

LAYEREDGS: EFFICIENT DYNAMIC SCENE RENDERING AND POINT TRACKING WITH MULTI-LAYER DEFORMABLE GAUSSIAN SPLATTING

SUPPLEMENTARY MATERIAL

Anonymous authors

Paper under double-blind review

In this supplementary material, we offer additional details regarding our FastParticle and Panoptic datasets, which provide the necessary context for our experiments. We delve into our method for articulated objects segmentation, presenting full qualitative results that demonstrate its effectiveness across various scenarios. Additionally, we clarify our rationale for maintaining the same number of iterations in our comparisons and present a comparison under equal wall-clock time, showing that our method still outperforms Dynamic3DGS Luiten et al. (2023). We also include visualizations illustrating the multi-layer clustering structure we employ, as well as the manually annotated tracking labels used for evaluating 2D tracking results. Furthermore, we discuss our approach to learning the deformation, emphasizing the two-phase training strategy. Finally, we reflect on limitations, identifying potential areas for future improvement.

A FASTPARTICLE AND PANOPTIC DATASETS

In this section, we introduce the FastParticle and Panoptic datasets used in our experiments in details. The real-world Panoptic dataset includes six scenes: Basketball, Boxes, Football, Juggle, Softball, and Tennis. Each frame in these scenes comes with segmentation provided by the original authors. Following Luiten et al. (2023), we distinguish between foreground and background in these scenes and utilize background loss and floor loss accordingly. Each scene in this dataset contains 150 frames captured by a total of 31 cameras, with 27 cameras used for training and 4 for testing.

The synthetic FastParticle dataset, which we have accelerated, contains six dynamic scenes: Robot, Spring, Wheel, Pendulums, Robot-Task, and Cloth. After acceleration, these scenes respectively have 35, 18, 38, 24, 35, and 35 frames. As illustrated in fig. I, we show the dynamic evolution of some scenes, highlighting the significant changes between frames. This dataset includes 40 cameras in total, from which we randomly select 4 as testing cameras and the remaining 36 as training cameras.

For all experiments, we provide the same static checkpoints to all baselines. For the 12 scenes across the two datasets, we train for 20,000 iterations to obtain the checkpoints. Due to the varying complexity of the static scenes, 3,000 iterations are sufficient for most FastParticle scenes.

B ARTICULATED OBJECTS SEGMENTATION

As mentioned in Sec. 5.1. The intuition behind the KMeans design is that, (1) Gaussians belong to the same part of the object should be close to each other at all time, and (2) the rotations of Gaussians within the same rigid part should be the same. The first one can be trivial, here we provide more explanations about the second point. As shown in fig. II, suppose we have a rigid body with its centroid denoted as C_0 . This rigid body can be considered as a combination of two smaller rigid bodies, with their centroids denoted as C_1 and C_2 , respectively. After rotation, C_1 and C_2 move to C'_1 and C'_2 . Taking C_0 as the origin of the coordinate system, the movement of the rigid body can only be a rotation R around C_0 , and the two smaller rigid bodies move accordingly. When considering the left smaller rigid body alone, its motion should consist of a translation of its centroid C_1 and a rotation R_1 around C_1 . We aim to prove that $R_1 = R$. Therefore, consider a point P on the left rigid body, which moves to point P' after the movement. From the perspective of C_0 , we




Figure I: This figure shows the evolution of three scenes from the FastParticle dataset, demonstrating the high dynamic characteristics of the accelerated dataset.




Figure II: Illustration of a rigid body rotating R around its centroid. When considering the rigid body as composed of two smaller rigid bodies, it can be shown that the rotation of each smaller rigid body around its own centroid is the same with R .

have

$$\overrightarrow{C_0 P'} = R \overrightarrow{C_0 P}. \quad (1)$$

Also, from the perspective of C_1 , we can have

$$\begin{aligned} \overrightarrow{C_0 P'} &= R_1 \overrightarrow{C_1 P} + \overrightarrow{C_0 C_1'} \\ &= R_1 \overrightarrow{C_1 P} + R \overrightarrow{C_0 C_1}. \end{aligned} \quad (2)$$

Therefore, we have

$$R \overrightarrow{C_1 P} = R_1 \overrightarrow{C_1 P}. \quad (3)$$

Since the choice of P is arbitrary, we can conclude that $R_1 = R$. Similarly, we can prove that the rotation of the smaller rigid body on the right is also R . The above demonstrates the case where the rigid body is divided into two parts. This conclusion can be generalized to any case of multiple divisions, meaning that all parts of the same rigid body have the same rotation. Returning to our problem, since the rotation of Gaussians is around their centroids, the Gaussians belonging to the same rigid body should have the same rotation.

C FULL QUALITATIVE RESULTS




Figure III: Qualitative results on FastParticle

In this section, we provide qualitative results on all 12 scenes from the two datasets. As shown in fig. III and fig. IV, both our method and Luiten et al. (2023) are trained 100 iterations between two consecutive frames.

D SAME WALL-CLOCK TIME COMPARISONS

In our experiments, we use the same number of iterations across different methods for consistency. While wall-clock time may vary depending on the specific implementation (e.g., whether CUDA acceleration is employed), the number of iterations reflects the convergence speed of the algorithms. A lower number of iterations indicates faster convergence, showing that the optimization problem is




Figure IV: Qualitative results on Panoptic

Metrics	Method	FastParticle					
		Robot	Spring	Wheel	Pendulums	Robot-Task	Cloth
PSNR \uparrow	Ours ₁₀₀	29.46	30.28	27.95	30.60	27.67	31.68
	Dynamic3DGS ₃₀₀ Luiten et al. (2023)	27.66	27.16	26.67	29.57	26.79	30.41
SSIM \uparrow	Ours ₁₀₀	0.96	0.97	0.94	0.97	0.95	0.97
	Dynamic3DGS ₃₀₀ Luiten et al. (2023)	0.95	0.95	0.93	0.96	0.95	0.97
LPIPS \downarrow	Ours ₁₀₀	0.09	0.04	0.07	0.06	0.10	0.06
	Dynamic3DGS ₃₀₀ Luiten et al. (2023)	0.10	0.06	0.08	0.06	0.10	0.07

Table I: Comparison of our method trained with 100 iterations per time frame against Dynamic3DGS.




Figure V: Coarse-to-fine multi-layer clustering structures for two objects in the FastParticle dataset.

easier to solve. This practice is commonly used in the evaluation of online methods, as demonstrated in the Dynamic3DGS Luiten et al. (2023) comparison (see Table 1 in their paper), where different methods are also compared using the same number of iterations.

Even when comparing with equivalent wall time, our method remains superior. To further illustrate this, we provide a comparison of our method trained for 100 iterations per frame versus Dynamic3DGS Luiten et al. (2023) trained for 300 iterations per frame on the FastParticle dataset. The results show that our method has an average training speed per iteration approximately twice as fast as Dynamic3DGS Luiten et al. (2023). As seen in table I, despite the difference in iteration count, our method still outperforms Dynamic3DGS Luiten et al. (2023) in terms of both efficiency and final performance.

E ILLUSTRATION OF THE MULTI-LAYER STRUCTURE

In fig. V, we show the coarse-to-fine multi-layer clustering structures for two objects in the FastParticle dataset. Different colors in the figure represent different clusters, and for clarification, the same color in different layers does not indicate any correlation between the clusters.

F TRACKING LABELS

Here, as shown in fig. VI, we present all manually annotated 2D tracking ground truths. Since the human eye can only track points with distinct features across multiple frames, we only selected such points for annotation.

G LEARNING THE DEFORMATION

algorithm 1 summarizes our training process. Initially, we train our Gaussians on the static scene using observations from the first frame. Subsequently, we perform multilevel coarse-to-fine clustering for the centroids of the Gaussians. For each input in every time frame, we use an optimization approach to backpropagate loss and subsequently update our deformation functions.

For potential negative impacts, since LayeredGS can learn deformation information and be used for creating new motions or inserting objects, such applications can be used for fake news to convince people by multi-view renderings. More censorship needs to be established in such cases.




Figure VI: Illustration of our manually annotated tracking ground truths.

Algorithm 1: Deformation-based Dynamic Scene Reconstruction Algorithm

Input: Images from all frames
 $\Theta_{\text{prev}} \leftarrow$ Initialization stage (Static Gaussian Splatting);
Do Clustering;
for t **in** time_frames **do**
 Initialize the Deformation D ;
 for iter **in** max_iters **do**
 $\Theta_{\text{curr}} \leftarrow D(\Theta_{\text{prev}})$;
 Images $\leftarrow \text{Render}(\Theta_{\text{curr}})$;
 loss $\leftarrow \text{Loss}(\text{gt_Images}, \text{Images})$;
 Backpropagate(loss);
 end
end

H LIMITATIONS

While our method significantly reduces training iterations to 100 per frame, achieving real-time training and rendering remains a challenge. Additionally, the learned deformation information is not fully utilized, and the presented articulated object segmentation results are not well refined. Future work will focus on addressing these limitations by exploring real-time training approaches, refining deformation utilization techniques, and developing more sophisticated segmentation methods.

REFERENCES

Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023.