

SharpTimeGS: Sharp and Stable Dynamic Gaussian Splatting via Lifespan Modulation

Supplementary Material

6.1. Method Details

Definition of E and L_e . E represents the accumulated error of Gaussian primitives from images. For each Gaussian primitive, we assign an additional scalar e_k to it and render a corresponding map R_e . We then introduce the auxiliary loss:

$$L_e = \sum \nabla[\mathcal{E}(p)] R_e(p), \quad (9)$$

$$\mathcal{E} = \left\| \tilde{I} - I_{gt} \right\|_1,$$

where $\nabla[\cdot]$, \mathcal{E} , \tilde{I} , and I_{gt} are the stop-gradient operation, the error map, the rendered image, and the ground-truth image, respectively. We initialize e_k to 0 for each Gaussian primitive and never update it during training. In this way, $L_e = 0$ and other parameters of Gaussian primitives are left invariant by L_e . The gradient with respect to e_k yields which is the per-Gaussian-primitive error E_k for one view and E is defined as the maximum E_k for each densification:

$$\frac{\partial L_e}{\partial e_k} = \sum \mathcal{E}(p) w(p) = E_k, \quad (10)$$

$$E = \max_k \{E_k\}.$$

Finally, E provided a comprehensive evaluation of whether the Gaussian primitives contained sufficient error for densification. It was then utilized to calculate the scoring metric s .

6.2. More Experiments Results

More visual comparison on the Neural3DV Dataset [19]. Additional qualitative comparisons on the Neural3DV [19] dataset are shown in Figure 7. As shown in Figure 7, our method can achieve better results in both static regions and dynamic regions (e.g., window and knife). These results indicate that the proposed 4D representation can represent both static and dynamic regions accurately. More sequences are available in the supplementary video.

More visual comparison on the ENeRF-Outdoor Dataset [38]. More qualitative comparisons on the ENeRF-Outdoor Dataset [38] dataset are shown in Figure 8. Our approach better preserves static geometry (e.g., wall) and produces sharper, more consistent reconstructions of dynamic content (e.g., watermelon), demonstrating strong robustness under large outdoor motions. Please refer to the supplementary video for more results.

More visual comparison on the SelfCap Dataset [47]. Additional qualitative results on the SelfCap dataset [47], which features rapid and complex human motions, are provided in the supplementary video. SharpTimeGS maintains static background fidelity while capturing fast, fine-grained movements more accurately than existing methods.

6.3. Rendering Speed Analysis

Our approach demonstrates the capability of achieving real-time rendering at 4K resolution, exceeding 100 frames per second (FPS) on a single NVIDIA RTX 4090 graphics card. We show the results of our interactive rendering at the end of the supplementary video and Figure 6. As can be seen in the video, our method can achieve high speed rendering at 4K resolution. This remarkable performance showcases the efficiency and optimization of our rendering technique, making it feasible for applications requiring high visual fidelity and responsiveness.

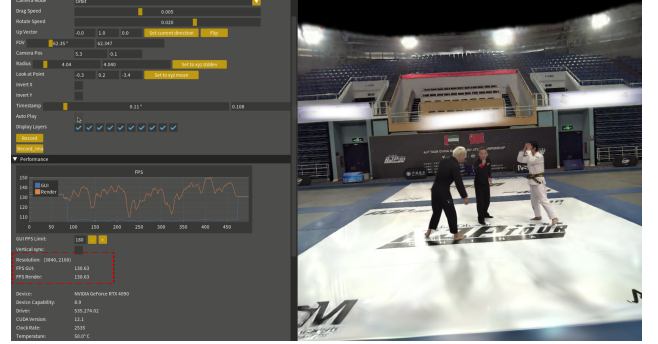


Figure 6. SharpTimeGS achieves real-time rendering at 4K resolution with over 100 FPS on a single RTX 4090. The system maintains both high image fidelity and stable frame rate during free-viewpoint navigation. For anonymity compliance required by double-blind policy, identity-related regions have been intentionally obfuscated.

6.4. Efficiency.

The results are reported in Tab. 3. Our method consistently outperforms prior works. These gains are enabled by our representation, which achieves high visual quality with fewer Gaussian primitives, resulting in faster training convergence, higher rendering speed, and smaller model size.




Figure 7. More visual results on the Neural3DV Dataset [19]. Our method can achieve better performance.




Figure 8. More visual results on the ENeRF-Outdoor Dataset [38]. Our method can achieve better performance.

References

- [1] Benjamin Attal, Jia-Bin Huang, Christian Richardt, Michael Zollhoefer, Johannes Kopf, Matthew O’Toole, and Changil

Kim. Hyperreel: High-fidelity 6-dof video with ray-conditioned sampling. In *CVPR*, pages 16610–16620, 2023.

- [2] Jeongmin Bae, Seoha Kim, Youngsik Yun, Hahyun Lee, Gun

Table 3. Efficiency comparison on SelfCap Dataset. We report PSNR, total cost time (preprocessing time + training time; 60 frames), rendering speed (on single NVIDIA RTX 4090 GPU), and model size.

Method	PSNR \uparrow	Cost Time (Min) \downarrow	Speed (fps) \uparrow	Size (MB) \downarrow
Deformable-3DGS	25.85	93 (3 + 90)	57	92
4DGS	25.86	163 (3 + 160)	142	827
STGS	24.77	107 (15 + 92)	65	87
FreeTimeGS	27.50	88 (15 + 73)	267	96
Ours	28.14	62 (17 + 45)	327	78

- Bang, and Youngjung Uh. Per-gaussian embedding-based deformation for deformable 3d gaussian splatting. In *ECCV*, pages 321–335. Springer, 2024. 3
- [3] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *CVPR*, pages 130–141, 2023. 1, 2, 3
- [4] Dan Casas, Marco Volino, John Collomosse, and Adrian Hilton. 4d video textures for interactive character appearance. In *Computer Graphics Forum*, pages 371–380. Wiley Online Library, 2014. 1
- [5] Danpeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. Pgsr: Planar-based gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE TVCG*, 2024. 5
- [6] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM TOG*, 34(4):1–13, 2015. 1
- [7] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM TOG*, 35(4):1–13, 2016. 1
- [8] Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4d-rotor gaussian splatting: towards efficient novel view synthesis for dynamic scenes. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 1, 3
- [9] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *CVPR*, pages 12479–12488, 2023. 1, 2, 3
- [10] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE TPAMI*, 32(8):1362–1376, 2009. 1
- [11] Zhongpai Gao, Benjamin Planche, Meng Zheng, Anwesa Choudhuri, Terrence Chen, and Ziyang Wu. 7dgs: Unified spatial-temporal-angular gaussian splatting. *arXiv preprint arXiv:2503.07946*, 2025. 3
- [12] Zhiyang Guo, Wengang Zhou, Li Li, Min Wang, and Houqiang Li. Motion-aware 3d gaussian splatting for efficient dynamic scene reconstruction. *IEEE TCSVT*, 2024. 3
- [13] Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. Humanrf: High-fidelity neural radiance fields for humans in motion. *ACM TOG*, 42(4):1–12, 2023. 2
- [14] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 1, 3, 4, 5, 7
- [15] Mijeong Kim, Jongwoo Lim, and Bohyung Han. 4d gaussian splatting in the wild with uncertainty-aware regularization. *Advances in Neural Information Processing Systems*, 37:129209–129226, 2024. 3
- [16] Seoha Kim, Jeongmin Bae, Youngsik Yun, Hahyun Lee, Gun Bang, and Youngjung Uh. Sync-nerf: Generalizing dynamic nerfs to unsynchronized videos. In *AAAI*, pages 2777–2785, 2024. 2
- [17] Isaac Labe, Noam Issachar, Itai Lang, and Sagie Benaim. Dgd: Dynamic 3d gaussians distillation. In *European Conference on Computer Vision*, pages 361–378. Springer, 2024. 3
- [18] Junoh Lee, ChangYeon Won, Hyunjun Jung, Inhwon Bae, and Hae-Gon Jeon. Fully explicit dynamic gaussian splatting. *Advances in Neural Information Processing Systems*, 37:5384–5409, 2024. 3, 5, 8
- [19] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *CVPR*, pages 5521–5531, 2022. 1, 2, 5, 8, 10, 11
- [20] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *CVPR*, pages 8508–8520, 2024. 1, 3, 5, 7, 8
- [21] Yiqing Liang, Numair Khan, Zhengqin Li, Thu Nguyen-Phuoc, Douglas Lanman, James Tompkin, and Lei Xiao. Gafre: Gaussian deformation fields for real-time dynamic novel view synthesis. pages 2642–2652. *IEEE*, 2025. 1, 3
- [22] Haotong Lin, Sida Peng, Zhen Xu, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Efficient neural radiance fields for interactive free-viewpoint video. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 2, 5, 7, 8
- [23] Christian Lipski, Christian Linz, Kai Berger, and Marcus Magnor. Virtual video camera: Image-based viewpoint navigation through space and time. In *SIGGRAPH'09: Posters*, pages 1–1, 2009. 1
- [24] Zhicheng Lu, Xiang Guo, Le Hui, Tianrui Chen, Min Yang, Xiao Tang, Feng Zhu, and Yuchao Dai. 3d geometry-aware deformable gaussian splatting for dynamic view synthesis. In *CVPR*, pages 8900–8910, 2024. 3
- [25] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *International Conference on 3D Vision (3DV)*, 2024. 3
- [26] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2
- [27] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, pages 5845–5854, 2021. 2
- [28] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-

- Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ACM TOG*, 40(6), 2021. 1, 2
- [29] Sida Peng, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Representing volumetric videos as dynamic mlp maps. In *CVPR*, pages 4252–4262, 2023.
- [30] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, pages 10318–10327, 2021. 1, 2
- [31] LIU Qingming, Yuan Liu, Jiepeng Wang, Xianqiang Lyu, Peng Wang, Wenping Wang, and Junhui Hou. Modgs: Dynamic gaussian splatting from casually-captured monocular videos with depth priors. In *ICLR*, 2025. 3
- [32] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 5
- [33] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. Revising densification in gaussian splatting. In *ECCV*, pages 347–362. Springer, 2024. 4
- [34] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, pages 4104–4113, 2016. 1, 5
- [35] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *CVPR*, pages 16632–16642, 2023. 2, 3
- [36] Richard Shaw, Michal Nazarczuk, Jifei Song, Arthur Moreau, Sibi Catley-Chandar, Helisa Dharmo, and Eduardo Pérez-Pellitero. Swings: sliding windows for dynamic 3d gaussian splatting. In *ECCV*, pages 37–54. Springer, 2024. 3
- [37] Liangchen Song, Anpei Chen, Zhong Li, Zhang Chen, Lele Chen, Junsong Yuan, Yi Xu, and Andreas Geiger. Nerf-player: A streamable dynamic scene representation with decomposed neural radiance fields. *IEEE TVCG*, 29(5): 2732–2742, 2023. 2
- [38] Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. Fenerf: Face editing in neural radiance fields. In *CVPR*, pages 7662–7672, 2022. 10, 11
- [39] Cong Wang, Di Kang, Yanpei Cao, Linchao Bao, Ying Shan, and Song-Hai Zhang. Neural point-based volumetric avatar: Surface-guided neural points for efficient and photorealistic volumetric head avatar. In *ACM SIGGRAPH Asia 2023 Conference Proceedings*, 2023. 2
- [40] Feng Wang, Sinan Tan, Xinghang Li, Zeyue Tian, Yafei Song, and Huaping Liu. Mixed neural voxels for fast multi-view video synthesis. In *CVPR*, pages 19706–19716, 2023. 2
- [41] Yihan Wang, Lahav Lipson, and Jia Deng. Sea-raft: Simple, efficient, accurate raft for optical flow. In *ECCV*, pages 36–54. Springer, 2024. 5
- [42] Yifan Wang, Peishan Yang, Zhen Xu, Jiaming Sun, Zhanhua Zhang, Yong Chen, Hujun Bao, Sida Peng, and Xiaowei Zhou. Freetimegs: Free gaussian primitives at anytime anywhere for dynamic scene reconstruction. In *CVPR*, pages 21750–21760, 2025. 1, 2, 3, 5, 7, 8
- [43] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. 5, 8
- [44] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *CVPR*, pages 20310–20320, 2024. 1
- [45] Jiawei Xu, Zexin Fan, Jian Yang, and Jin Xie. Grid4d: 4d decomposed hash encoding for high-fidelity dynamic gaussian splatting. *Advances in Neural Information Processing Systems*, 37:123787–123811, 2024. 3
- [46] Zhen Xu, Sida Peng, Haotong Lin, Guangzhao He, Jiaming Sun, Yujun Shen, Hujun Bao, and Xiaowei Zhou. 4k4d: Real-time 4d view synthesis at 4k resolution. In *CVPR*, pages 20029–20040, 2024. 3
- [47] Zhen Xu, Yinghao Xu, Zhiyuan Yu, Sida Peng, Jiaming Sun, Hujun Bao, and Xiaowei Zhou. Representing long volumetric video with temporal gaussian hierarchy. *ACM TOG*, 43(6):1–18, 2024. 2, 3, 6, 7, 8, 10
- [48] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *CVPR*, 2024. 1, 3, 5, 8
- [49] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. In *ICLR*, 2024. 1, 3, 5, 7, 8
- [50] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. Absgs: Recovering fine details in 3d gaussian splatting. In *ACM MM*, pages 1053–1061, 2024. 4
- [51] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. 5, 8
- [52] Ruijie Zhu, Yanzhe Liang, Hanzhi Chang, Jiacheng Deng, Jiahao Lu, Wenfei Yang, Tianzhu Zhang, and Yongdong Zhang. Motiongs: Exploring explicit motion guidance for deformable 3d gaussian splatting. In *NeurIPS*, pages 101790–101817, 2024. 3