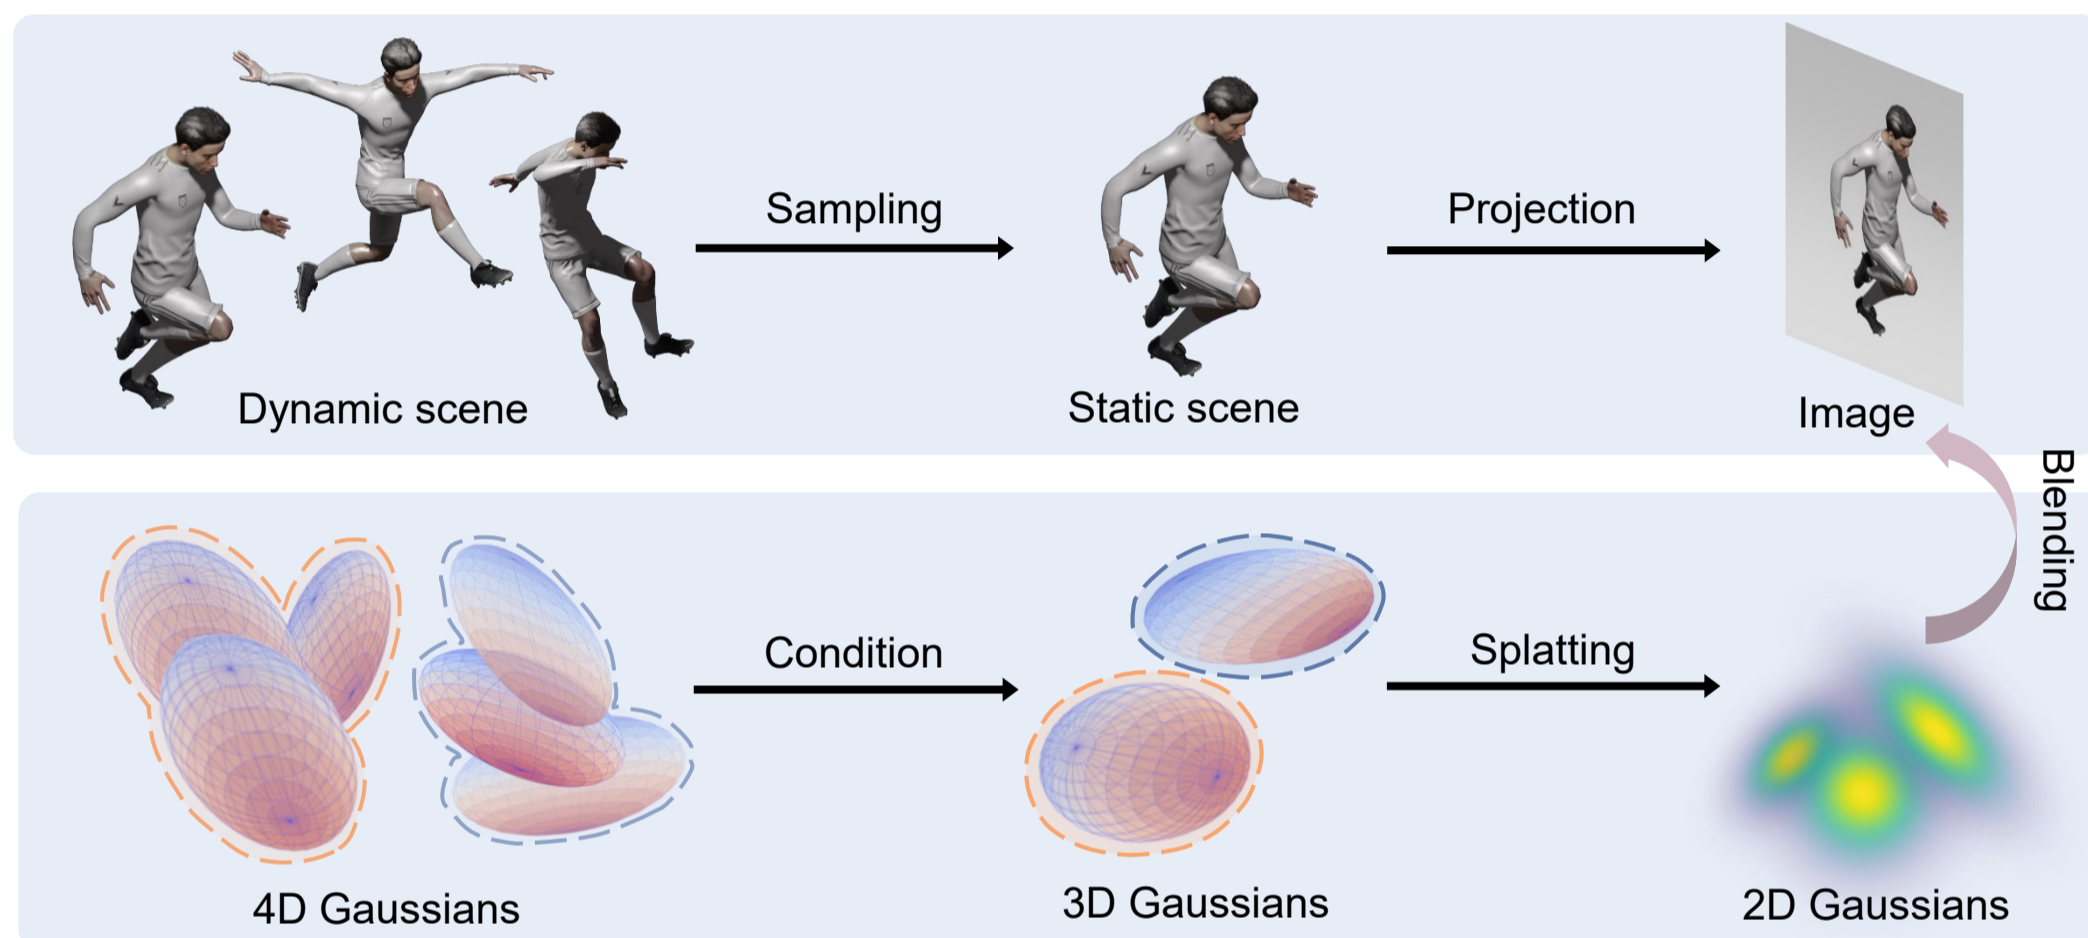





Zeyu Yang, Hongye Yang, Zijie Pan, Li Zhang
Fudan University

Overview


- The reconstruction of a dynamic scene can be recast as optimizing a series of native 4D Gaussians to fit its underlying spatiotemporal 4D volume.
- Efficient end-to-end optimization.
- Fully interpretable representation, which inherits the merits of 3D Gaussian Splatting, thus friendly for editing and composition.



- The rendering process of 4D Gaussian Splatting has conceptual parallels with the imaging process of a dynamic scenes:



Parameterization of 4D Gaussian



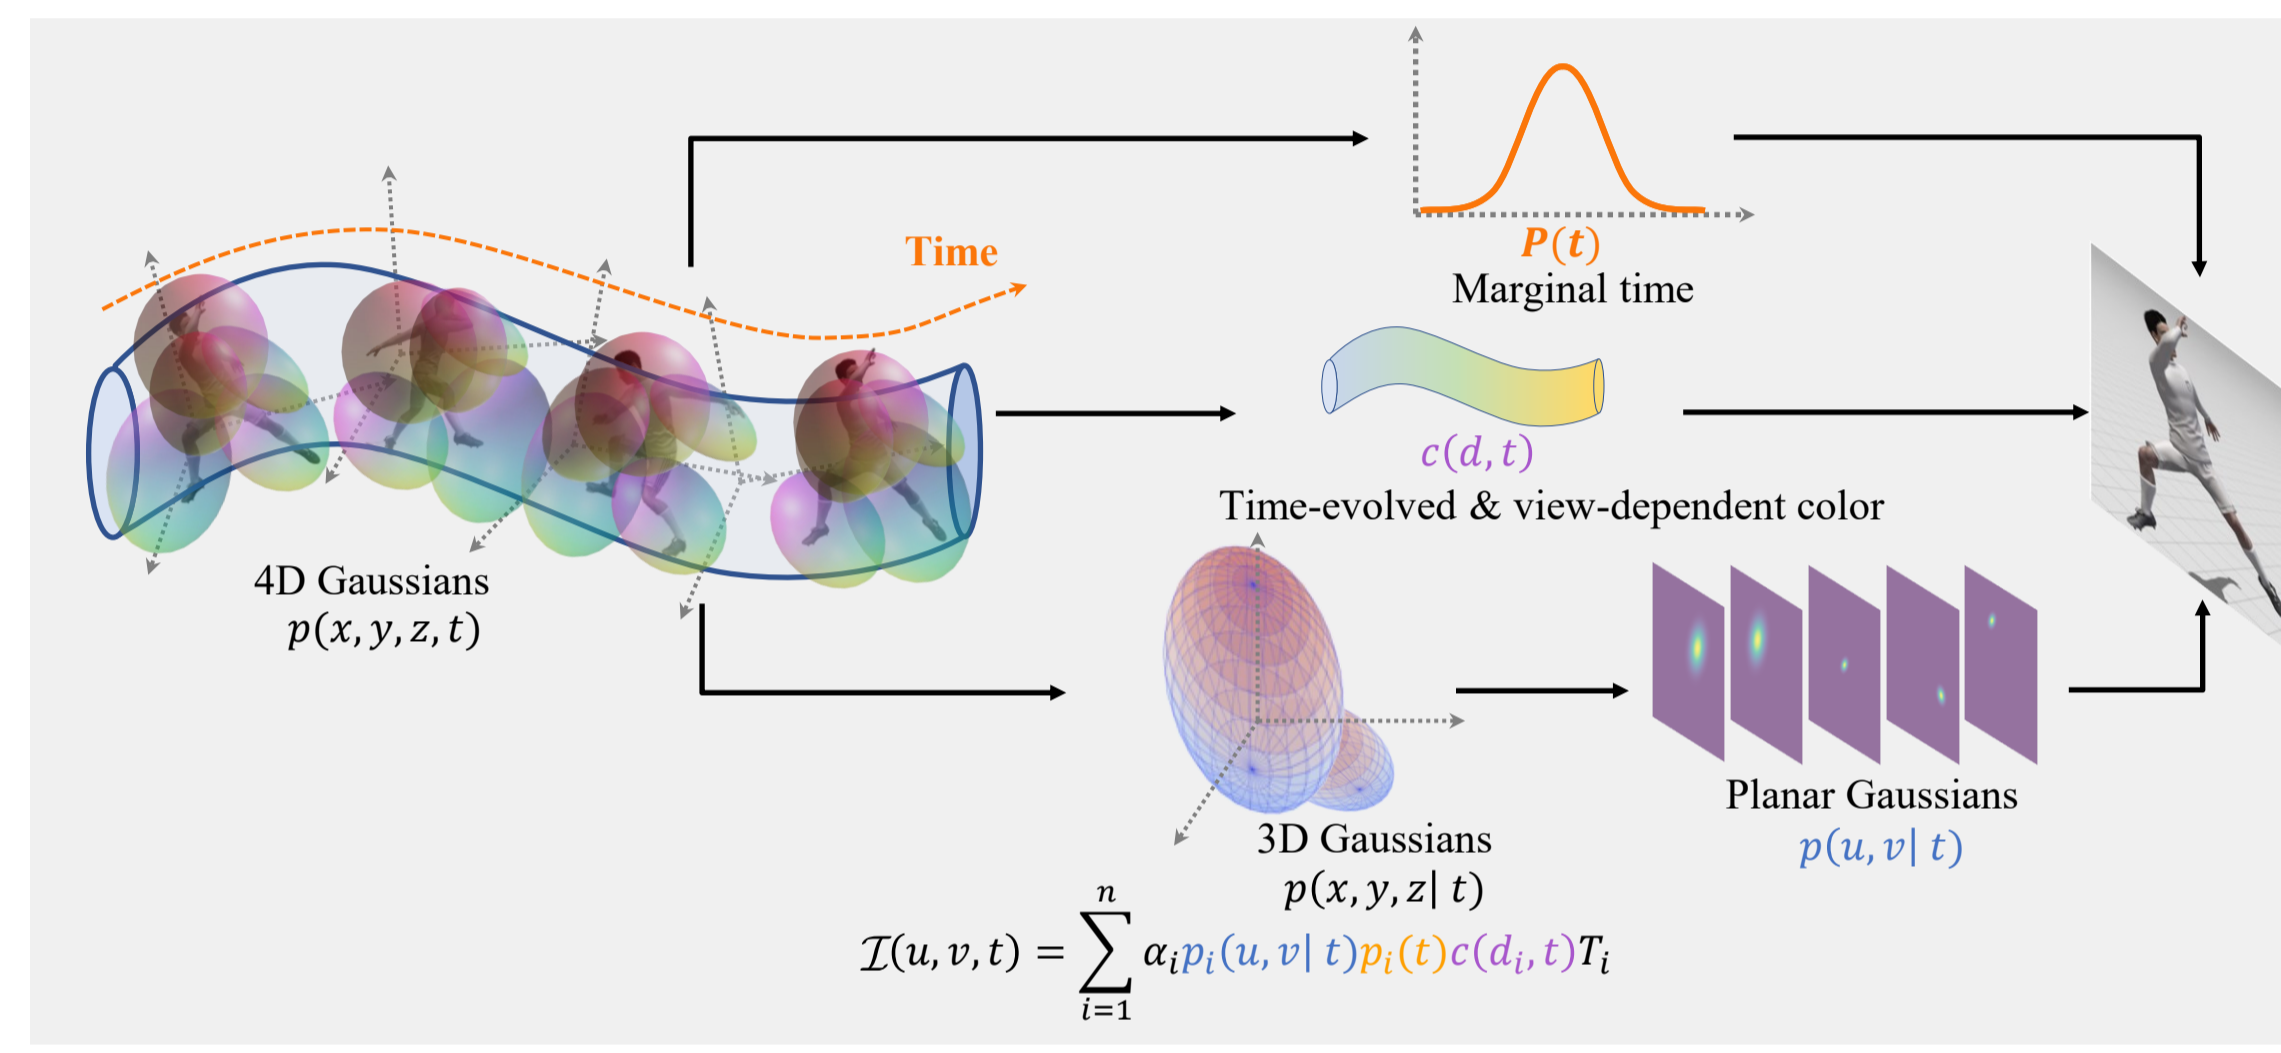
$$p = \begin{bmatrix} p_x \\ p_y \\ p_z \\ p_t \end{bmatrix} \quad \text{opacity} \quad \text{SH}$$

$$\Sigma = \begin{bmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} & \sigma_{xt} \\ \sigma_{yy} & \sigma_{yz} & \sigma_{yt} \\ \sigma_{zz} & \sigma_{zt} \\ \sigma_{tt} \end{bmatrix} = RSS^T R^T \quad S = \text{diag}(s_x, s_y, s_z, s_t)$$

$$R = L(q_l)R(q_r) = \begin{bmatrix} a & -b & -c & -d \\ b & a & -d & c \\ c & d & a & -b \\ d & -c & b & a \end{bmatrix} \begin{bmatrix} p & -q & -r & -s \\ q & p & s & -r \\ r & -s & p & q \\ s & r & -q & p \end{bmatrix} \quad q_l = (a, b, c, d) \\ q_r = (p, q, r, s)$$

Efficient rendering via 4D Gaussian splatting

- The influence of a 4D Gaussian on a pixel at time t can be evaluated by first decomposing it into its spatial conditional distribution and temporal marginal distribution, which are 3D and 1D Gaussians respectively, and then splatting the conditional 3D Gaussian weighted by the marginal Gaussian at given viewpoint.



Results

- State-of-the-art performance and efficiency for neural volumetric video.



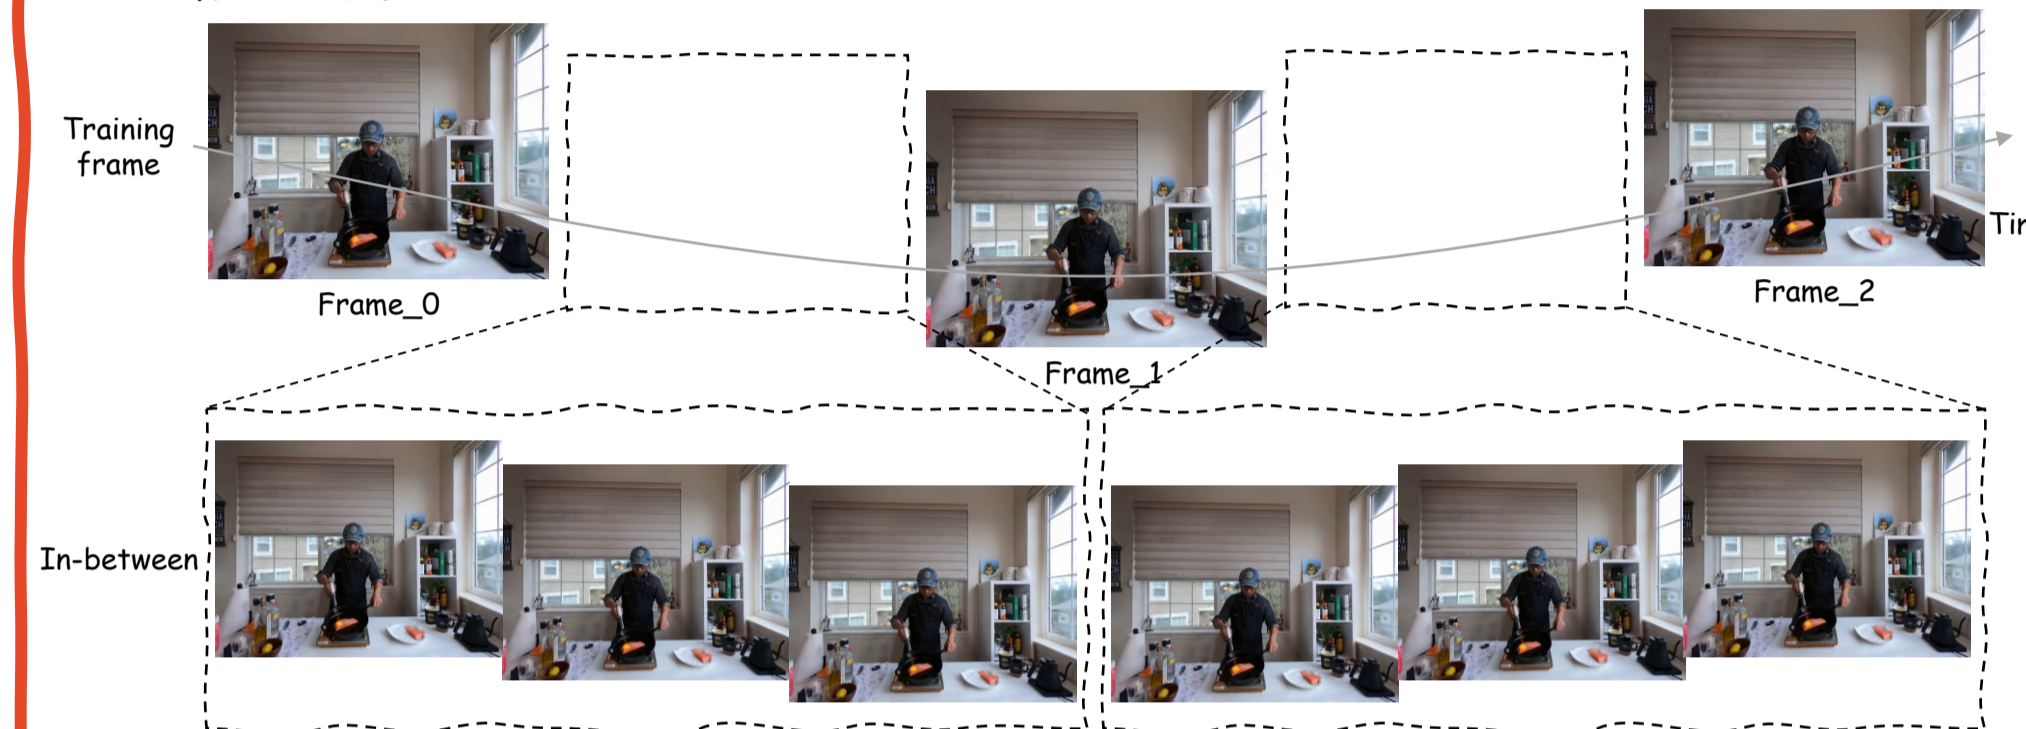
Method	PSNR ↑	DSSIM ↓	LPIPS ↓	FPS ↑
<i>- Plenoptic Video (real, multi-view)</i>				
Neural Volumes (Lombardi et al., 2019) ¹	22.80	0.062	0.295	-
LLFF (Mildenhall et al., 2019) ¹	23.24	0.076	0.235	-
DyNeRF (Li et al., 2022b) ²	29.58	0.020	0.099	0.015
HexPlane (Cao & Johnson, 2023)	31.70	0.014	0.075	0.56 ³
K-Planes-explicit (Fridovich-Keil et al., 2023)	30.88	0.020	-	0.23 ³
K-Planes-hybrid (Fridovich-Keil et al., 2023)	31.63	0.018	-	-
MixVoxels-L (Wang et al., 2023)	30.80	0.020	0.126	16.7
StreamRF (Li et al., 2022a) ¹	29.58	-	-	8.3
NeRFPlayer (Song et al., 2023)	30.69	0.035 ²	0.111	0.045
HyperReel (Attal et al., 2023)	31.10	0.037 ²	0.096	2.00
4DGS (Wu et al., 2023) ⁴	31.02	0.030	0.150	36
4DGS (Ours)	32.01	0.014	0.055	114

Method	PSNR ↑	SSIM ↑	LPIPS ↓
<i>- D-NeRF (synthetic, monocular)</i>			
T-NeRF (Pumarola et al., 2021)	29.51	0.95	0.08
D-NeRF (Pumarola et al., 2021)	29.67	0.95	0.07
TiNeuVox (Fang et al., 2022)	32.67	0.97	0.04
HexPlanes (Cao & Johnson, 2023)	31.04	0.97	0.04
K-Planes-explicit (Fridovich-Keil et al., 2023)	31.05	0.97	-
K-Planes-hybrid (Fridovich-Keil et al., 2023)	31.61	0.97	-
V4D (Gan et al., 2023)	33.72	0.98	0.02
4DGS (Wu et al., 2023) ¹	33.30	0.98	0.03
4DGS (Ours)	34.09	0.98	0.02


* Videos corresponding to all the images above are available at: <https://fudan-zvg.github.io/4d-gaussian-splatting>.

Results

- 4D Gaussian Splatting can correctly model the smooth inter-frame motions.



- 4D Gaussian Splatting can capture the underlying dynamics with only photometric loss.



- 4D Gaussian Splatting can be also applied for the reconstruction of large-scale urban scenes without relying on the laborious foreground annotations.

