **The latent regularization** penalizes high-frequency oscillations in the learned ODE function. Given  
286 the latent trajectory  $z(t)$  evolved by the neural ODE with velocity field  $\dot{z}(t) = f_\theta(z(t))$ , we approxi-  
287 mate the latent acceleration through finite differencing:

$$288 \quad 289 \quad 290 \quad 291 \quad \mathcal{R}_{\text{latent}} = \frac{1}{N_e - 1} \sum_{j=1}^{N_e - 1} \left\| \frac{f_\theta(z(t_{j+1})) - f_\theta(z(t_j))}{\Delta t_j} \right\|_2^2 \quad (6)$$

292 where  $\{t_j\}_{j=1}^{N_e}$  are the extrapolation timestamps,  $\Delta t_j = t_{j+1} - t_j$  the step size,  $f_\theta$  the neural ODE.

293 **The trajectory regularization** enforces smoothness directly in 3D space by penalizing accelerations  
294 of the Gaussian positions. For each Gaussian  $G_k(t)$ ,  $\mu_k(t) \in \mathbb{R}^3$  is its position at time  $t$ . We compute:

$$295 \quad 296 \quad 297 \quad 298 \quad \mathcal{R}_{\text{traj}} = \frac{1}{MN_e} \sum_{k=1}^M \sum_{j=1}^{N_e - 2} \left\| \frac{v_k(t_{j+1}) - v_k(t_j)}{\Delta t_j} \right\|_2^2 \quad (7)$$

299 where  $M$  is the number of Gaussians, and the velocity is approximated as:

$$300 \quad 301 \quad 302 \quad 303 \quad v_k(t_j) = \frac{\mu_k(t_{j+1}) - \mu_k(t_j)}{\Delta t_j} \quad (8)$$

304 **Adaptive weighting for regularization.** In early stages, strong regularization may inhibit the model  
305 from learning the meaningful dynamics. To address this, we introduce an *adaptive regularization*  
306 *weighting mechanism* that dynamically adjusts the contribution of regularization throughout training.  
307 At each iteration, we estimate the model’s convergence state using the Exponential Moving Average  
308 (EMA) of the trajectory prediction loss, which provides a stable signal compared to the raw loss  
309 that may fluctuate due to oscillations. As training progresses and the extrapolation loss decreases,  
310 the regularization weight is increased. This gradually biases the model toward selecting smoother  
311 trajectories among the many plausible solutions, thereby guiding fine-grained trajectory predictions  
312 toward stable convergence.

313 The final training loss then becomes:

$$314 \quad 315 \quad 316 \quad 317 \quad \mathcal{L} = \mathcal{L}_e + s_t (\lambda_{\text{latent}} \mathcal{R}_{\text{latent}} + \lambda_{\text{traj}} \mathcal{R}_{\text{traj}}) \quad (9)$$

318 where  $\lambda_{\text{latent}}$  and  $\lambda_{\text{traj}}$  are hyperparameters controlling the regularization strength, and  $s_t$  is an adaptive  
319 weighting term. For more details, refer to Sec. A.0.1.

320 4 EXPERIMENTS  
321

322 In this section, we evaluate the effectiveness of ODE-GS for extrapolating dynamic 3D scenes.  
323 We begin by presenting quantitative results in Section 4.1 to assess rendering quality on unseen  
324 future timestamps across D-NeRF (Pumarola et al., 2021), NVFi (Li et al., 2023), and HyperNeRF  
325 (Park et al., 2021b) followed by qualitative evaluations in Section 4.2 to demonstrate quality of the  
326 extrapolated dynamics through visual coherence.

Table 1: Quantitative extrapolation results on the D-NeRF dataset. Metrics reported include PSNR, SSIM, and LPIPS-vgg. The best metric is highlighted in red, and second best is highlighted in orange.

Method	Lego			Mutant			Standup			Trex		
	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)
TiNeuVox-B	23.34	.9102	.0942	24.40	.9282	.0700	21.77	.9169	.0927	20.72	.9284	.0751
4D-GS	<b>24.25</b>	.9150	.0810	22.48	.9300	.0520	18.61	.9180	.0840	<b>23.83</b>	<b>9460</b>	.0510
Deformable-GS	23.25	<b>.9349</b>	<b>.0579</b>	24.45	.9310	.0461	21.37	.9124	.0844	20.74	.9421	<b>.0465</b>
GaussianPredict	12.25	.7594	.2325	<b>.27.12</b>	<b>.9514</b>	<b>.0285</b>	<b>26.91</b>	<b>.9456</b>	<b>.0465</b>	21.52	.9443	<b>.0437</b>
<b>Ours</b>	<b>25.74</b>	<b>.9378</b>	<b>.0547</b>	<b>34.53</b>	<b>.9804</b>	<b>.0126</b>	<b>28.91</b>	<b>.9557</b>	<b>.0360</b>	<b>22.04</b>	<b>.9475</b>	.0485
Jumpingjacks				Bouncingballs			Hellwarrior			Hook		
Method	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)
	19.87	.9115	.0954	25.92	.9677	.0853	29.36	.9097	.1138	21.05	.8817	.1033
4D-GS	19.95	<b>.9270</b>	<b>.0770</b>	<b>29.55</b>	<b>.9790</b>	.0340	16.84	.8790	.1250	22.03	.9090	.0670
Deformable-GS	<b>20.32</b>	.9162	.0790	<b>29.49</b>	<b>.9804</b>	<b>.0237</b>	30.15	.9172	.0799	21.60	.8876	.0820
GaussianPredict	20.12	.9150	.0811	28.09	.9759	<b>.0322</b>	<b>30.75</b>	<b>.9264</b>	<b>.0767</b>	<b>23.75</b>	<b>.9112</b>	<b>.0553</b>
<b>Ours</b>	<b>22.18</b>	<b>.9243</b>	<b>.0715</b>	24.91	.9660	.0472	<b>31.80</b>	<b>.9365</b>	<b>.0686</b>	<b>28.33</b>	<b>.9493</b>	<b>.0343</b>

Table 2: Quantitative extrapolation results on NVFi dataset scenes. Metrics reported include PSNR, SSIM, and LPIPS. The best metric is highlighted in red, second-best is highlighted in orange.

Method	fan			whale			shark			bat			telescope		
	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)									
TiNeuVox	26.91	9.915	.0643	27.20	.9430	.0579	30.95	.9656	.0367	28.65	.9434	.0663	27.04	.9297	.0507
Deformable-GS	23.75	9.924	.0519	26.58	.9605	.0386	29.11	.9672	.0273	27.07	.9456	.0482	22.92	.9346	.0459
4D-GS	24.78	.9565	.0417	22.31	.9638	.0370	22.56	.9648	.0333	19.83	.9565	.0496	22.77	.9414	.0432
GaussianPredict	<b>30.21</b>	<b>.9682</b>	<b>.0324</b>	25.11	.9610	.0442	29.91	.9695	.0295	22.96	.9587	.0761	21.94	.9381	.0453
NVFi	27.17	.9630	.0370	26.03	.9780	.0290	28.87	<b>.9820</b>	<b>.0210</b>	25.02	.9680	.0420	27.10	.9630	.0460
4D-Rotor-GS	30.16	.9594	.0346	<b>33.45</b>	<b>.9850</b>	<b>.0211</b>	<b>35.08</b>	.9760	.0287	<b>31.76</b>	<b>.9688</b>	<b>.0375</b>	<b>30.00</b>	<b>.9714</b>	<b>.0237</b>
<b>Ours</b>	<b>33.49</b>	<b>.9711</b>	<b>.0303</b>	<b>33.86</b>	<b>.9859</b>	<b>.0135</b>	<b>38.73</b>	<b>.9892</b>	<b>.0082</b>	<b>36.68</b>	<b>.9825</b>	<b>.0176</b>	<b>36.57</b>	<b>.9884</b>	<b>.0057</b>
fallingball				chessboard				dormroom				dining			
Method	PSNR(↑)	SSIM(↑)	LPIPS(↓)	PSNR(↑)	SSIM(↑)	LPIPS(↓)									
	<b>30.00</b>	<b>.9500</b>	<b>.0400</b>	21.76	.7567	.2421	24.01	.7400	<b>.1813</b>	23.56	.8443	<b>.1288</b>	25.36	.8222	<b>.1372</b>
TiNeuVox	24.50	.9200	.0600	20.28	.7866	.2227	22.56	.7423	.2232	20.99	.7922	.2168	23.63	.8107	.1924
Deformable-GS	22.00	.9100	.0700	20.71	.8199	.3444	21.99	.7375	.4036	22.08	.8499	.2800	23.42	.8252	.3356
4D-GS	21.50	.9000	.0800	20.12	.7283	.3168	20.01	.6571	.3840	18.01	.6845	.3811	21.11	.8390	.2708
GaussianPredict	<b>31.37</b>	<b>.9780</b>	<b>.0410</b>	<b>27.84</b>	<b>.8720</b>	<b>.2100</b>	<b>30.41</b>	<b>.8260</b>	.2730	<b>29.01</b>	<b>.8890</b>	.1710	<b>.31.72</b>	<b>.9080</b>	.1540
<b>Ours</b>	22.62	.9244	.0684	<b>33.38</b>	<b>.9266</b>	<b>.0976</b>	<b>34.16</b>	<b>.9019</b>	<b>.1155</b>	<b>30.30</b>	<b>.8829</b>	<b>.1451</b>	<b>.34.53</b>	<b>.9183</b>	<b>.1006</b>

## 4.1 QUANTITATIVE RESULTS

**D-NeRF.** On the D-NeRF dataset (Table 1), ODE-GS consistently surpasses both interpolation-based baselines (Deformable-GS, 4D-GS) and extrapolation-oriented methods (GaussianPredict). Our model averages 27.30 Peak Signal-to-Noise Ratio (PSNR), 0.9497 SSIM, and 0.0467 LPIPS-vgg (Zhang et al., 2018), increasing metric performance against previous SOTA GaussianPrediction by 18.6%. We have especially large margins in Mutant (+10 dB) and Standup (+7 dB), where scene motion are very smooth and follow a straightforward trajectory. These gains directly validate our motivation: interpolation methods degrade when extrapolated beyond the training window due to their timestamp-conditioned design, while our Latent ODE approach enables extrapolation.

**NVFi.** The NVFi benchmark (Table 2) emphasizes robustness under both controlled single-object motion (fan, shark, telescope) and complex multi-object indoor dynamics (factory, darkroom, chessboard). ODE-GS achieves new state of the art across nearly all sequences, averaging 33.43 PSNR, 0.9471 SSIM, and 0.0603 LPIPS, improving previous SOTA method NVFi by 20%, and GaussianPrediction by 39.6% on averaging across scenes and metrics. Notably, in cluttered or occluded settings like factory and darkroom, our model reduces perceptual error (LPIPS) by more than 40% compared to baselines, demonstrating the benefit of decoupling reconstruction from forecasting.

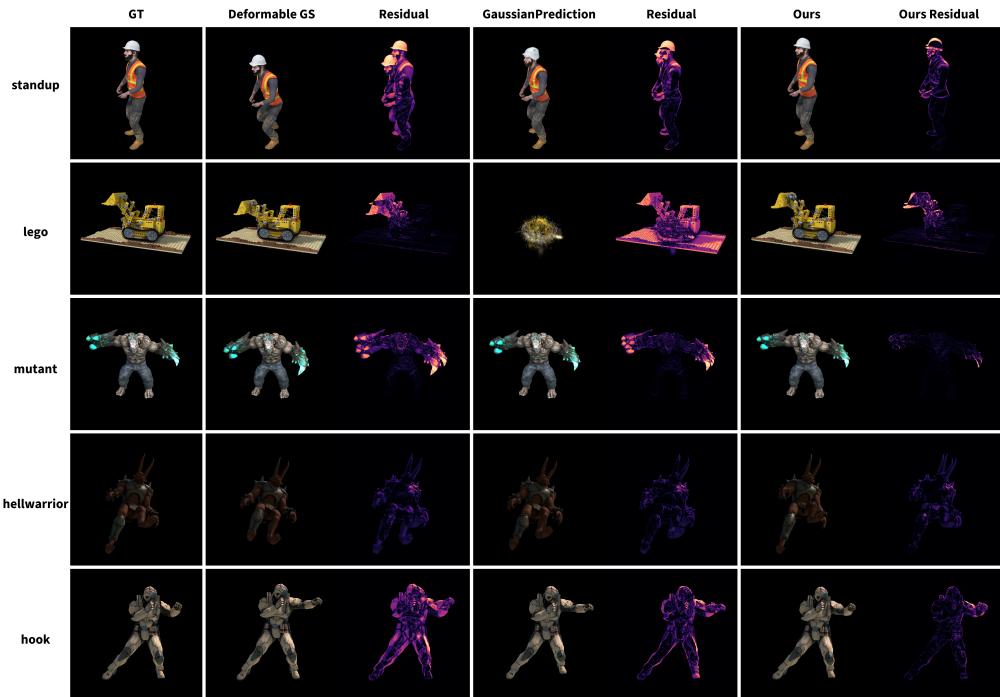
**HyperNeRF.** On real-world HyperNeRF scenes (Table 3), ODE-GS establishes consistent improvements over baselines despite the higher noise and irregular dynamics of captured videos. Our method delivers lower LPIPS in split-cookie, slice-banana, and cut-lemon, and outperforms in both PSNR and SSIM in chickchicken. These results show that ODE-GS avoids overfitting to noisy timestamp signals and remains stable under real-world non-idealities.

## 4.2 QUALITATIVE RESULTS

To provide a more intuitive understanding of our model's performance, we present a qualitative analysis of the extrapolated renderings, shown by Figure 3, which offers a compelling side-by-side

378 Table 3: Quantitative extrapolation results on HyperNeRF. Metrics reported include PSNR, SSIM,  
 379 and LPIPS. The best metric is highlighted in bold, second-best is underlined.  
 380

381 Method	382 split-cookie			383 slice-banana			384 cut-lemon		
	385 PSNR( $\uparrow$ )	386 SSIM( $\uparrow$ )	387 LPIPS( $\downarrow$ )	388 PSNR( $\uparrow$ )	389 SSIM( $\uparrow$ )	390 LPIPS( $\downarrow$ )	391 PSNR( $\uparrow$ )	392 SSIM( $\uparrow$ )	393 LPIPS( $\downarrow$ )
TiNeuVox	16.67	<b>.6135</b>	.4778	18.44	.6242	.6119	18.84	.6228	.5743
Deformable-GS	<b>17.84</b>	.5698	<b>.2945</b>	<b>21.73</b>	<b>.6530</b>	<b>.3241</b>	<b>21.36</b>	<b>.6950</b>	<b>.3207</b>
GaussianPredict	16.93	.5604	.3336	<b>21.97</b>	.6110	.3749	20.91	.6137	.3220
<b>Ours</b>	<b>20.72</b>	<b>.6593</b>	<b>.2406</b>	21.29	<b>.6437</b>	<b>.3230</b>	<b>21.69</b>	<b>.6964</b>	<b>.3098</b>
387 Method	388 keyboard			389 3dprinter			390 chickchicken		
	391 PSNR( $\uparrow$ )	392 SSIM( $\uparrow$ )	393 LPIPS( $\downarrow$ )	394 PSNR( $\uparrow$ )	395 SSIM( $\uparrow$ )	396 LPIPS( $\downarrow$ )	397 PSNR( $\uparrow$ )	398 SSIM( $\uparrow$ )	399 LPIPS( $\downarrow$ )
TiNeuVox	<b>19.03</b>	.6823	.4665	<b>18.16</b>	.5878	.4949	15.23	.6438	.5888
Deformable-GS	19.98	.6957	<b>.2497</b>	19.89	<b>.6775</b>	.2378	19.03	<b>.6996</b>	<b>.3167</b>
GaussianPredict	20.13	<b>.6999</b>	.2511	19.96	.6468	<b>.2209</b>	<b>21.94</b>	.6903	.3284
<b>Ours</b>	<b>21.06</b>	<b>.7399</b>	<b>.2327</b>	<b>20.22</b>	<b>.6946</b>	<b>.2242</b>	<b>20.29</b>	<b>.7254</b>	<b>.3023</b>



416 Figure 3: Qualitative visualization on 5 scenes from DNeRF dataset, from left to right are the ground  
 417 truth image, rendered result from Deformable GS(Yang et al., 2024), residual of Deformable GS  
 418 against GT, GaussianPrediction(Zhao et al., 2024), residual of GaussianPrediction against GT,  
 419 and finally Our as well as Ours residual against GT.

420  
 421 comparison on five scenes from the D-NeRF dataset. While all models are tasked with predicting the  
 422 scene at a future, unseen timestamp, the results vary dramatically. We show the difference in predicted  
 423 motion via residual error maps, which visualize the pixel-wise error between the rendered images  
 424 and the ground truth, where brighter regions indicate high error and dark regions indicate low error.  
 425 The error maps for both baselines show bright, structured residuals concentrated around the central  
 426 object or character, revealing substantial inaccuracies in both shape and position. Conversely, the  
 427 residual map for ODE-GS is significantly darker and less structured, providing clear visual evidence  
 428 of a much lower prediction error. This demonstrates that by learning the underlying dynamics in a  
 429 continuous latent space, ODE-GS not only preserves the high-frequency details but also forecasts  
 430 motion more accurately for photorealistic novel-view synthesis in future frames. In particular, the  
 431 scene Lego is recognized to have inaccurate poses by previous studies (Zhao et al., 2024; Yang et al.,  
 432 2024), but our method can still extrapolate the scene that matches the ground truth image with low

432 error, while methods like Gaussian Prediction (Zhao et al., 2024) can be less robust to such noises  
 433 and fail completely at extrapolating the scene. For additional qualitative on NVFi see Figure 6.  
 434

#### 435 4.2.1 ABLATION STUDY

437 We compare ODE-GS against four different ablation  
 438 settings: Without ODE formulation (pure au-  
 439 toregressive Transformer baseline trained under the  
 440 same input–output supervision), without regulariza-  
 441 tions, without adaptive regularization scaling, and  
 442 without dynamic sampling. The results are sum-  
 443 marized through average metrics in Table 4. In  
 444 general, our full method outperforms all ablation settings.  
 445 For our ablation setting without ODE, we use the  
 446 Transformer encoder and decoder in a pure au-  
 447 toregressive style. Specifically, this model variation has  
 448 the same architecture only without the neural ODE.  
 449 However, the autoregressive Transformer directly pre-  
 450 dicts the next Gaussian parameters in discrete fixed  
 451 steps, as each predicted output is then fed back as  
 452 input for subsequent predictions. As shown in  
 453 Table 4, the autoregressive baseline significantly under-  
 454 performs, almost doubling the LPIPS metric. This  
 455 highlights a limitation of discrete autoregressive mod-  
 456 els also discussed in previous works (Chen et al.,  
 457 2018; Rubanova et al., 2019): they lack the inherent  
 458 smoothness prior naturally expressed by the ODE  
 459 formulation. Therefore, abrupt jumps or oscillations  
 460 in the predicted dynamics can occur. For the setting  
 461 without regularization, we used only the extrapolation  
 462 loss  $L_e$  for training objective. As shown on Table  
 463 4, the additional regularization narrows down solution  
 464 space further, which results in higher average  
 465 metrics. In Figure 4, we observe these additional  
 466 constraints helps the model improve on scenes with  
 467 more complex and diverse motion such as the dining  
 468 and hell-warrior scene. For per-scene results, refer to  
 469 Table 8 in the Appendix.

## 470 5 LIMITATIONS AND CONCLUSION

471 Our model inherits the quality of the underlying interpolation model used to generate Gaussian  
 472 trajectories. If the interpolation stage fails to accurately reconstruct the scene within the observed  
 473 window (for instance, in cases of fast-moving small objects like the *fallingball* scene) then the  
 474 subsequent extrapolation will propagate these errors forward. In scenarios like *bouncingballs* and  
 475 *trex* where the evolution of the scene undergoes abrupt changes, discontinuities, or fundamentally  
 476 novel behaviors not foreshadowed by the past, the model’s predictions degrade. One possible way  
 477 to address this limitation is to utilize data-driven priors for extrapolation that can generalize across  
 478 scenes instead of overfitting to the observed dynamics of specific scenes. We have conducted a  
 479 preliminary experiment on our model’s generalization capability across scenes B.2.

480 In summary, this work introduces ODE-GS, a method that integrates 3D Gaussian Splatting with  
 481 Transformer-based latent neural ODEs to achieve dynamic scene extrapolation. By decoupling scene  
 482 reconstruction from temporal forecasting, enforcing smoothness through continuous-time latent  
 483 dynamics, and incorporating adaptive regularization and dynamic sampling, ODE-GS establishes  
 484 state-of-the-art performance on synthetic (D-NeRF, NVFi) and real-world (HyperNeRF) benchmarks.  
 485 Our results demonstrate that modeling scene evolution as latent flows mitigates out-of-distribution  
 486 failures common in timestamp-conditioned methods and enables accurate extrapolation of scene  
 487 dynamics beyond the observed time window.

Table 4: Ablation study average results over the NVFi dataset.

Method	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
w/o ODE	23.71	.879	.113
w/o Regularization	32.90	.943	.066
w/o Adaptive reg.	32.19	.938	.068
w/o Dynamic sampling	31.35	.935	.069
<b>Ours</b>	<b>33.43</b>	<b>.947</b>	<b>.060</b>

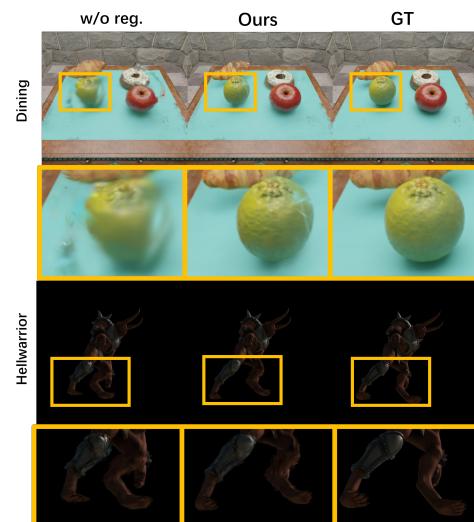
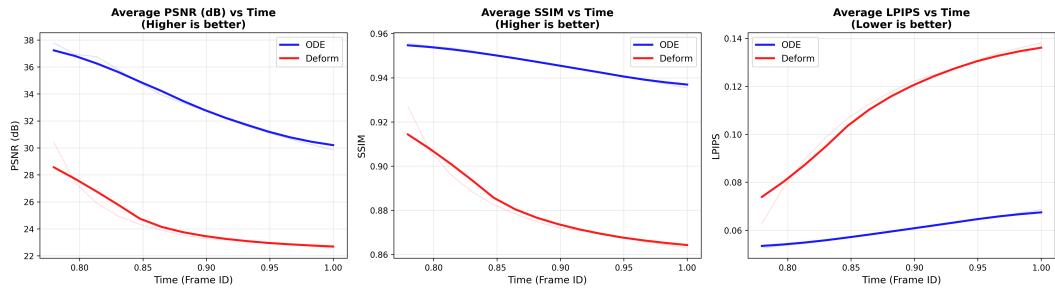


Figure 4: Qualitative comparison on ODE-GS trained using latent and trajectory regularization vs. using only extrapolation loss on two selected scenes. We highlight the areas within each scene with highest visual disparity.

486  
 487 Table 5: Quantitative extrapolation results on the D-NeRF dataset also comparing with 4D-Rotor-  
 488 Gaussians (Duan et al., 2024) as well as using additional projection loss during training. Metrics  
 489 reported include PSNR, SSIM, and LPIPS-vgg. The best metric is highlighted in red, and second  
 490 best is highlighted in orange.

491 Method	492 Lego			493 Mutant			494 Standup			495 Trex		
	496 PSNR( $\uparrow$ )	497 SSIM( $\uparrow$ )	498 LPIPS( $\downarrow$ )	499 PSNR( $\uparrow$ )	500 SSIM( $\uparrow$ )	501 LPIPS( $\downarrow$ )	502 PSNR( $\uparrow$ )	503 SSIM( $\uparrow$ )	504 LPIPS( $\downarrow$ )	505 PSNR( $\uparrow$ )	506 SSIM( $\uparrow$ )	507 LPIPS( $\downarrow$ )
TiNeuVox-B	23.34	.9102	.0942	24.40	.9282	.0700	21.77	.9169	.0927	20.72	.9284	.0751
4D-GS	24.25	.9150	.0810	22.48	.9300	.0520	18.61	.9180	.0840	23.83	.9460	.0510
Deformable-GS	23.25	.9349	.0579	24.45	.9310	.0461	21.37	.9124	.0844	20.74	.9421	.0465
GaussianPredict	12.25	.7594	.2325	27.12	.9514	.0285	26.91	.9456	.0465	21.52	.9443	.0437
4D-Rotor-Gaussians	22.32	.9178	.0705	30.62	.9686	.0252	25.79	.9305	.0632	20.21	.9425	.0692
<b>Ours</b>	<b>25.74</b>	<b>.9378</b>	<b>.0547</b>	<b>34.53</b>	<b>.9804</b>	<b>.0126</b>	<b>28.91</b>	<b>.9557</b>	<b>.0360</b>	<b>22.04</b>	<b>.9475</b>	<b>.0485</b>
<b>Ours_with_projection</b>	<b>25.63</b>	<b>.9368</b>	<b>.0549</b>	29.72	.9671	<b>.0219</b>	<b>29.27</b>	<b>.9546</b>	<b>.0380</b>	22.02	<b>.9469</b>	.0496
498 Method	499 Jumpingjacks			500 Bouncingballs			501 Hellwarrior			502 Hook		
	503 PSNR( $\uparrow$ )	504 SSIM( $\uparrow$ )	505 LPIPS( $\downarrow$ )	506 PSNR( $\uparrow$ )	507 SSIM( $\uparrow$ )	508 LPIPS( $\downarrow$ )	509 PSNR( $\uparrow$ )	510 SSIM( $\uparrow$ )	511 LPIPS( $\downarrow$ )	512 PSNR( $\uparrow$ )	513 SSIM( $\uparrow$ )	514 LPIPS( $\downarrow$ )
TiNeuVox-B	19.87	.9115	.0954	25.92	.9677	.0853	29.36	.9097	.1138	21.05	.8817	.1033
4D-GS	19.95	<b>.9270</b>	.0770	<b>29.55</b>	<b>.9790</b>	.0340	16.84	.8790	.1250	22.03	.9090	.0670
Deformable-GS	<b>20.32</b>	.9162	.0790	<b>29.49</b>	<b>.9804</b>	<b>.0237</b>	30.15	.9172	.0799	21.60	.8876	.0820
GaussianPredict	20.12	.9150	.0811	28.09	.9759	<b>.0322</b>	30.75	.9264	.0767	23.75	.9112	.0553
4D-Rotor-Gaussians	20.93	.9063	.1007	24.05	.9401	.0731	28.97	.9006	.1252	23.56	.9156	.0729
<b>Ours</b>	<b>22.18</b>	<b>.9243</b>	<b>.0715</b>	24.91	.9660	.0472	<b>31.80</b>	<b>.9365</b>	<b>.0686</b>	<b>28.33</b>	<b>.9493</b>	<b>.0343</b>
<b>Ours_with_projection</b>	<b>21.32</b>	.9205	<b>.0758</b>	25.96	.9684	.0436	<b>31.48</b>	<b>.9331</b>	<b>.0704</b>	<b>27.80</b>	<b>.9454</b>	<b>.0368</b>

## 506 6 ADDITIONAL EXPERIMENTS



**Projection loss joint training** To verify the effectiveness of training solely on the generated trajectories of the interpolation model, we have also experimented with adding the image projection loss to the training of the extrapolation model. Specifically, for every iteration, we sample a random camera with ground truth image, and project the gaussians onto the image via the rendering pipeline mentioned in 3. We follow the original 3DGS (Kerbl et al., 2023) and apply L1 loss to the rendered image against the Ground Truth. As shown in table 10, the projection loss does not increase the performance of the extrapolation model trained on pure generated data. This is likely due to the interpolation model were trained using the projection loss and has already accurately represented the training data, thus being able to fit the trajectory for extrapolation is already sufficient as training supervision.

**Performance degrade analysis** Shown in figure 5 as well as per-scene figure 8,7,11,12,14,10,15,13, we analyzed the performance degrade of the rendered results during extrapolation period. On the NVFi dataset, our model not only shows an overall high metric performance, but also a slower degrade rate in time. Especially the SSIM and LPIPS measures degrade much slower than the baseline Yang et al. (2024), showing our method’s capability of maintaining plausible motion and object appearance as extrapolation length becomes longer.

540 ETHICS STATEMENT  
541

542 This work advances methods for forecasting dynamics in 3D scenes using Gaussian splatting and  
543 latent ODEs. Given the scope of the work, we do not identify immediate ethical risks associated  
544 with the approach itself. However, as with any machine learning system, outcomes depend on the  
545 quality and diversity of the training data. If deployed in downstream applications such as robotics  
546 or autonomous navigation, inaccurate forecasts of scene dynamics may pose safety risks, and care  
547 should be taken to validate performance in those contexts.

548  
549 REPRODUCIBILITY STATEMENT  
550

551 We provide detailed descriptions of our approach in Sec. 3, including architecture design, sampling  
552 strategies, and training objectives. Experimental settings, dataset usage, and evaluation metrics are  
553 reported in Sec. B.1, with additional implementation details included in Appendix B. To ensure  
554 reproducibility, we plan to release our source code, pretrained checkpoints, and configuration files.

555 REFERENCES  
556

558 Benjamin Attal, Jia-Bin Huang, Christian Richardt, Michael Zollhoefer, Johannes Kopf, Matthew  
559 O’Toole, and Changil Kim. Hyperreel: High-fidelity 6-dof video with ray-conditioned sampling.  
560 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
561 16610–16620, 2023.

562 Peter W Battaglia, Jessica B Hamrick, and Joshua B Tenenbaum. Simulation as an engine of physical  
563 scene understanding. *Proceedings of the national academy of sciences*, 110(45):18327–18332,  
564 2013.

565 Andrew Bond, Jui-Hsien Wang, Long Mai, Erkut Erdem, and Aykut Erdem. Gaussianvideo: Efficient  
566 video representation via hierarchical gaussian splatting. *arXiv preprint arXiv:2501.04782*, 2025.

567 Michael Broxton, John Flynn, Ryan Overbeck, Daniel Erickson, Peter Hedman, Matthew Duvall,  
568 Jason Dourgarian, Jay Busch, Matt Whalen, and Paul Debevec. Immersive light field video with a  
569 layered mesh representation. *ACM Transactions on Graphics (TOG)*, 39(4):86–1, 2020.

570 Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Proceedings of  
571 the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 130–141, 2023.

572 Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K. Duvenaud. Neural ordinary dif-  
573 ferential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31,  
574 pp. 6571–6583, 2018.

575 Edward De Brouwer, Jaak Simm, Adam Arany, and Yves Moreau. GRU-ODE-Bayes: Continuous  
576 modeling of sporadically-observed time series. In *Advances in Neural Information Processing  
577 Systems (NeurIPS)*, volume 32, pp. 7366–7376, 2019.

578 Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh  
579 Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d:  
580 Real-time performance capture of challenging scenes. *ACM Transactions on Graphics (ToG)*, 35  
(4):1–13, 2016.

581 Yilun Du, Yinan Zhang, Hong-Xing Yu, Joshua B Tenenbaum, and Jiajun Wu. Neural radiance  
582 flow for 4d view synthesis and video processing. In *2021 IEEE/CVF International Conference on  
583 Computer Vision (ICCV)*, pp. 14304–14314. IEEE Computer Society, 2021.

584 Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4d-rotor  
585 gaussian splatting: towards efficient novel view synthesis for dynamic scenes. In *ACM SIGGRAPH  
586 2024 Conference Papers*, pp. 1–11, 2024.

587 Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner,  
588 and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia  
589 2022 Conference Papers*, pp. 1–9, 2022.

594 Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic  
 595 monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*,  
 596 pp. 5712–5721, 2021.

597

598 Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs:  
 599 Sparse-controlled gaussian splatting for editable dynamic scenes. In *Proceedings of the IEEE/CVF*  
 600 *conference on computer vision and pattern recognition*, pp. 4220–4230, 2024.

601 Kacper Kania, Kwang Moo Yi, Marek Kowalski, Tomasz Trzcinski, and Andrea Tagliasacchi. Conerf:  
 602 Controllable neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer*  
 603 *Vision and Pattern Recognition*, pp. 18623–18632, 2022.

604

605 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuhler, and George Drettakis. 3d gaussian splatting  
 606 for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

607

608 Mina A Khoei, Guillaume S Masson, and Laurent U Perrinet. The flash-lag effect as a motion-based  
 609 predictive shift. *PLoS computational biology*, 13(1):e1005068, 2017.

610

611 Patrick Kidger, Ricky T. Q. Chen, and Miles Cranmer. torchode: A parallel ode solver for pytorch.  
 612 <https://github.com/patrick-kidger/torchode>, 2021.

613

614 Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013.  
 615 URL <https://api.semanticscholar.org/CorpusID:216078090>.

616

617 Jinxi Li, Ziyang Song, and Bo Yang. Nvfi: Neural velocity fields for 3d physics learning from  
 618 dynamic videos. *Advances in Neural Information Processing Systems*, 36:34723–34751, 2023.

619

620 Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time  
 621 dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*  
 622 *Pattern Recognition*, pp. 8508–8520, 2024.

623

624 Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time  
 625 view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer*  
 626 *Vision and Pattern Recognition*, pp. 6498–6508, 2021.

627

628 Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser  
 629 Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint*  
 630 *arXiv:1906.07751*, 2019.

631

632 Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians:  
 633 Tracking by persistent dynamic view synthesis. In *2024 International Conference on 3D Vision*  
 634 (*3DV*), pp. 800–809. IEEE, 2024.

635

636 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and  
 637 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*  
 638 *of the ACM*, 65(1):99–106, 2021.

639

640 Leigh A Mrotek and John F Soechting. Predicting curvilinear target motion through an occlusion.  
 641 *Experimental Brain Research*, 178(1):99–114, 2007.

642

643 Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking  
 644 of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and*  
 645 *pattern recognition*, pp. 343–352, 2015.

646

647 Sergio Orts-Escalano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yury  
 648 Degtyarev, David Kim, Philip L Davidson, Sameh Khamis, Mingsong Dou, et al. Holoportion:  
 649 Virtual 3d teleportation in real-time. In *Proceedings of the 29th annual symposium on user interface*  
 650 *software and technology*, pp. 741–754, 2016.

651

652 Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for  
 653 dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 654 *Recognition*, pp. 2856–2865, 2021.

648 Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz,  
 649 and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the*  
 650 *IEEE/CVF international conference on computer vision*, pp. 5865–5874, 2021a.

651

652 Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman,  
 653 Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for  
 654 topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021b.

655 Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun  
 656 Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *Proceedings of*  
 657 *the IEEE/CVF International Conference on Computer Vision*, pp. 14314–14323, 2021a.

658

659 Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei  
 660 Zhou. Neural body: Implicit neural representations with structured latent codes for novel view  
 661 synthesis of dynamic humans. In *Proceedings of the IEEE/CVF conference on computer vision*  
 662 and pattern recognition, pp. 9054–9063, 2021b.

663 Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural  
 664 radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF conference on computer*  
 665 *vision and pattern recognition*, pp. 10318–10327, 2021.

666 Rajesh PN Rao and Dana H Ballard. Predictive coding in the visual cortex: a functional interpretation  
 667 of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.

668

669 Yulia Rubanova, Ricky T. Q. Chen, and David K. Duvenaud. Latent ordinary differential equations for  
 670 irregularly-sampled time series. In *Advances in Neural Information Processing Systems (NeurIPS)*,  
 671 volume 32, pp. 5321–5331, 2019.

672 Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and  
 673 Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a  
 674 dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference*  
 675 *on Computer Vision*, pp. 12959–12970, 2021.

676

677 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz  
 678 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*  
 679 *systems*, 30, 2017.

680 Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian,  
 681 and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings*  
 682 *of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 20310–20320, 2024.

683 Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable  
 684 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the*  
 685 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 20331–20341, 2024.

686

687 Caglar Yildiz, Markus Heimonen, and Harri Lähdesmäki. ODE2VAE: Deep generative second order  
 688 ODEs with bayesian neural networks. In *Advances in Neural Information Processing Systems*  
 689 (*NeurIPS*), volume 32, pp. 10280–10290, 2019.

690 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable  
 691 effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.

692

693 Boming Zhao, Yuan Li, Ziyu Sun, Lin Zeng, Yujun Shen, Rui Ma, Yinda Zhang, Hujun Bao, and  
 694 Zhaopeng Cui. Gaussianprediction: Dynamic 3d gaussian prediction for motion extrapolation and  
 695 free view synthesis. In *ACM SIGGRAPH 2024 Conference Papers*, pp. 1–12, 2024.

696

697

698

699

700

701

702 **A APPENDIX**  
 703

704 **A.0.1 DETAILS ON ADAPTIVE WEIGHTING FOR REGULARIZATION**  
 705

706 This mechanism operates at each training iteration by using the Exponential Moving Average (EMA)  
 707 of the trajectory prediction loss as a proxy for the model’s convergence state. The regularization  
 708 weight increased as the extrapolation loss decreases, allowing the regularizers to more strongly guide  
 709 the final, fine-grained trajectory predictions.

710 The scaling factor  $s_t$  at iteration  $t$  is computed as:  
 711

$$712 s_t = \exp \left( -\frac{1}{\tau} \cdot \text{clip} \left( \frac{\mathcal{L}_{\text{EMA}}(t) - \mathcal{L}_{\text{end}}}{\mathcal{L}_{\text{init}} - \mathcal{L}_{\text{end}}}, 0, 1 \right) \right) \quad (10)$$

714 where:  
 715

- 716 •  $\mathcal{L}_{\text{EMA}}(t)$  is the EMA of the trajectory prediction loss at iteration  $t$ .  
 717
- 718 •  $\mathcal{L}_{\text{init}}$  and  $\mathcal{L}_{\text{end}}$  are hyperparameters representing the expected initial and final extrapolation  
 719 loss values.
- 720 •  $\tau$  is a temperature hyperparameter that controls the decay rate of the scaling. A lower  $\tau$   
 721 leads to a faster decay and a more aggressive increase in regularization.
- 722 • The ‘clip’ function normalizes the loss into the range  $[0, 1]$ , ensuring the scaling factor  
 723 remains between  $(0, 1]$ .

725 The EMA of the trajectory prediction loss,  $\mathcal{L}_{\text{EMA}}(t)$ , is calculated as:  
 726

$$727 \mathcal{L}_{\text{EMA}}(t) = \alpha \cdot \mathcal{L}_{\text{EMA}}(t-1) + (1 - \alpha) \cdot \mathcal{L}_{\text{e}}(t) \quad (11)$$

728 where  $\mathcal{L}_{\text{e}}(t)$  is the extrapolation loss at the current iteration and  $\alpha$  is the EMA decay rate. The final  
 729 regularization weights for the latent and trajectory regularizers are then scaled by  $s_t$  at each iteration  
 730 before being added to the total loss.

731 **A.0.2 PROBABILISTIC FORECASTING WITH A VARIATIONAL LATENT ODE**  
 732

733 Many past Latent ODE works that focus on extrapolation and forecasting has been using a variational  
 734 formulation for the model (Rubanova et al., 2019). Specifically, they follow the Variational Autoencoder  
 735 (Kingma & Welling, 2013) approach. We provide details on our model variant which uses this  
 736 approach instead of deterministic modeling, as well as quantitative results compare the two methods  
 737 on the NVFi dataset, as shown in table 6

738 **Variational Variant for the Transformer ODE architecture.** Instead of mapping the observed  
 739 history to a single initial state for the ODE, prior works often formulate trajectory forecasting  
 740 as a probabilistic problem (Chen et al., 2018). Given a historical trajectory  $\gamma_k = \{G_k(t_j)\}_{j=1}^{T_c}$   
 741 for a Gaussian  $k$ , the Transformer encoder produces a latent vector  $h_k$ . A projection head then  
 742 parameterizes a Gaussian posterior distribution over the initial latent state,  
 743

$$744 q_{\phi}(z_k(t_0) | \gamma_k) = \mathcal{N}(\mu_{z_k}, \text{diag}(\sigma_{z_k}^2)), \quad (12)$$

746 where  $z_k(t_0)$  is the latent state at the start of extrapolation. During training, we sample from this  
 747 distribution via the reparameterization trick:

$$748 z_k(t_0) = \mu_{z_k} + \sigma_{z_k} \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I). \quad (13)$$

750 The sampled latent state  $z_k(t_0)$  is then evolved forward in time by numerically solving the latent  
 751 ODE,  
 752

$$753 z_k(t) = \text{ODESolve}(f_{\theta}, z_k(t_0), t), \quad (14)$$

754 and subsequently decoded to predicted Gaussian parameters  
 755

$$\hat{G}_k(t) = \delta_{\psi}(z_k(t)). \quad (15)$$

756 Table 6: Comparison between our deterministic and variational formulations on NVFi dataset scenes.  
 757 Metrics reported include PSNR, SSIM, and LPIPS.  
 758

759 Scene	760 Ours (Deterministic)			761 Ours (Variational)		
	762 PSNR( $\uparrow$ )	763 SSIM( $\uparrow$ )	764 LPIPS( $\downarrow$ )	765 PSNR( $\uparrow$ )	766 SSIM( $\uparrow$ )	767 LPIPS( $\downarrow$ )
768 factory	769 34.53	770 .9183	771 .1006	772 22.57	773 .7851	774 .218
775 dining	776 30.30	777 .8829	778 .1451	779 16.19	780 .3981	781 .595
782 darkroom	783 34.16	784 .9019	785 .1155	786 22.83	787 .7098	788 .233
789 whale	790 33.86	791 .9859	792 .0135	793 26.49	794 .9601	795 .038
797 shark	798 38.73	799 .9892	800 .0082	801 29.43	802 .9652	803 .030
805 chessboard	806 33.38	807 .9266	808 .0976	809 19.69	810 .7667	811 .244
814 bat	815 36.68	816 .9825	817 .0176	818 28.01	819 .9554	820 .043
823 fan	824 33.49	825 .9711	826 .0303	827 21.67	828 .8589	829 .114
833 telescope	834 36.57	835 .9884	836 .0057	837 22.66	838 .9283	839 .053
844 fallingball	845 22.62	846 .9244	847 .0684	848 9.45	849 .7559	850 .206

772  
 773 **Objective.** The variational model is trained by maximizing the Evidence Lower Bound (ELBO),  
 774 which corresponds to minimizing

$$775 \mathcal{L}_{\text{var-e}} = \sum_{t \in \mathcal{T}_e} \underbrace{[-\log p_\sigma(G_k(t) | \hat{G}_k(t))]}_{\text{prediction NLL}} + \text{KL}\left[q_\phi(z_k(t_0) | \gamma_k) \parallel p(z_k(t_0))\right] \quad (16)$$

776 where  $\mathcal{T}_e$  denotes the extrapolation timestamps. The first term encourages predicted Gaussian  
 777 trajectories  $\hat{G}_k(t)$  to align with ground truth  $G_k(t)$ , the second regularizes the latent space toward a  
 778 unit Gaussian prior  $p(z_k(t_0))$ . The final objective is then to compose this loss with regularization as  
 779 discussed in 3.4.

780 While variational formulations provide a principled way to capture uncertainty in trajectory fore-  
 781 casting, our results 6 show that the deterministic version of ODE-GS significantly outperforms its  
 782 variational counterpart across all NVFi scenes. This gap arises because variational training introduces  
 783 additional complexity through posterior sampling and KL regularization, which can destabilize opti-  
 784 mization when data is limited and the ground-truth dynamics are relatively deterministic. In practice,  
 785 the model tends to underfit sharp motion patterns, producing blurred or averaged predictions that  
 786 reduce both PSNR and SSIM while inflating perceptual error (LPIPS). By contrast, the deterministic  
 787 formulation directly learns smooth latent flows aligned with observed trajectories, avoiding posterior  
 788 collapse and better exploiting the strong spatio-temporal smoothness priors inherent in dynamic 3D  
 789 scenes.

## 790 B DETAILS ON DYNAMIC TRAJECTORY SAMPLING

791 We provide additional information on our sampling strategy introduced in Section 3.3. Training data  
 792 for the extrapolation module are derived from the pre-trained interpolation model, which can generate  
 793 Gaussian trajectories at arbitrary real-valued timestamps through the deformation function  $\mathcal{D}_\omega$ .

794 Each training sample is indexed by a Gaussian  $k$  and a starting time  $t_0$ . We fix the number of context  
 795 steps  $N_c$ , target steps  $N_e$ , and the context span  $T_c$ .

796 The context sequence is uniformly sampled as

$$802 \Delta_c = \frac{T_c}{N_c-1}, \quad \gamma_c^{(i)} = \{G_k(t_0 + i\Delta_c)\}_{i=1}^{N_c}. \quad (17)$$

803 The target sequence begins immediately after the context window at  $t_c^{\text{end}} = t_0 + T_c$  and spans the  
 804 remaining horizon:

$$806 T_e = t_{\max} - t_c^{\text{end}}, \quad \Delta_e = \frac{T_e}{N_e}, \quad \gamma_e^{(i)} = \{G_k(t_c^{\text{end}} + i\Delta_e)\}_{i=1}^{N_e}. \quad (18)$$

807 This ensures  $N_c$  and  $N_e$  remain fixed across all samples, while  $T_e$  varies with  $t_0$ . The variability  
 808 in target length naturally generates prediction tasks of differing difficulty, ranging from short- to  
 809 long-term forecasts.

810  
811 Table 7: Per-scene ablation over NVFi of our full method against removal of ODE and removal of  
812 dynamic trajectory sampling. Metrics reported include PSNR, SSIM, and LPIPS.

Scene	w/o ODE			w/o dynamic trajectory sampling			Ours		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
factory	24.06	.814	.179	33.84	.915	.102	34.53	.918	.101
dining	20.98	.784	.217	25.47	.810	.204	30.30	.883	.145
darkroom	22.08	.731	.227	33.89	.900	.117	34.16	.902	.116
chessboard	19.84	.784	.225	32.27	.920	.100	33.38	.927	.098
bat	28.42	.960	.039	35.37	.979	.023	36.68	.983	.018
telescope	23.20	.937	.040	33.09	.983	.010	36.57	.988	.006
fallingball	17.14	.904	.094	26.22	.936	.064	22.62	.924	.068
fan	23.35	.931	.052	26.97	.948	.043	33.49	.971	.030
shark	29.47	.967	.026	33.17	.980	.015	38.73	.989	.008
whale	28.51	.973	.028	33.21	.984	.015	33.86	.986	.014
<b>Average</b>	<b>23.71</b>	<b>.879</b>	<b>.113</b>	<b>31.35</b>	<b>.935</b>	<b>.069</b>	<b>33.43</b>	<b>.947</b>	<b>.060</b>

824  
825 The complete dataset is defined as the union over all Gaussian indices and valid starting times:

$$\mathcal{D} = \bigcup_{k=1}^K \bigcup_{t_0 \in \mathcal{T}_k} \{ (\gamma_c^{(i)}, \gamma_e^{(i)}) \}. \quad (19)$$

830  
831 During training, the context  $\gamma_c^{(i)}$  is processed by the Transformer encoder–ODE module to predict  
832  $\hat{\gamma}_e^{(i)} = \{\hat{G}_k(t)\}_{t \in T_e}$ . The prediction is supervised with an  $L_1$  extrapolation loss:

$$\mathcal{L}_e = \frac{1}{N_e} \sum_{j=1}^{N_e} \|\hat{G}_k(t_j) - G_k(t_j)\|_1. \quad (20)$$

833  
834 At inference, we simply take the final context segment of length  $T_c$  from the observed window and  
835 apply the same extrapolation procedure.

## 836 B.1 IMPLEMENTATION DETAILS

837  
838 **Model architecture.** The Transformer encoder has  $d_{\text{model}} = 128$  latent dimensions with  $n_{\text{head}} = 8$   
839 attention heads and  $L_{\text{enc}} = 5$  encoder layers. For the latent ODE component, we set the latent  
840 dimension to  $d_{\text{latent}} = 64$  and use an MLP with  $L_{\text{ode}} = 4$  layers and  $d_{\text{hidden}} = 64$  hidden units per  
841 layer. The decoder network mirrors this structure with  $L_{\text{dec}} = 5$  layers and  $d_{\text{hidden}} = 128$  hidden units.  
842 For the ODE function, we use Tanh activations for smooth and bounded outputs. For the interpolation  
843 model, we follow Deformable GS (Yang et al., 2024) implementation.

844  
845 **Hyperparameters.** Our two-stage training pipeline is implemented in PyTorch with the original 3D  
846 Gaussian Splatting renderer (Kerbl et al., 2023). For the interpolation stage, we set the learning rate  
847 to  $8 \times 10^{-4}$  with a cosine annealing scheduler that decays to a minimum of  $1.6 \times 10^{-6}$ , training for  
848 40k iterations on images with timestamps prior to our dataset-specific train/test split (Sec. B.2.1).

849  
850 For the extrapolation stage, we solve latent ODEs using the `torchode` package (Kidger et al., 2021)  
851 with the adaptive DOPRI5 solver configured with tolerances  $rtol = 10^{-3}$  and  $atol = 10^{-4}$ . We train  
852 on NVIDIA GPUs (RTX 3090 or A6000) with a batch size of 512 for 40 epochs, using an initial  
853 learning rate of  $1 \times 10^{-3}$  and cosine annealing down to  $1 \times 10^{-6}$ . The input context sequence length  
854 is set to  $T_c = 30$  and the extrapolation length to  $T_e = 10$  during optimization. For adaptive trajectory  
855 sampling, the temperature parameter is set to 0.05 on D-NeRF and 0.15 on HyperNeRF and NVFi.  
856 For adaptive regularization, we use  $\mathcal{L}_{\text{init}} = 0.02$  for D-NeRF and NVFi and 0.05 for HyperNeRF, we  
857 set  $\mathcal{L}_{\text{end}} = 0.0$  over training on all datasets. Our EMA decay rate is set at 0.9. We use  $\lambda_{\text{traj}} = 10^{-1}$   
858 for the trajectory regularizer and  $\lambda_{\text{latent}} = 10^{-5}$  for the latent regularizer.

## 860 B.2 MULTI-SCENE GENERALIZATION EXPERIMENT

861  
862 We evaluate our extrapolation model’s capability to generalize across scenes in Table 9. Specifically,  
863 we set up a multi-scene training experiment where we collect all generated trajectories from the

864  
865  
866  
867 Table 8: Per-scene ablation over NVFi of our full method against removal of regularizers and removal  
868 of adaptive regularizer scaling. Metrics reported include PSNR, SSIM, and LPIPS.

Scene	w/o regularizers			w/o adaptive regularizer scaling			Ours		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
factory	34.68	.919	.102	33.53	.912	.104	34.53	.918	.101
dining	27.26	.841	.180	25.99	.828	.187	30.30	.883	.145
darkroom	33.81	.896	.123	32.32	.880	.127	34.16	.902	.116
chessboard	33.16	.929	.096	32.16	.921	.099	33.38	.927	.098
bat	37.43	.985	.017	36.34	.983	.018	36.68	.983	.018
telescope	36.23	.988	.006	35.68	.988	.006	36.57	.988	.006
fallingball	19.35	.918	.086	19.53	.918	.085	22.62	.924	.068
fan	34.96	.976	.027	35.23	.974	.028	33.49	.971	.030
shark	38.26	.989	.009	36.63	.986	.010	38.73	.989	.008
whale	33.88	.987	.013	34.52	.987	.013	33.86	.986	.014
<b>Average</b>	<b>32.90</b>	<b>.943</b>	<b>.066</b>	<b>32.19</b>	<b>.938</b>	<b>.068</b>	<b>33.43</b>	<b>.947</b>	<b>.060</b>

879  
880 respective interpolation models from each scene in the NVFi dataset, union them into the same dataset,  
881 and train one extrapolation model (The Transformer Latent ODE) on all trajectories simultaneously.  
882 This is equivalent to having the scenes share the same set of extrapolation model weights during  
883 training. To assess the generalization of the model, we intentionally hold out the *whale* scene’s  
884 observed trajectories to be excluded from this dataset union. The last rows (Ours Multi) of Table  
885 9 shows our quantitative results of extrapolation for each scene using this shared-weight model  
886 for extrapolation, conditioned on the last segment of Gaussian trajectories for each scene’s unique  
887 interpolation model. Our quantitative numbers show that our extrapolation model can fit to multiple  
888 scene trajectories at the same time, while being capable of generalizing to unseen scene dynamics with  
889 simple motion patterns like *whale*. Although simultaneously training on multiple scene trajectories  
890 does not make the extrapolation model exceed metric performs compared to Our default per-scene  
891 training, it shows the potential of such and approach for future works to explore.  
892

### B.2.1 DATASETS DETAILS

893  
894 We evaluate ODE-GS on three datasets: the D-NeRF dataset (Pumarola et al., 2021), the NVFi dataset  
895 (Li et al., 2023), and the HyperNeRF dataset (Park et al., 2021b). The D-NeRF dataset comprises eight  
896 synthetic dynamic scenes (Lego, Mutant, Standup, Trex, Jumpingjacks, Bouncingballs, Hellwarrior,  
897 Hook), each containing 100-200 training images and 20 test images with timestamps normalized  
898 from 0 to 1, rendered at 800×800 resolution with black backgrounds. The NVFi dataset is also a  
899 synthetic dataset that provides two subcategories: the simpler Dynamic Object Dataset featuring  
900 rotating objects (fan, whale, shark, bat, telescope) and the more challenging Dynamic Indoor Scene  
901 Dataset (chessboard, darkroom, dining, factory) containing multi-object scenes with occlusions and  
902 realistic lighting variations. The HyperNeRF dataset is a series of monocular videos on day-to-day  
903 scenes with varying motion complexity. We use the split-cookie, slice-banan, cut-lemon, keyboard,  
904 3dprinter, and chickchicken scenes, picked at random. For the D-Nerf dataset, we use 80% of the  
905 temporal sequence for training and reserve the final 20% as ground truth for extrapolation evaluation.  
906 For the HyperNeRF dataset, we use the first 90% of the temporal sequence. For the NVFi dataset, we  
907 follow the original train-test split where 75% of the temporal sequence is used for training and the  
908 other 25% is reserved for testing. Our model’s input time span  $T_c$  are set differently for each dataset,  
909 with  $T_c = 0.6$  for D-NeRF,  $T_c = 0.5$  for NVFi and HyperNeRF.  
910  
911  
912  
913  
914  
915  
916  
917

918  
919  
920  
921  
922  
923Table 9: Quantitative extrapolation results on NVFi dataset scenes for training the extrapolation model on multi-scene setting. Scene *whale* is not included in the training data. Metrics reported include PSNR, SSIM, and LPIPS. The best metric is highlighted in red, second-best is highlighted in orange.

Method	fan			dining			factory		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
TiNeuVox	26.91	.9315	.0643	23.56	.8443	<b>.1288</b>	25.36	.8222	.1372
Deformable-GS	23.75	.9274	.0519	20.99	.7922	.2168	23.63	.8107	.1924
4D-GS	24.78	.9565	.0417	22.08	.8499	.2800	23.42	.8252	.3356
GaussianPredict	<b>30.21</b>	<b>.9682</b>	<b>.0324</b>	18.01	.6845	.3811	21.11	.8390	.2708
NVFi	27.17	.9630	.0370	<b>29.01</b>	<b>.8980</b>	.1710	31.72	<b>.9080</b>	.1540
Ours Det.	<b>33.49</b>	<b>.9711</b>	<b>.0303</b>	<b>30.30</b>	<b>.8829</b>	<b>.1451</b>	<b>34.53</b>	<b>.9183</b>	<b>.1006</b>
Ours Multi.	27.30	.9494	.0422	27.61	.8530	.1591	<b>31.84</b>	.8956	<b>.1110</b>
Method	shark			bat			telescope		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
TiNeuVox	30.95	.9656	.0367	28.65	.9434	.0663	27.04	.9297	.0507
Deformable-GS	29.11	.9672	.0273	27.07	.9456	.0482	22.92	.9346	.0459
4D-GS	22.56	.9648	.0333	19.83	.9565	.0496	22.77	.9414	.0432
GaussianPredict	29.91	.9695	.0295	22.96	.9587	.0761	21.94	.9381	.0453
NVFi	28.87	.9820	.0210	25.02	.9680	.0420	27.10	.9630	.0460
Ours	<b>38.73</b>	<b>.9892</b>	<b>.0082</b>	<b>36.68</b>	<b>.9825</b>	<b>.0176</b>	<b>36.57</b>	<b>.9884</b>	<b>.0057</b>
Ours Multi.	<b>37.31</b>	<b>.9869</b>	<b>.0094</b>	<b>36.72</b>	<b>.9818</b>	<b>.0188</b>	<b>34.19</b>	<b>.9838</b>	<b>.0085</b>
Method	fallingball			chessboard			darkroom		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
TiNeuVox	<b>30.00</b>	<b>.9500</b>	<b>.0400</b>	21.76	.7567	.2421	24.01	.7400	.1813
Deformable-GS	24.50	.9200	.0600	20.28	.7866	.2227	22.56	.7423	.2232
4D-GS	22.00	.9100	.0700	20.71	.8199	.3444	21.99	.7375	.4036
GaussianPredict	21.50	.9000	.0800	20.12	.7283	.3168	20.01	.6371	.3840
NVFi	<b>31.37</b>	<b>.9780</b>	<b>.0410</b>	27.84	.8720	.2100	<b>30.41</b>	.8260	.2730
Ours	22.62	.9244	.0684	<b>33.38</b>	<b>.9266</b>	<b>.0976</b>	<b>34.16</b>	<b>.9019</b>	<b>.1155</b>
Ours Multi.	23.54	.9219	.0779	<b>29.51</b>	<b>.8949</b>	<b>.1116</b>	29.92	<b>.8489</b>	<b>.1364</b>
Method	whale			average					
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )			
TiNeuVox	27.20	.9430	.0579	26.54	.8826	.1005			
Deformable-GS	26.58	.9605	.0386	24.14	.8787	.1127			
4D-GS	22.31	.9638	.0370	22.25	.8926	.1638			
GaussianPredict	25.11	.9610	.0442	23.09	.8584	.1660			
NVFi	26.03	<b>.9780</b>	.0290	28.45	<b>.9336</b>	.1024			
Ours	<b>33.86</b>	<b>.9859</b>	<b>.0135</b>	<b>33.43</b>	<b>.9471</b>	<b>.0603</b>			
Ours Multi.	<b>29.22</b>	.9693	<b>.0260</b>	<b>30.72</b>	.9285	<b>.0701</b>			

958  
959  
960Table 10: Quantitative extrapolation results on the **bouncingballs** scene of D-NeRF dataset. Comparison between Deformable-GS and Ours across different data splits. The best metric is highlighted in red. For Ours, we set the input time span to be a shorten span of 0.1 instead of the default 0.6.

Method	Split 65%			Split 70%			Split 75%			Split 80%		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
Deformable-GS	24.11	.9614	.0533	23.30	.9600	.0550	24.08	.9642	.0478	<b>.29.49</b>	<b>.9804</b>	<b>.0237</b>
Ours	<b>26.55</b>	<b>.9690</b>	<b>.0130</b>	<b>26.27</b>	<b>.9679</b>	<b>.0158</b>	<b>26.84</b>	<b>.9715</b>	<b>.0146</b>	28.02	.9761	.0357
Method	Split 85%			Split 90%			Split 95%			Average		
	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
Deformable-GS	31.44	.9842	.0197	34.50	.9886	.0148	34.31	.9889	.0131	28.89	.9772	.0325
Ours	<b>32.97</b>	<b>.9860</b>	<b>.0054</b>	<b>36.23</b>	<b>.9913</b>	<b>.0039</b>	<b>39.18</b>	<b>.9944</b>	<b>.0035</b>	<b>31.00</b>	<b>.9796</b>	<b>.0138</b>

971

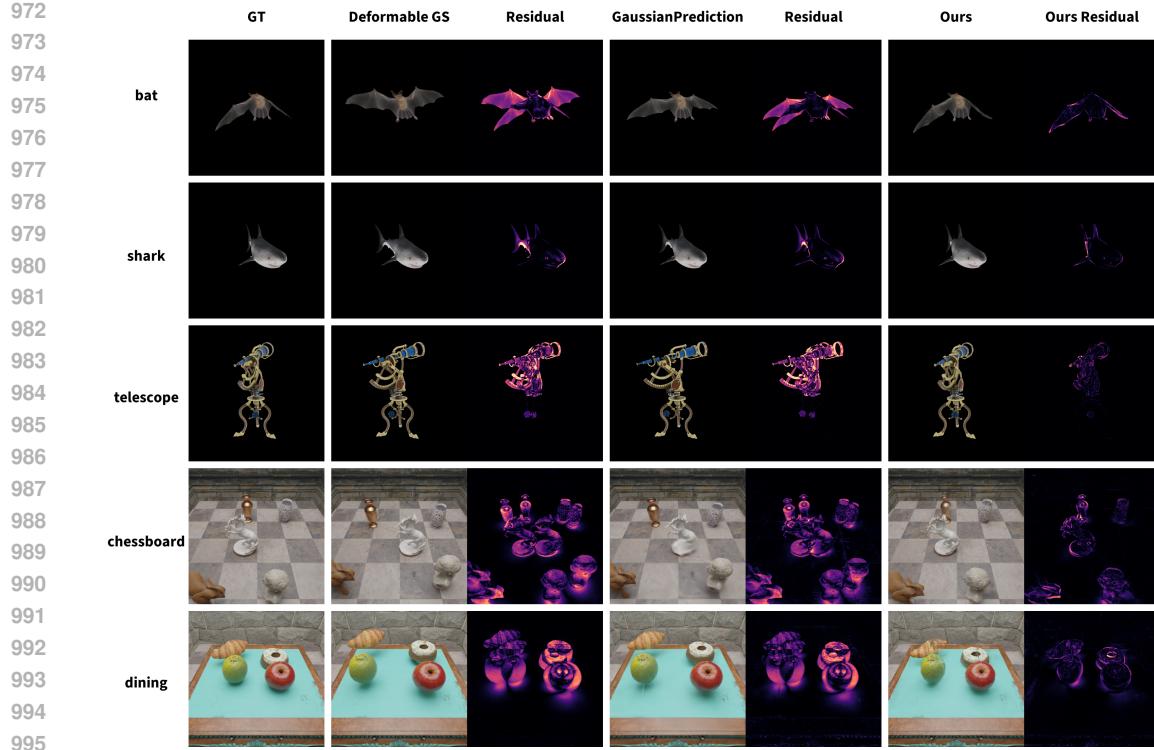


Figure 6: Qualitative results on 5 scenes from the NVFI (Li et al., 2023) dataset, from left to right are the ground truth image, rendered result from Deformable GS(Yang et al., 2024), residual of Deformable GS against GT, GaussianPrediction(Zhao et al., 2024), residual of GaussianPrediction against GT, and finally Our as well as Ours residual against GT.

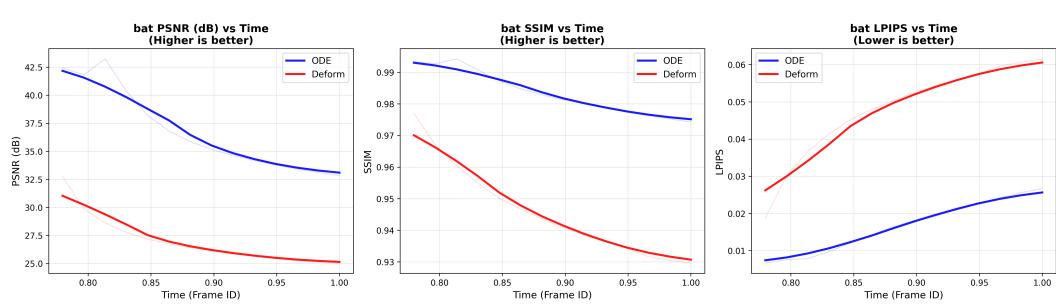


Figure 7: Bat Metrics Comparison

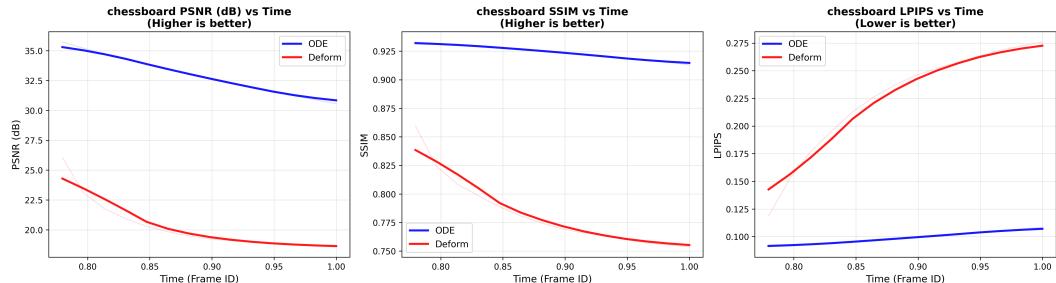


Figure 8: Chessboard Metrics Comparison

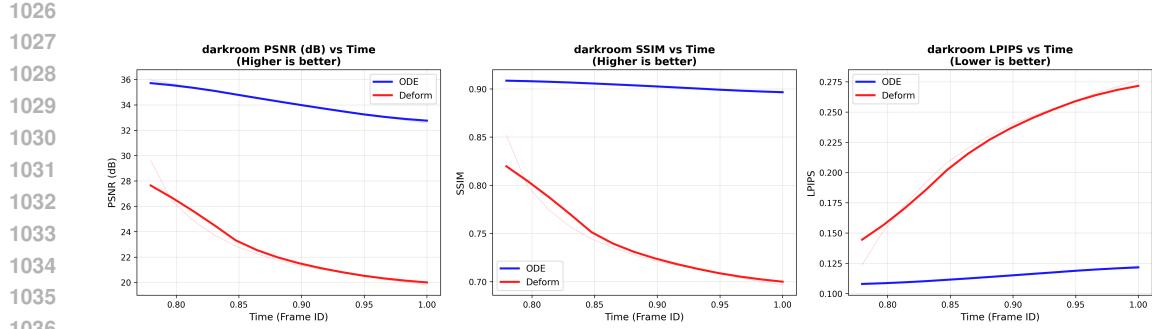


Figure 9: Darkroom Metrics Comparison

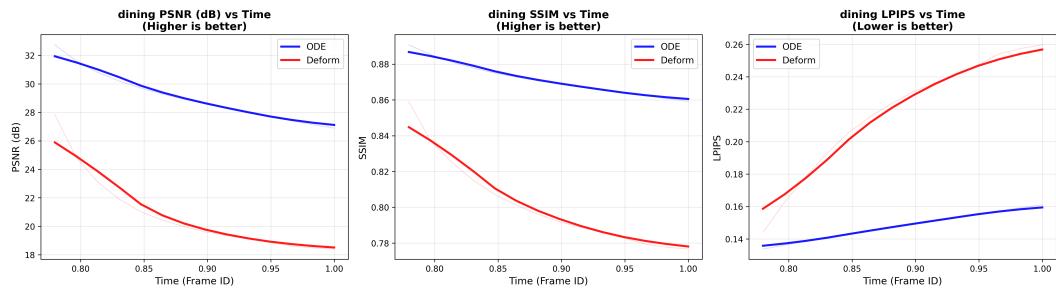


Figure 10: Dining Metrics Comparison

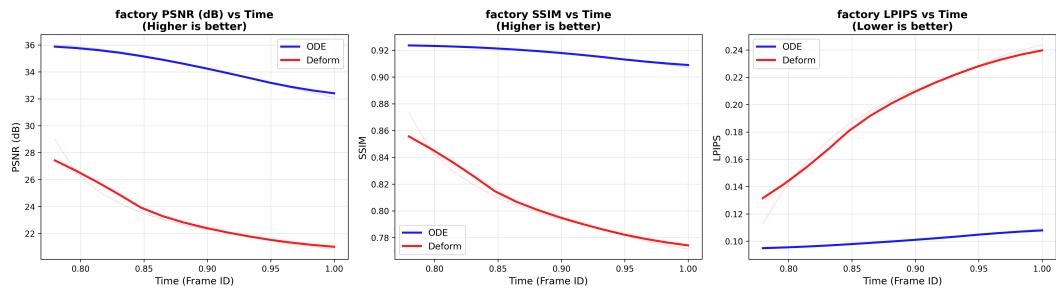


Figure 11: Factory Metrics Comparison

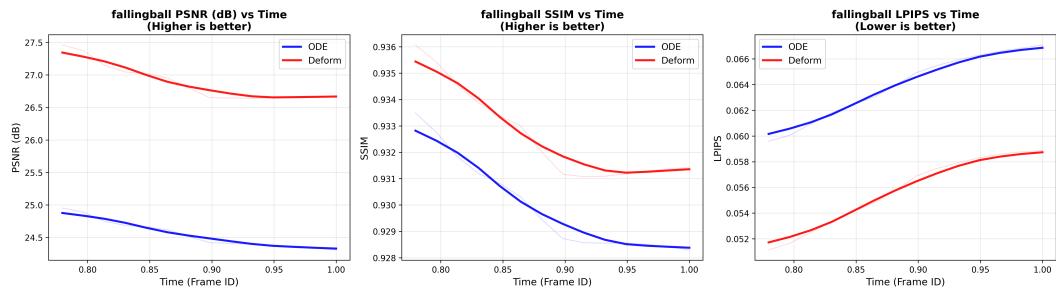


Figure 12: Falling Ball Metrics Comparison

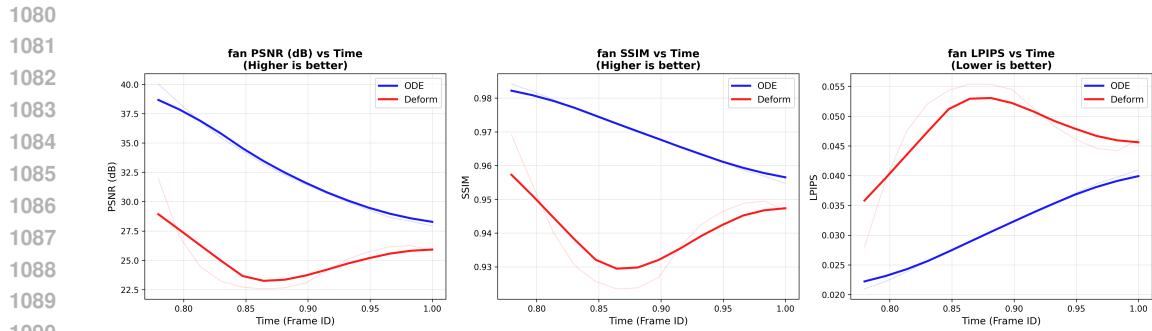


Figure 13: Fan Metrics Comparison

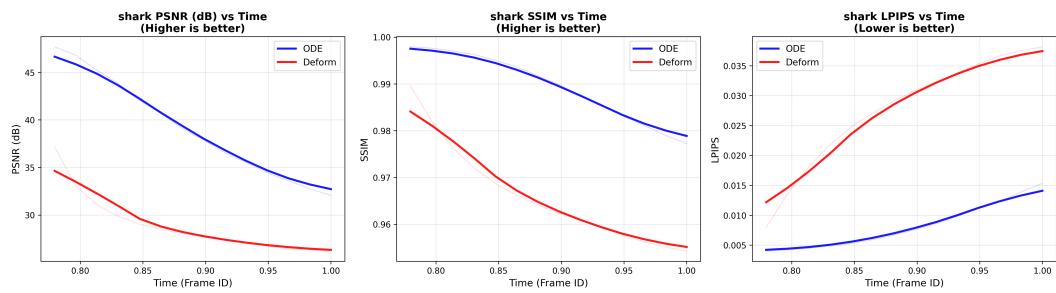


Figure 14: Shark Metrics Comparison

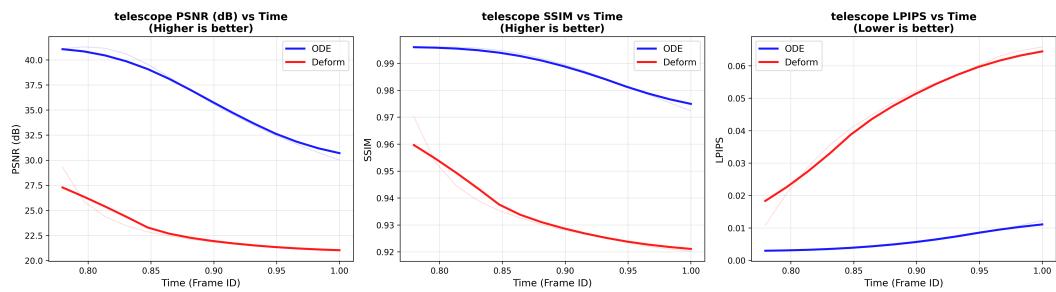


Figure 15: Telescope Metrics Comparison

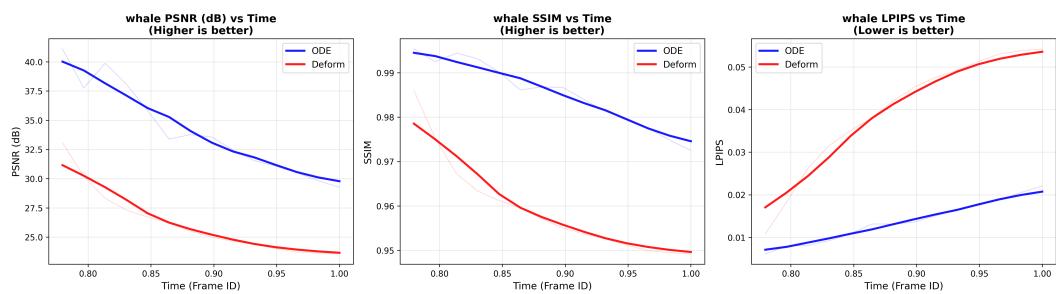


Figure 16: Whale Metrics Comparison

1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141



1142 Figure 17: Qualitative comparison of last frame from scene bat of NVFi by only using an autoregressive  
1143 transformer without the ODE component. From left to right is the GT image, the transformer's  
1144 rendered image, the baseline (Yang et al., 2024)'s result, the difference between the GT and trans-  
1145 former's rendered image, and finally the difference between the GT and the baseline's rendered  
1146 image.

1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155



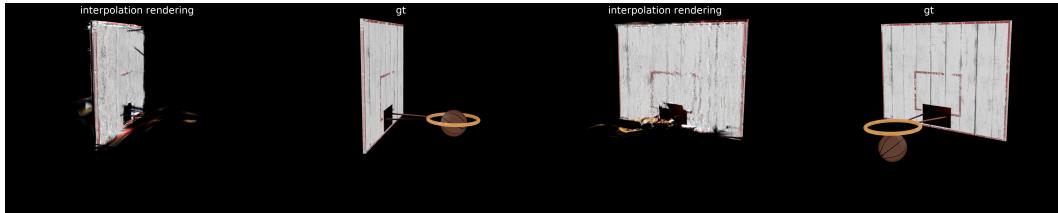
1156 Figure 18: Qualitative comparison of last frame from scene chess of NVFi by only using an au-  
1157 toregressive transformer without the ODE component. From left to right is the GT image, the  
1158 transformer's rendered image, the baseline (Yang et al., 2024)'s result, the difference between the  
1159 GT and transformer's rendered image, and finally the difference between the GT and the baseline's  
1160 rendered image.

1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169



1170 Figure 19: Qualitative comparison of last frame from scene factory of NVFi by only using an au-  
1171 toregressive transformer without the ODE component. From left to right is the GT image, the  
1172 transformer's rendered image, the baseline (Yang et al., 2024)'s result, the difference between the  
1173 GT and transformer's rendered image, and finally the difference between the GT and the baseline's  
1174 rendered image.

1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187



1184 Figure 20: Qualitative comparison of the fallingball scene in NVFi on the interpolation model and the  
1185 ground truth. As shown, the interpolation model fails to capture meaningful dynamics of the scene  
1186 while missing the fallingball object.