



**Finish Training Within ~100 Seconds**

*Compute **Where It Matters***


# TurboGS : Accelerating 3D Gaussian Splatting via **Error-Guided Sparse Pixel** Sampling and Optimization

Zheng Dong<sup>†1</sup>, Daifei Qiu<sup>2</sup>, Pinxuan Dai<sup>3</sup>, Ke Xu<sup>4</sup>, Jiamin Xu<sup>5</sup>, Lili He<sup>1</sup>, Rynson W. H. Lau<sup>4</sup>, Weiwei Xu<sup>†3</sup>



# Background : Why Do We Need **Fast 3DGS Optimization** ?

*Rome Images*



*3D Gaussian Scene*




*Ours : ~ 92s*



**> 12x Faster Training than Vanilla 3DGS**



*Digital Twin*



*Robotics Scene*



*Digital Human / Object*



*Content Creation*

**For Frequently Updated Scenes, Optimization Latency Becomes the Bottleneck**

# Baseline 3DGS : Inefficient Computation Allocation



**Gradient-Dominated** (No efficient density control)

3DGS : All Pixels Treated **Equally**



1.0 seond





10.0 seonds



20.0 seonds

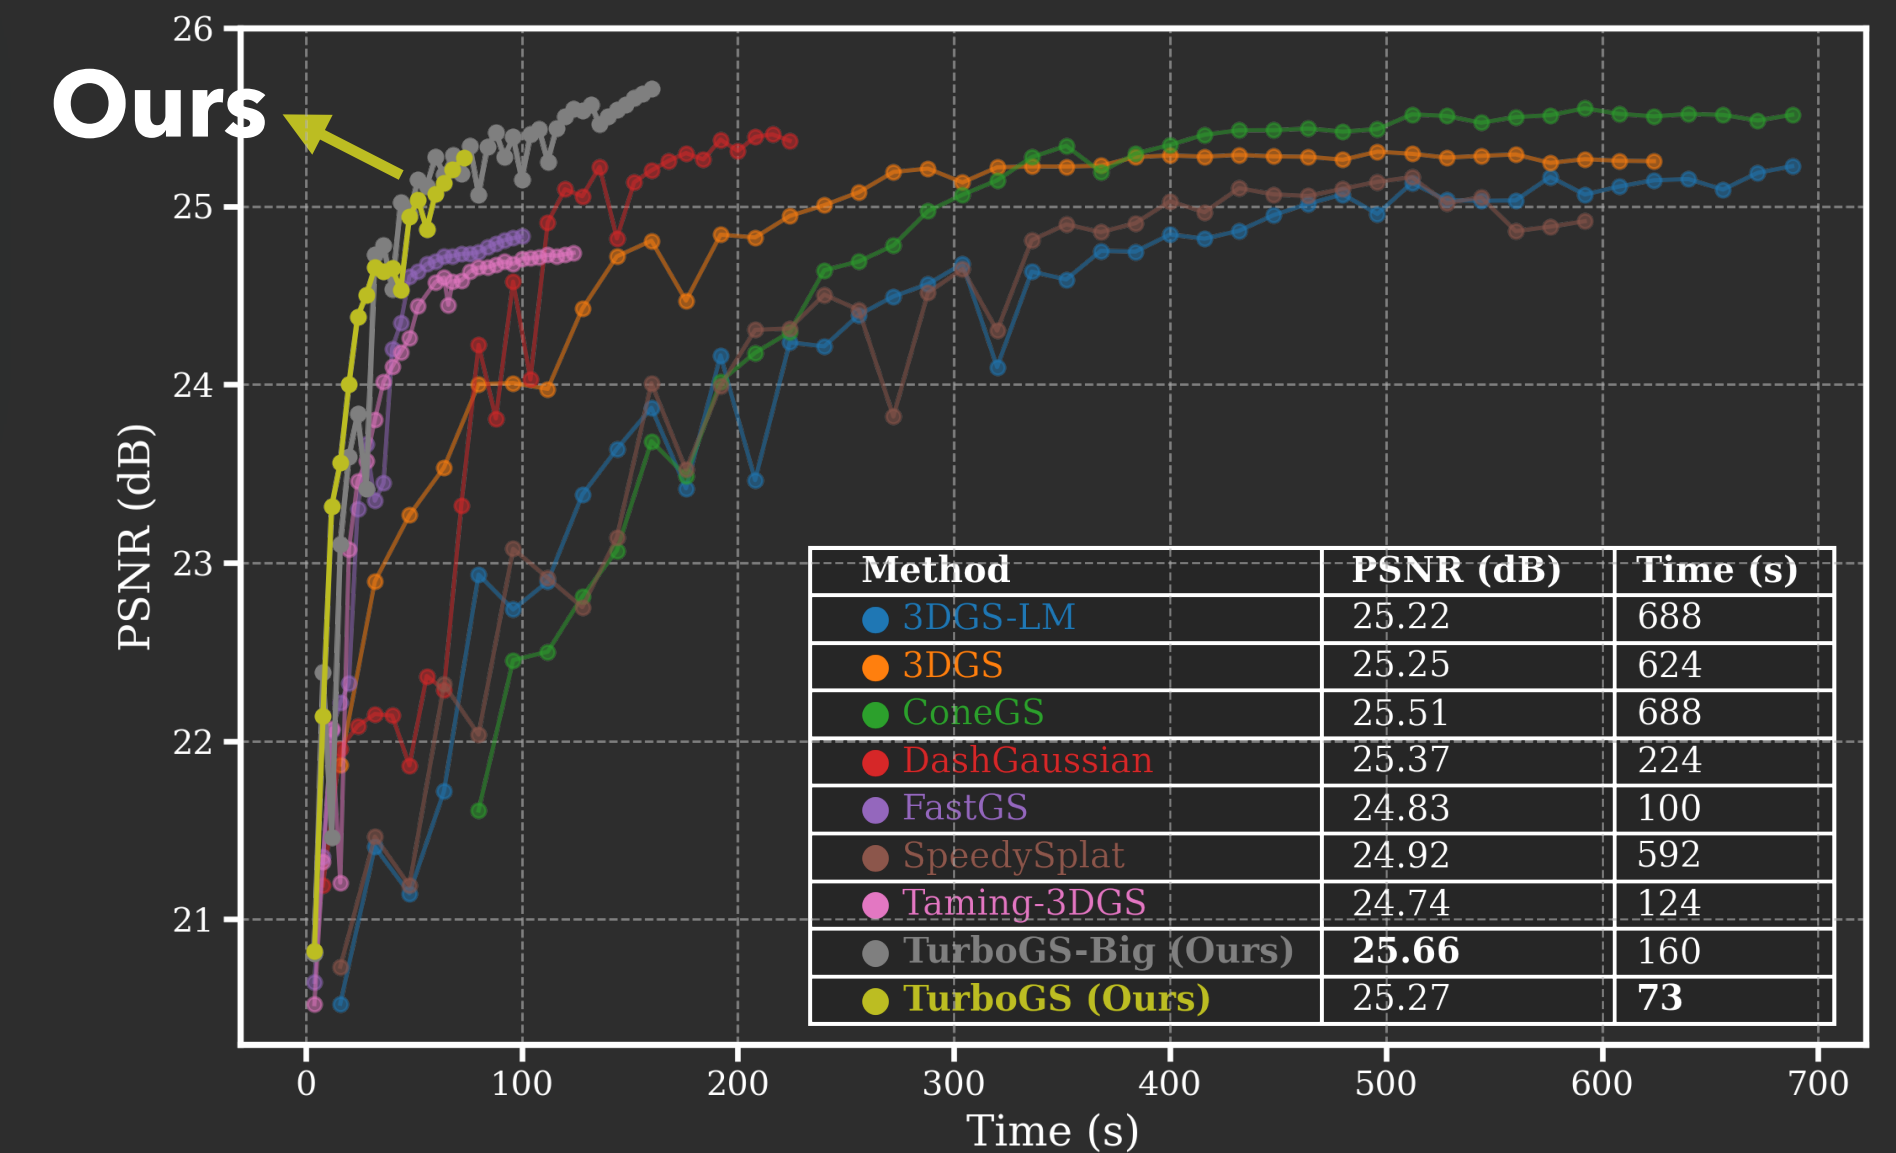
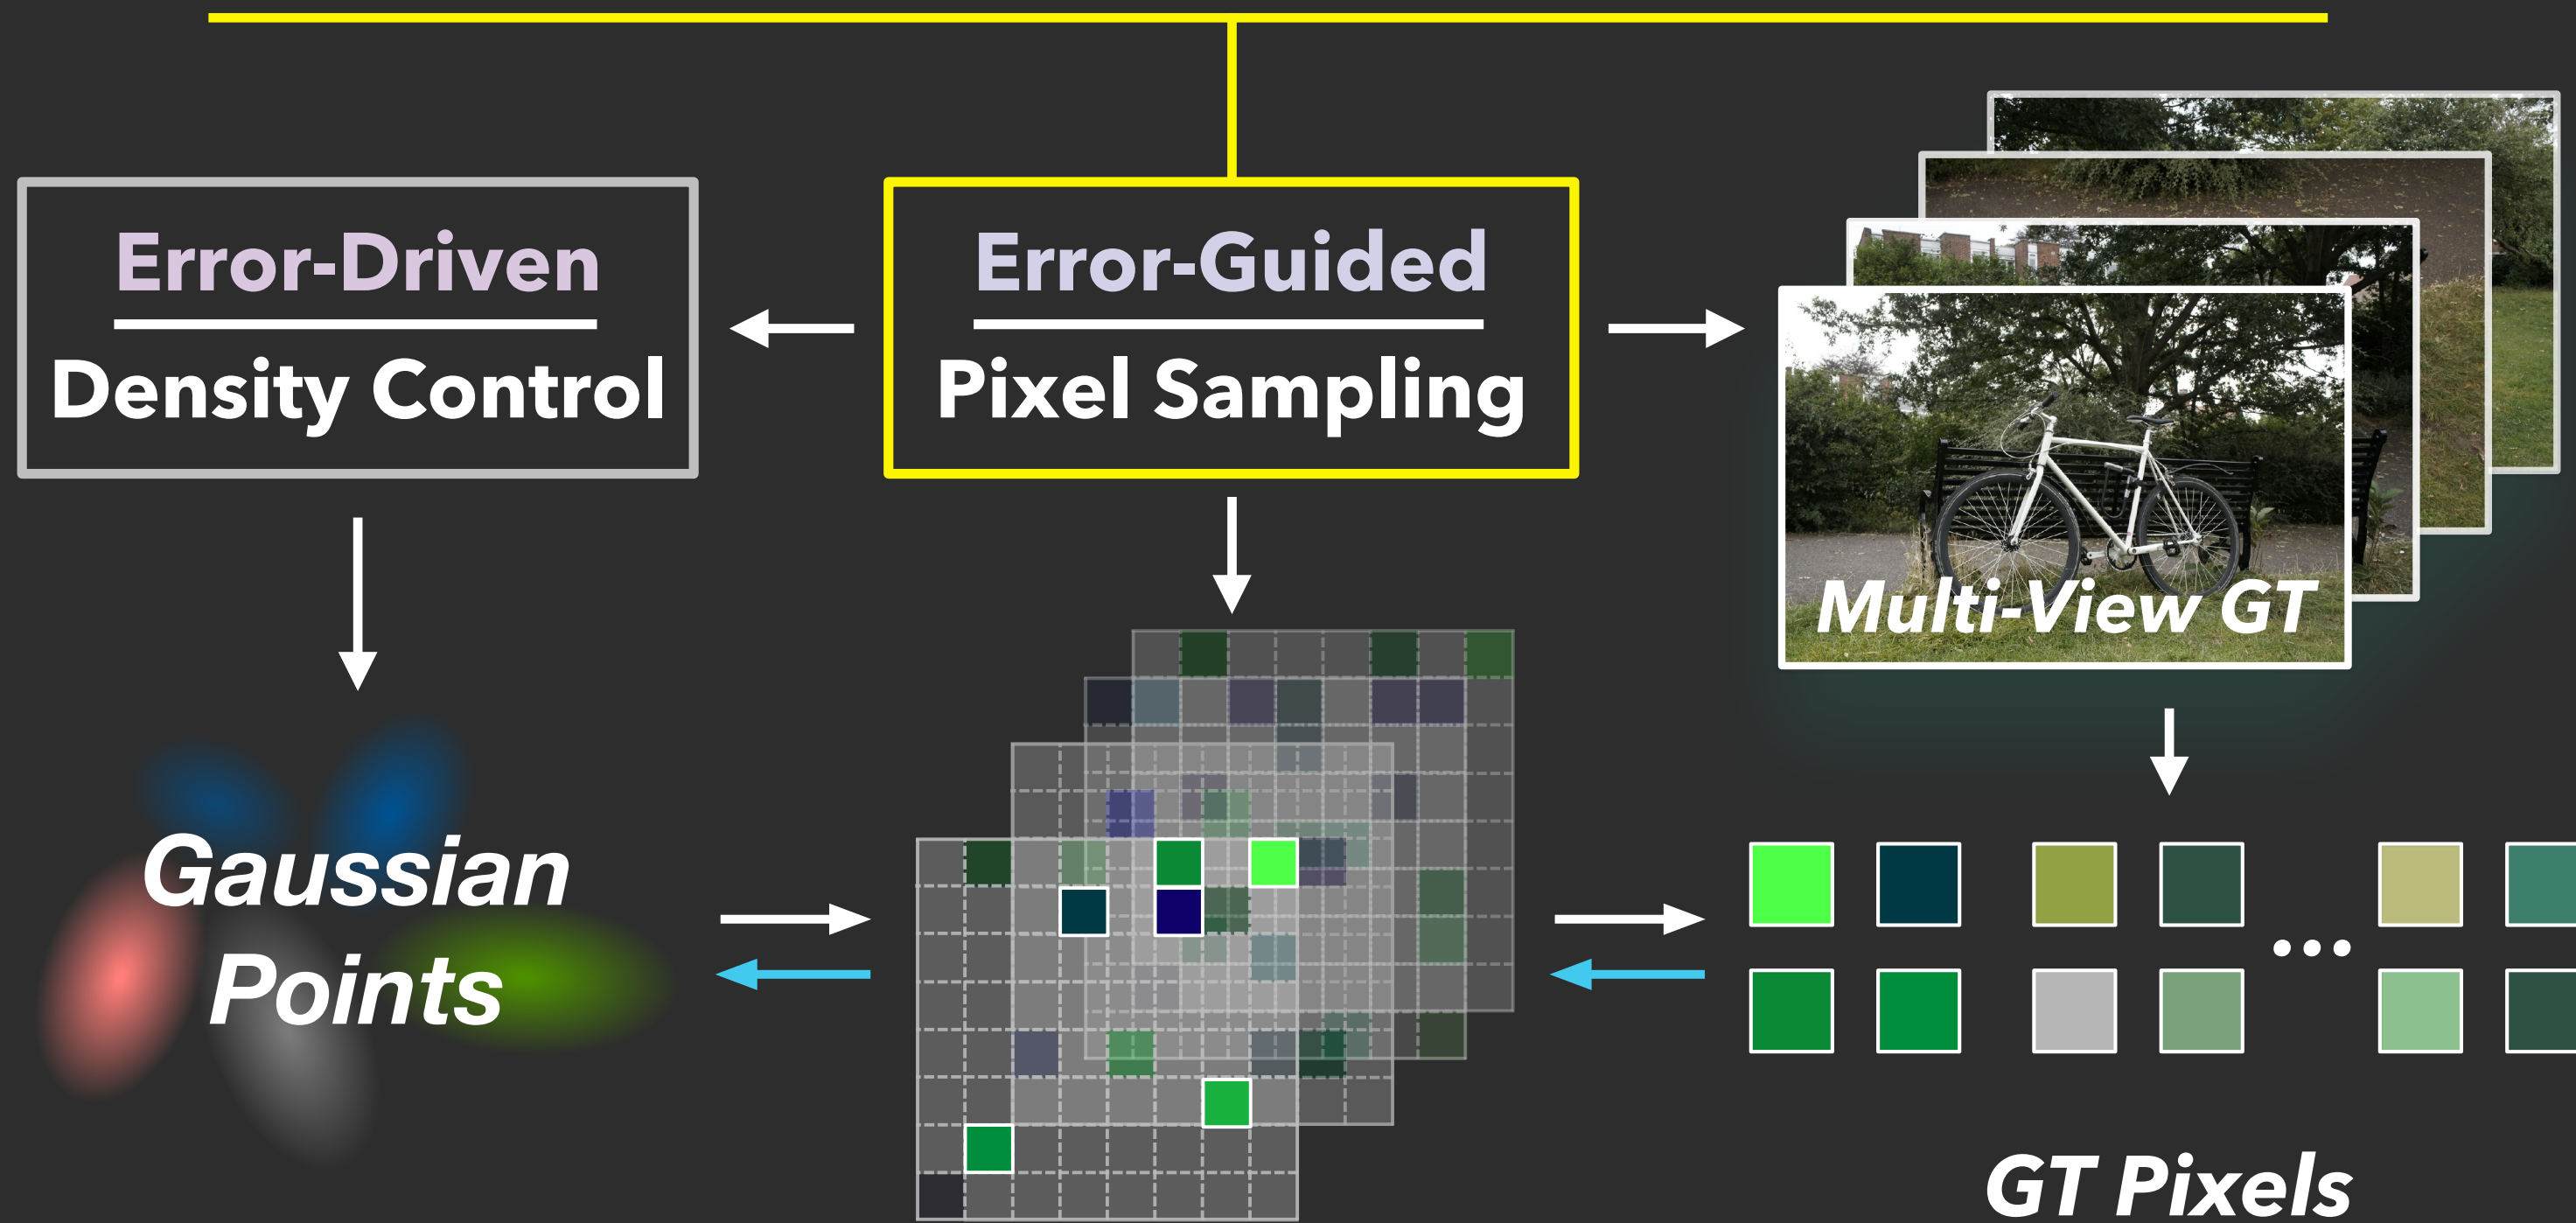
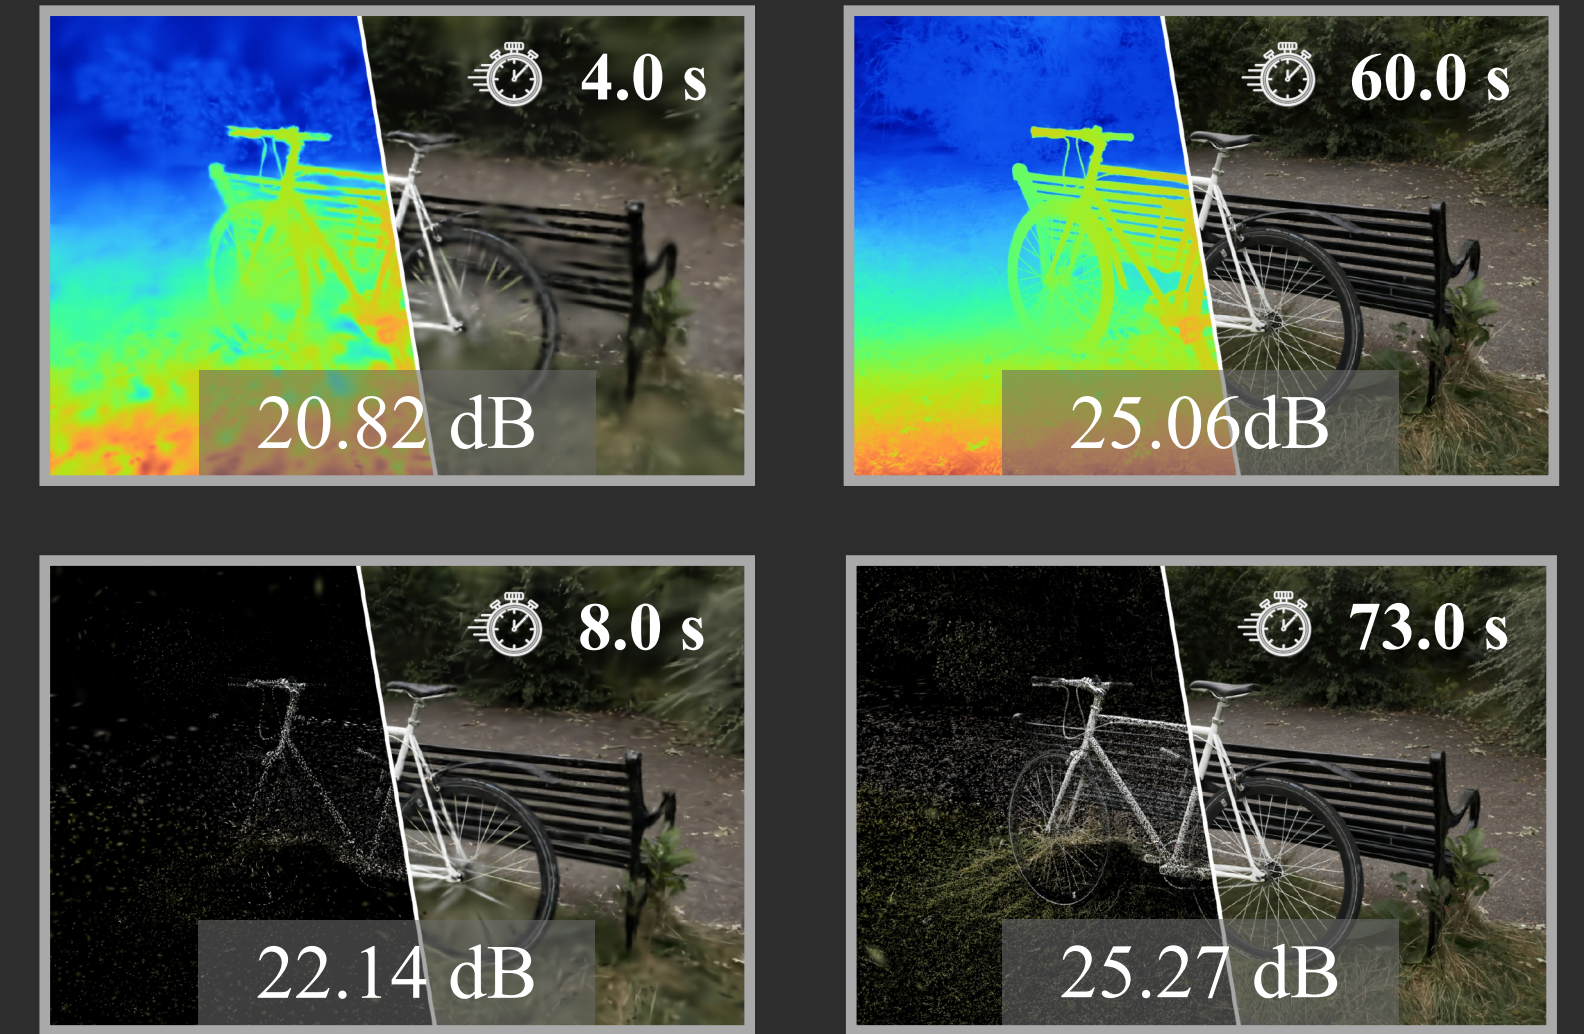



40.0 seonds ...



Most Pixels are **Rapidly Well Reconstructed**, yet They Still **Consume Large Computation**

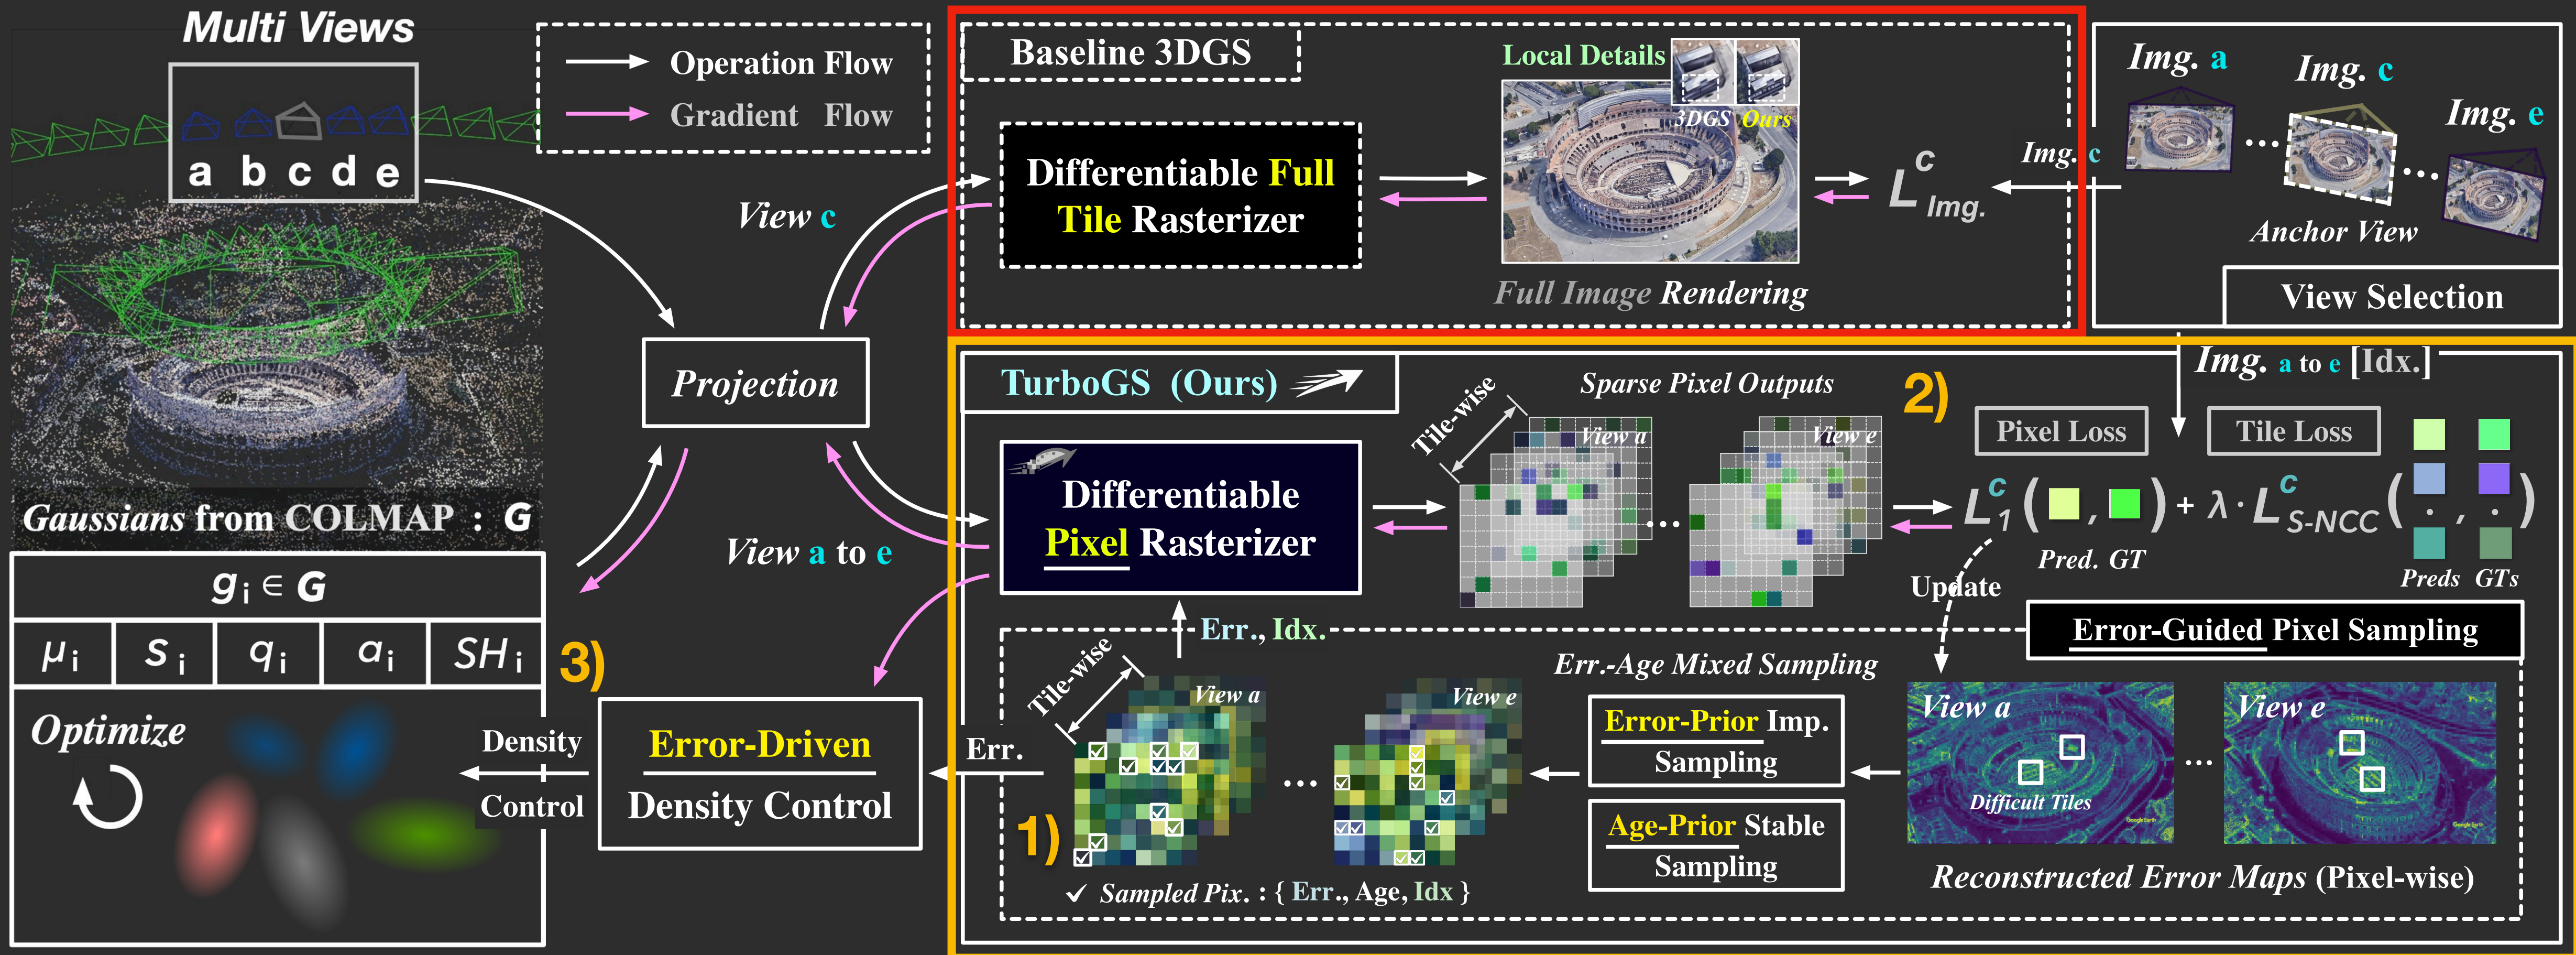
# Motivation : Compute Where It Matters



**PSNR - TIME on Bicycle Data**

# TurboGS : Framework Overview

3DGS : All Pixels and Gaussian Points Are Treated Equally

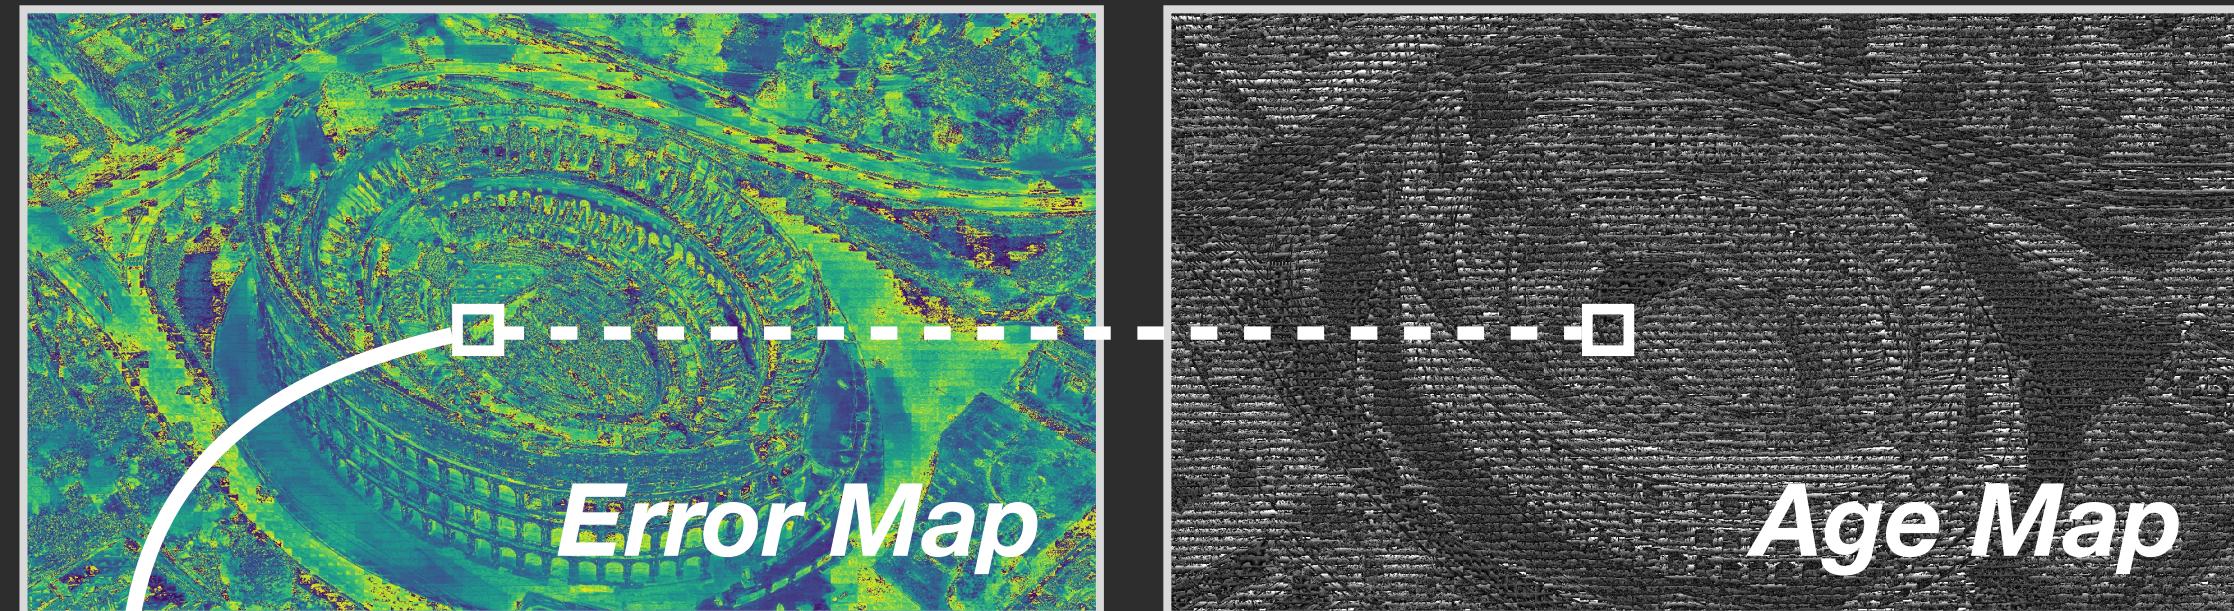


TurboGS Focuses on 1) Where to Optimize ? 2) How to Optimize ? 3) What to Optimize ?

# Adaptive **Error-Guided Sparse Pixel Sampling**

Persistent Informative Maps

 5.0 s



*Determine Where to Optimize ?*

**Pixel Loss as Error Signal :**  $e_{t,v}(\mathbf{p}) = \|\hat{\mathbf{I}}_{t,v}(\mathbf{p}) - \mathbf{I}_v(\mathbf{p})\|_1$

**EMA Update of Pixel Errors :**

$$\mathbf{E}_v(\mathbf{p}) \leftarrow (1 - \beta) \mathbf{E}_v(\mathbf{p}) + \beta e_{t,v}(\mathbf{p})$$

**Pixel Age Update :**  $\mathbf{A}_v(\mathbf{p}) \leftarrow \begin{cases} 0, & \mathbf{p} \text{ sampled,} \\ \mathbf{A}_v(\mathbf{p}) + 1, & \text{otherwise.} \end{cases}$

**Error-Prior Importance Sampling :**

$$\pi_E(\mathbf{p} \mid \tau, v) = \frac{\mathbf{E}_v(\mathbf{p})}{\sum_{\mathbf{q} \in \tau} \mathbf{E}_v(\mathbf{q}) + \epsilon}, \quad \mathbf{p} \in \tau,$$

**Age-Prior Stable Sampling :**

$$\mathcal{S}_{t,v}(\tau) = \text{TopK}_{\mathbf{p} \in \tau \setminus \mathcal{H}_{t,v}(\tau)}(\mathbf{A}_v(\mathbf{p}), N_{\text{stable}}),$$

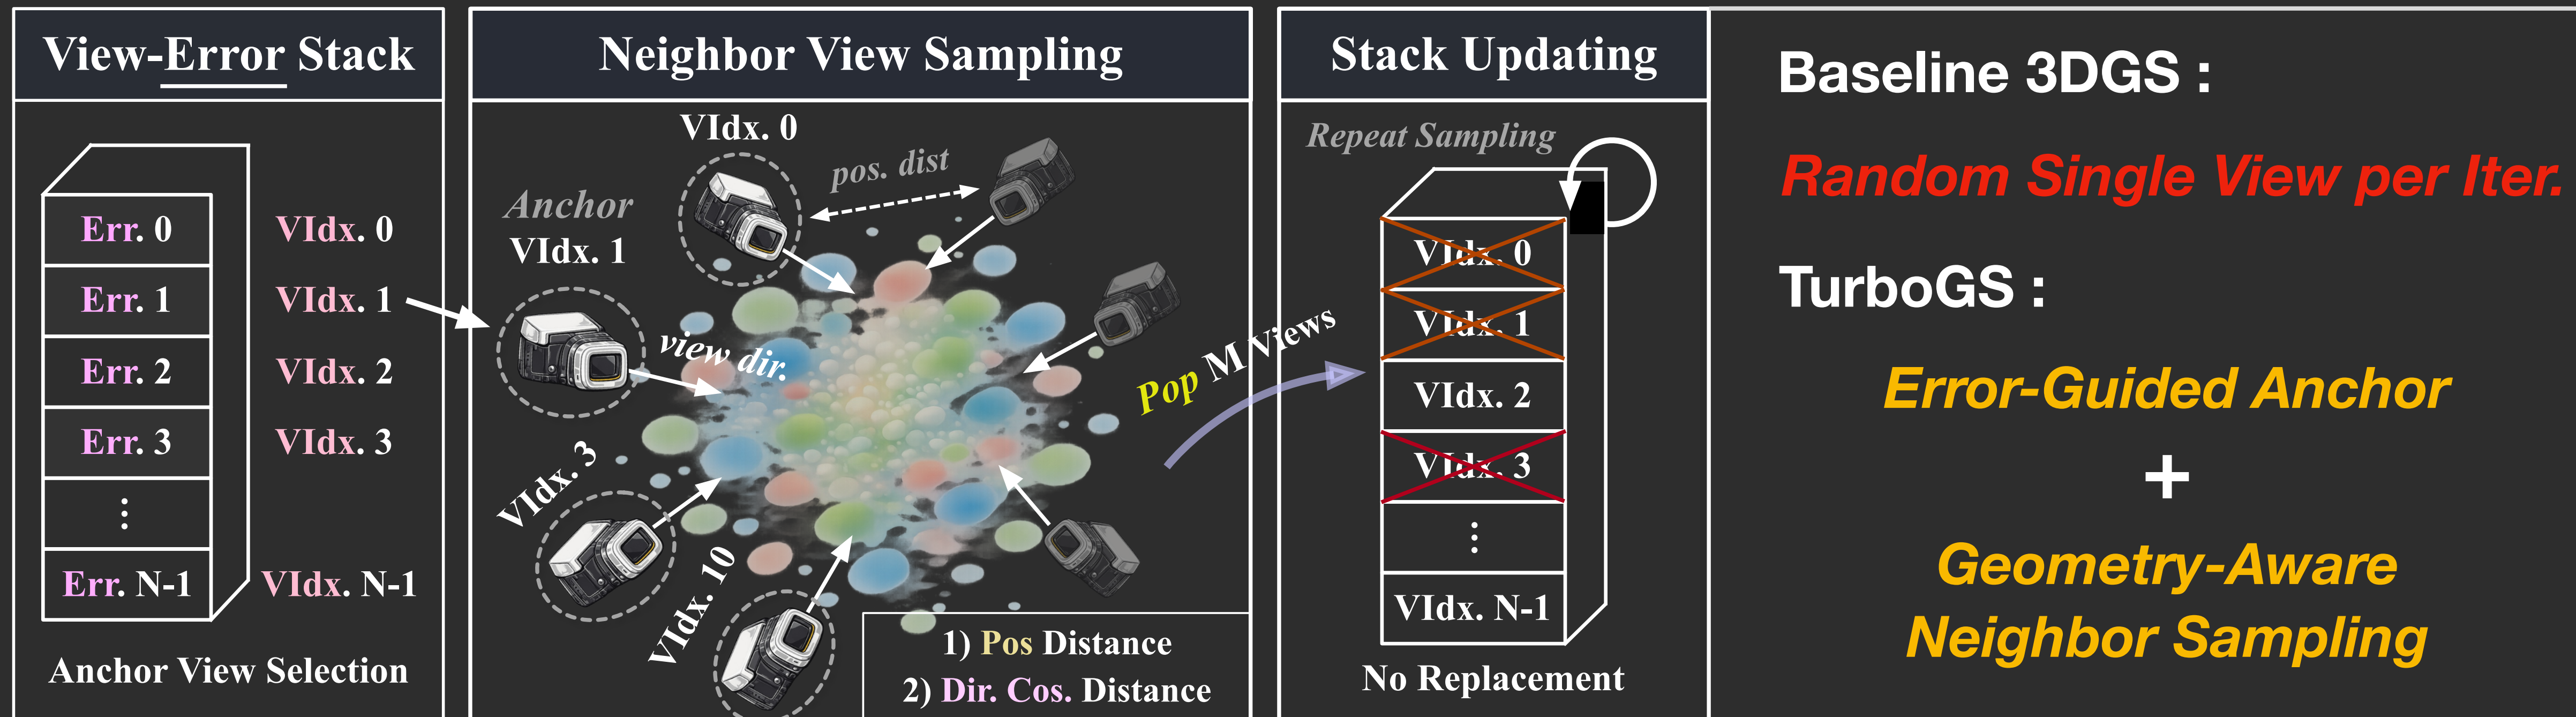
**Full Sampled Set :**  $\mathcal{P}_{t,v}(\tau) = \mathcal{H}_{t,v}(\tau) \cup \mathcal{S}_{t,v}(\tau)$



Focus on **Informative Pixels** with **Higher Recon. Error**

# Adaptive **Geometry-Aware** View Sampling

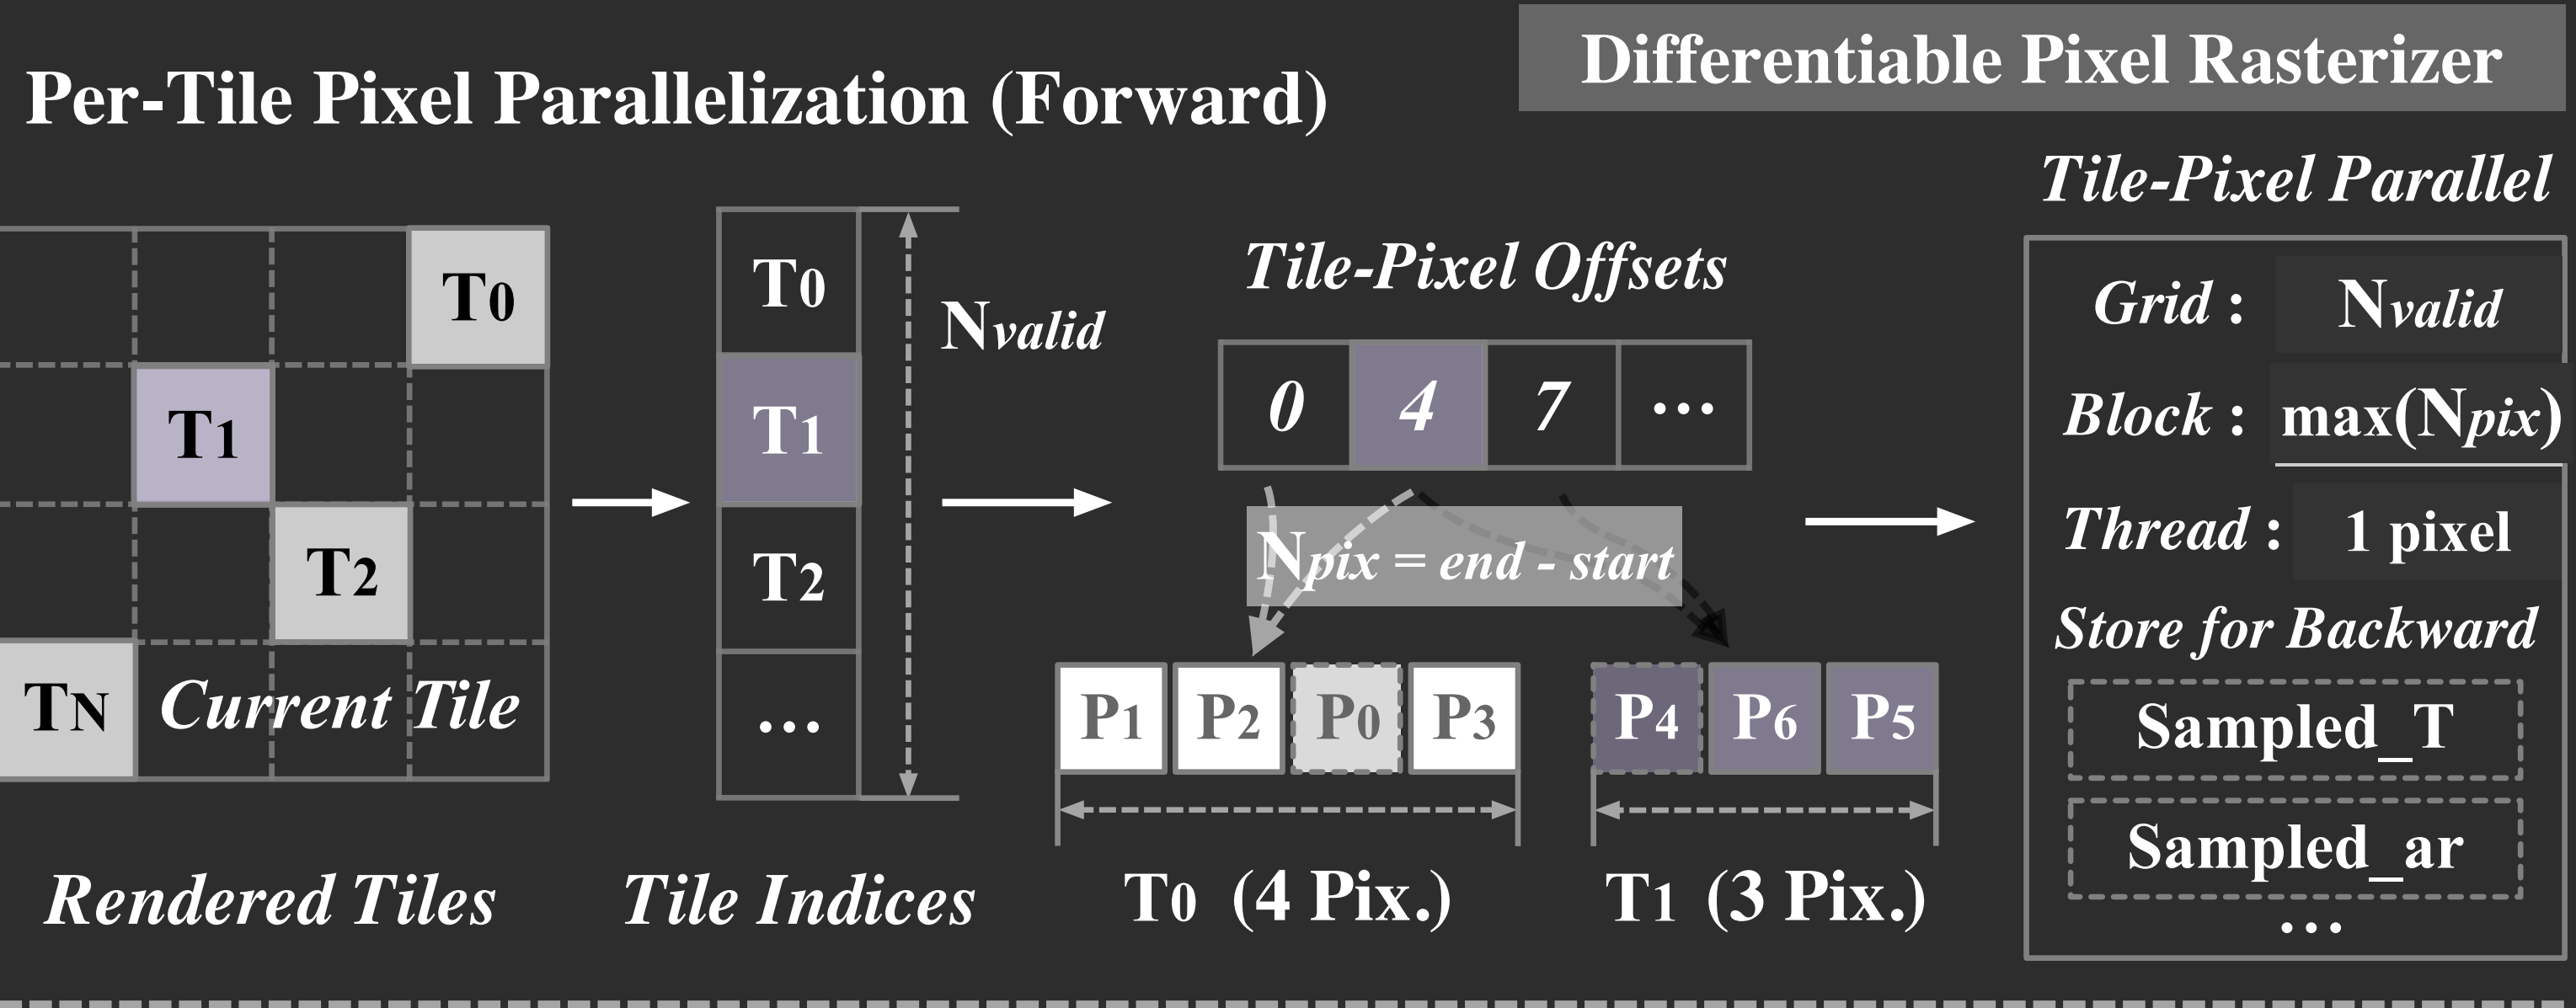
*Where to Optimize Globally ?* Multi-View Supervision with Local Geometric Perception



**Error Map Information Statistics :** 
$$\bar{E}_v = \frac{1}{|\Omega|} \sum_{p \in \Omega} E_v(p), \quad v^* = \arg \max_v \bar{E}_v, \quad (\text{Anchor View})$$

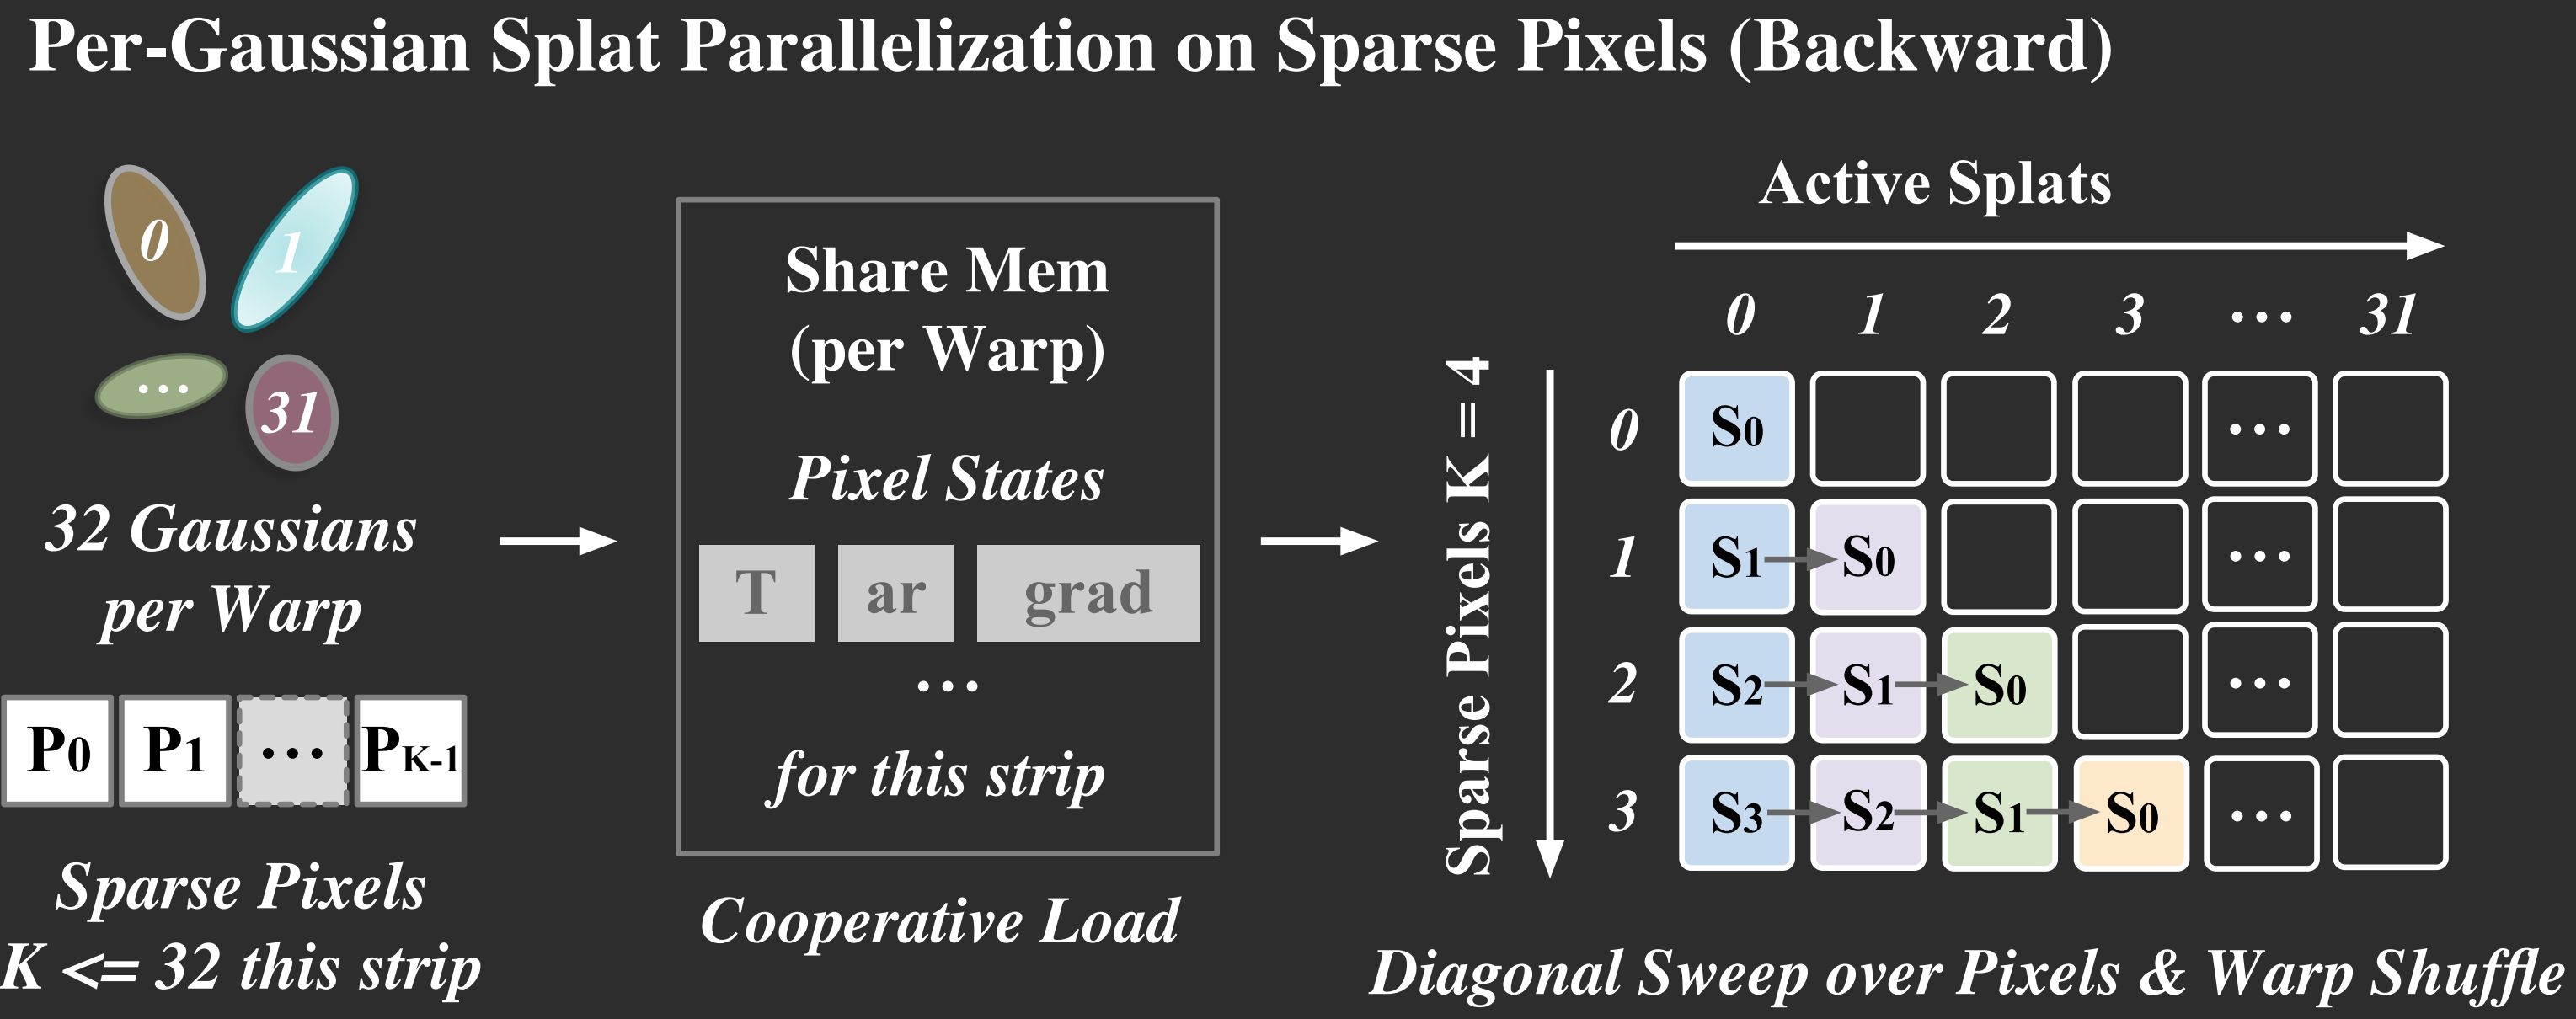
**Neighbor-View Searching :** 
$$\pi(u | v^*) \propto \exp\left(-\lambda_d \|c_u - c_{v^*}\|_2 - \lambda_\theta (1 - \langle d_u, d_{v^*} \rangle)\right), \quad (\text{K Sampled Views})$$

# Efficient Sparse Pixel Rasterization



*Forward Pixel Rasterization :*

- 1) Maintain **Valid Tile Indices**
- 2) Calc. **Tile-Pixel Offsets**
- 3) Query Pixel Indices in **Each Block**




*Backward :*

- 1) Sampled Sparse Pixels
- 2) **Per-Gaussian** Warp Assignment
- 3) **Warp-Shuffle** Accumulation
- 4) Gradient **Atomic Reduction**

# Making Sparse Optimization Effective

How to Optimize ?



**Tile-Wise Structure Loss :**

$$\mathcal{L}_{S-NCC}^{\tau} = 1 - \text{NCC}(\hat{\mathbf{I}}_{t,v}(\mathcal{P}_{t,v}(\tau)), \mathbf{I}_v(\mathcal{P}_{t,v}(\tau))) .$$

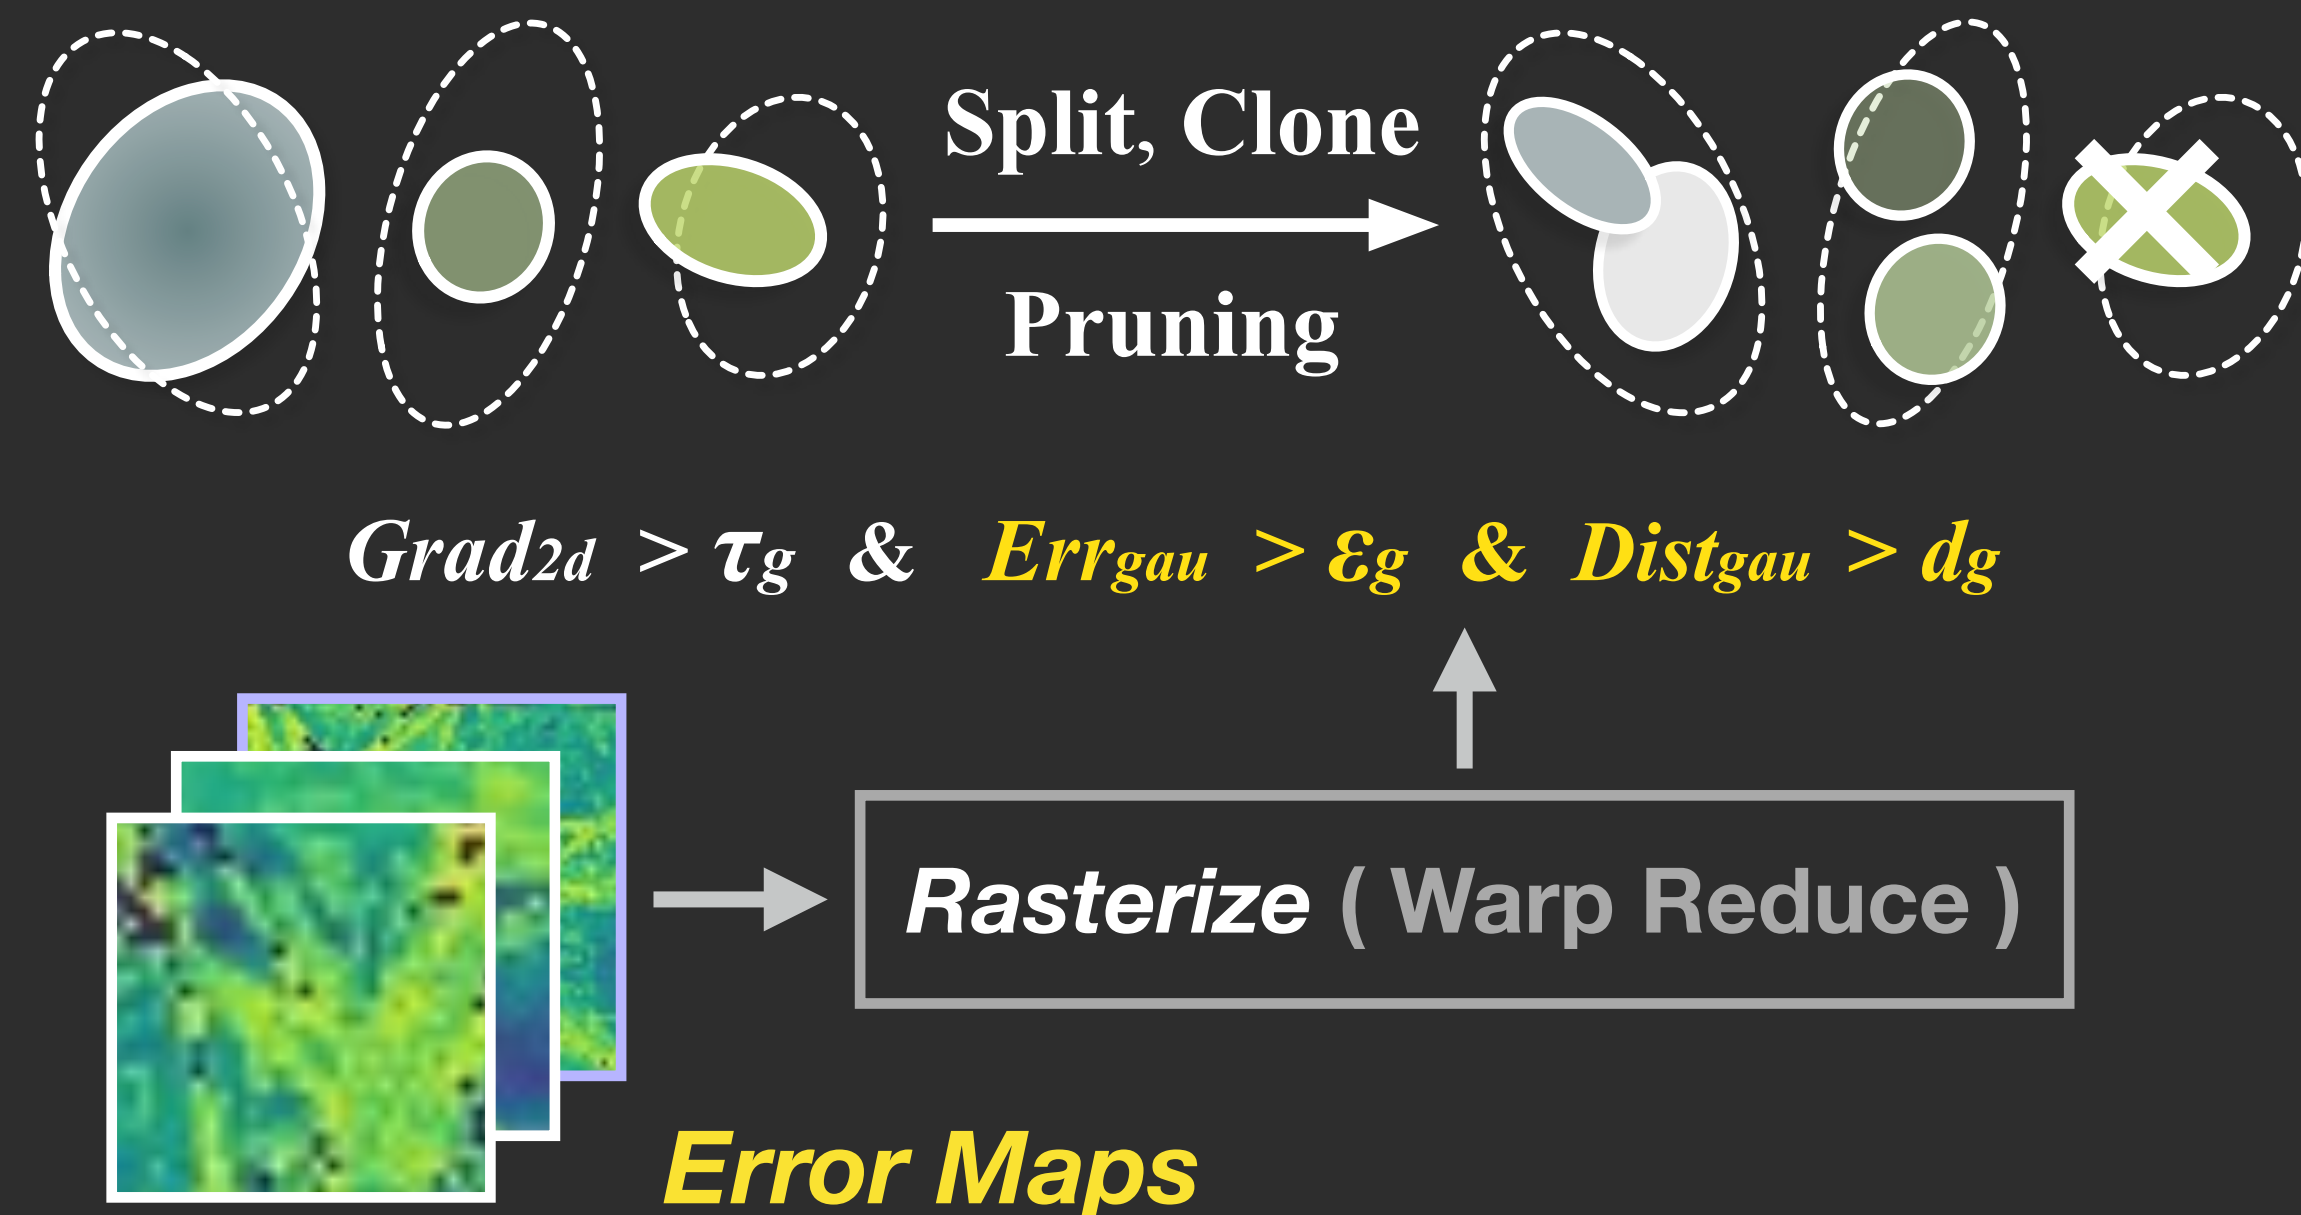
**Gradient-Balanced Weighting :**

$$\lambda_{ncc}^{\tau} = \text{clip}\left(\frac{\bar{g}_{\ell_1}^{\tau}}{\bar{g}_{S-NCC}^{\tau} + \epsilon}, 0, 1\right),$$

$$\bar{g}_k^{\tau} = \mathbb{E}_{\mathbf{p} \in \mathcal{P}_{t,v}(\tau)} \left[ \|\nabla_{\hat{\mathbf{I}}(\mathbf{p})} \mathcal{L}_k^{\tau}(\mathbf{p})\|_1 \right], \quad k \in \{\ell_1, S-NCC\} .$$

**Moment-LM Solver :**  $(\mathbf{J}_g^{\top} \mathbf{J}_g + \text{diag}(\sqrt{\mathbf{v}_g} + \epsilon)) \Delta_g = -\mathbf{m}_g$

What to Optimize ? **Compact Gaussians**



**Per-Gaussian Error & Weight Accumulation :**

$$\mathcal{E}_i^{(t)} = \sum_{v, \mathbf{p}} e_{t,v}(\mathbf{p}) w_{t,i}(\mathbf{p}), \quad \mathcal{W}_i^{(t)} = \sum_{v, \mathbf{p}} w_{t,i}(\mathbf{p}) .$$

**Normalization by Visibility :**

$$e_i^{(t)} = \mathcal{E}_i^{(t)} / (\mathcal{W}_i^{(t)} + \epsilon) \quad d_i^{(t)} = \mathcal{D}_i^{(t)} / (\mathcal{W}_i^{(t)} + \epsilon)$$


# Ablation Study : **WHY It Works**



3DGS (baseline)

+ Err-SPS. ... Err-DP. (Full)

Ground Truth



+ Err-SPS. & Geo-VS.

+ Sparse-NCC.


+ Err-DP. (Full)



Err-SPS. → RPS.

Geo-VS. → RVS.

Ours (Full)



S-NCC. → S-SSIM.

Sparse Adam (only)

Ours (Full)

## Ablation on *MipNeRF-360* Dataset

Ablated Method	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓
3DGS	891	<u>27.55</u>	0.816	<b>0.215</b>	2.73M
+ Err-SPS.	85	26.58	0.737	0.331	0.76M
+ Geo-VS.	89	26.73	0.743	0.321	1.05M
+ Sparse-NCC.	95	27.14	0.773	0.279	0.98M
+ Err-DP. (Full)	77	<b>27.57</b>	0.794	<u>0.256</u>	0.64M
Err-SPS. → RPS.	<b>75</b>	27.10	0.780	0.276	<u>0.56M</u>
Geo-VS. → RVS.	<u>76</u>	27.26	0.780	0.274	<b>0.47M</b>
Err-DP. → FastGS	91	<u>27.55</u>	<u>0.801</u>	0.258	0.60M
$\ell_1$ loss (only)	92	27.09	0.761	0.302	0.76M
S-NCC. → S-SSIM.	86	27.21	0.773	0.287	0.74M



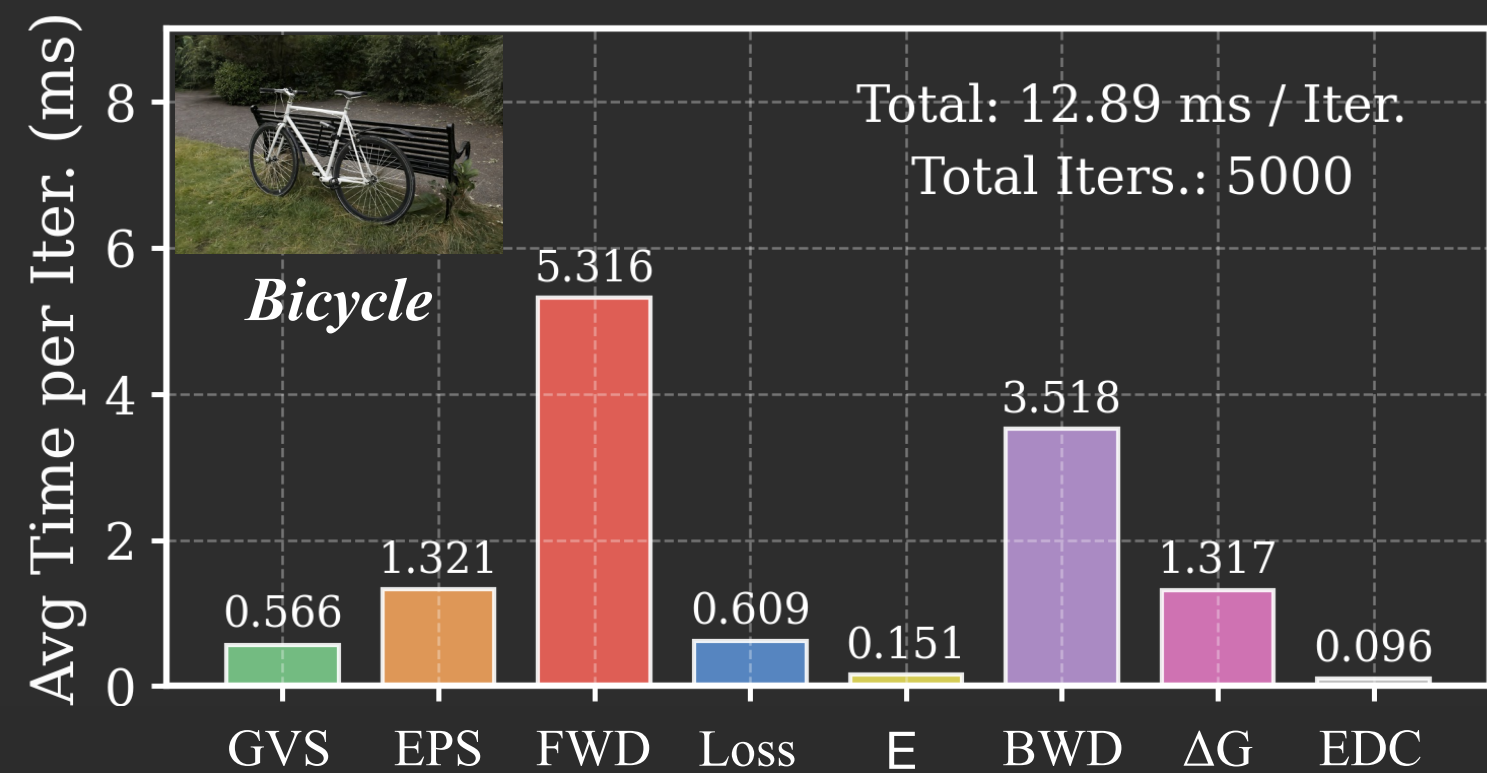
**Err-SPS.** *Dominates Major Acceleration*

**Geo-VS.** *Improved Multi-View Consistency*

**Sparse-NCC.** *Recovers Fine Structures*

**Err-DP.** *Adaptive Capacity Allocation*

# Analysis of TurboGS



**Completes training within ~80s in three representative datasets, while achieving competitive quality.**

Method	Mip-NeRF 360 ( <i>Barron et al., 2022</i> )						Tanks & Temples ( <i>Knapitsch et al., 2017</i> )						Deep Blending ( <i>Hedman et al., 2018</i> )					
	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑	Time↓	PSNR↑	SSIM↑	LPIPS↓	N <sub>GS</sub> ↓	FPS↑
3DGS	891	27.55	0.816	<b>0.215</b>	2.73M	199	583	23.84	<b>0.853</b>	<b>0.169</b>	1.58M	252	921	29.85	<b>0.907</b>	<b>0.238</b>	2.48M	208
Taming-3DGS	<b>148</b>	27.26	0.795	0.659	0.67M	<b>413</b>	<b>97</b>	23.77	0.836	0.211	0.32M	<b>612</b>	<b>100</b>	29.92	0.903	0.271	0.29M	601
Mini-Splatting	674	27.32	<b>0.822</b>	<b>0.217</b>	<b>0.49M</b>	279	479	23.24	0.836	0.202	<b>0.20M</b>	405	586	29.95	<b>0.907</b>	0.254	0.35M	364
Speedy-Splat	533	26.95	0.786	0.288	<b>0.32M</b>	<b>983</b>	292	23.41	0.820	0.239	<b>0.19M</b>	<b>1138</b>	497	29.55	0.903	0.268	<b>0.25M</b>	<b>1121</b>
3DGS-LM	622	27.41	0.814	0.220	3.35M	218	369	23.57	<b>0.844</b>	0.185	1.80M	285	572	29.58	<b>0.906</b>	<b>0.246</b>	2.89M	218
DashGaussian	169	<b>27.69</b>	<b>0.817</b>	0.220	2.24M	259	141	<b>24.04</b>	<b>0.851</b>	<b>0.181</b>	1.20M	334	112	<b>30.12</b>	<b>0.907</b>	<b>0.248</b>	1.95M	289
FastGS	<b>111</b>	27.51	0.805	0.261	<b>0.38M</b>	<b>1002</b>	<b>91</b>	<b>24.06</b>	0.842	0.209	<b>0.24M</b>	<b>1087</b>	<b>81</b>	30.06	0.905	0.267	<b>0.22M</b>	<b>1144</b>
ConeGS	778	<b>27.74</b>	<b>0.823</b>	<b>0.202</b>	0.87M	<b>427</b>	749	23.88	<b>0.853</b>	<b>0.160</b>	0.55M	564	805	<b>30.28</b>	<b>0.909</b>	<b>0.238</b>	0.52M	<b>655</b>
TurboGS (Ours)	<b>77</b>	27.57	0.794	0.256	0.64M	153	<b>73</b>	23.79	0.832	0.200	0.58M	170	<b>62</b>	30.00	0.900	0.277	0.49M	231
TurboGS-Big (Ours)	168	<b>27.99</b>	0.815	0.224	1.34M	134	146	<b>24.14</b>	0.841	<b>0.181</b>	0.91M	151	132	<b>30.26</b>	<b>0.909</b>	0.259	0.69M	216

# Optimization Progress Visualization



**OMMO-10 Part1**



TIME : 147s  
PSNR : 28.99dB  
SSIM : 0.845

**OMMO-10 Part2**



TIME : 147s  
PSNR : 28.99dB  
SSIM : 0.845

**Large-Scale Scene Results**

**1146 x 856**



**Rubble Full Data**



**Zoom in**



**Zoom in**

**4580 x 3430**



**Rubble Sub Region (Re-Run COLMAP)**

TIME 367.00s  
PSNR 25.88  
SSIM 0.775  
NGau 2.78M


## TurboGS (Ours)



TIME 100.00s  
PSNR 25.13  
SSIM 0.772  
NGau 0.66M

Especially with fine details

## 3DGS



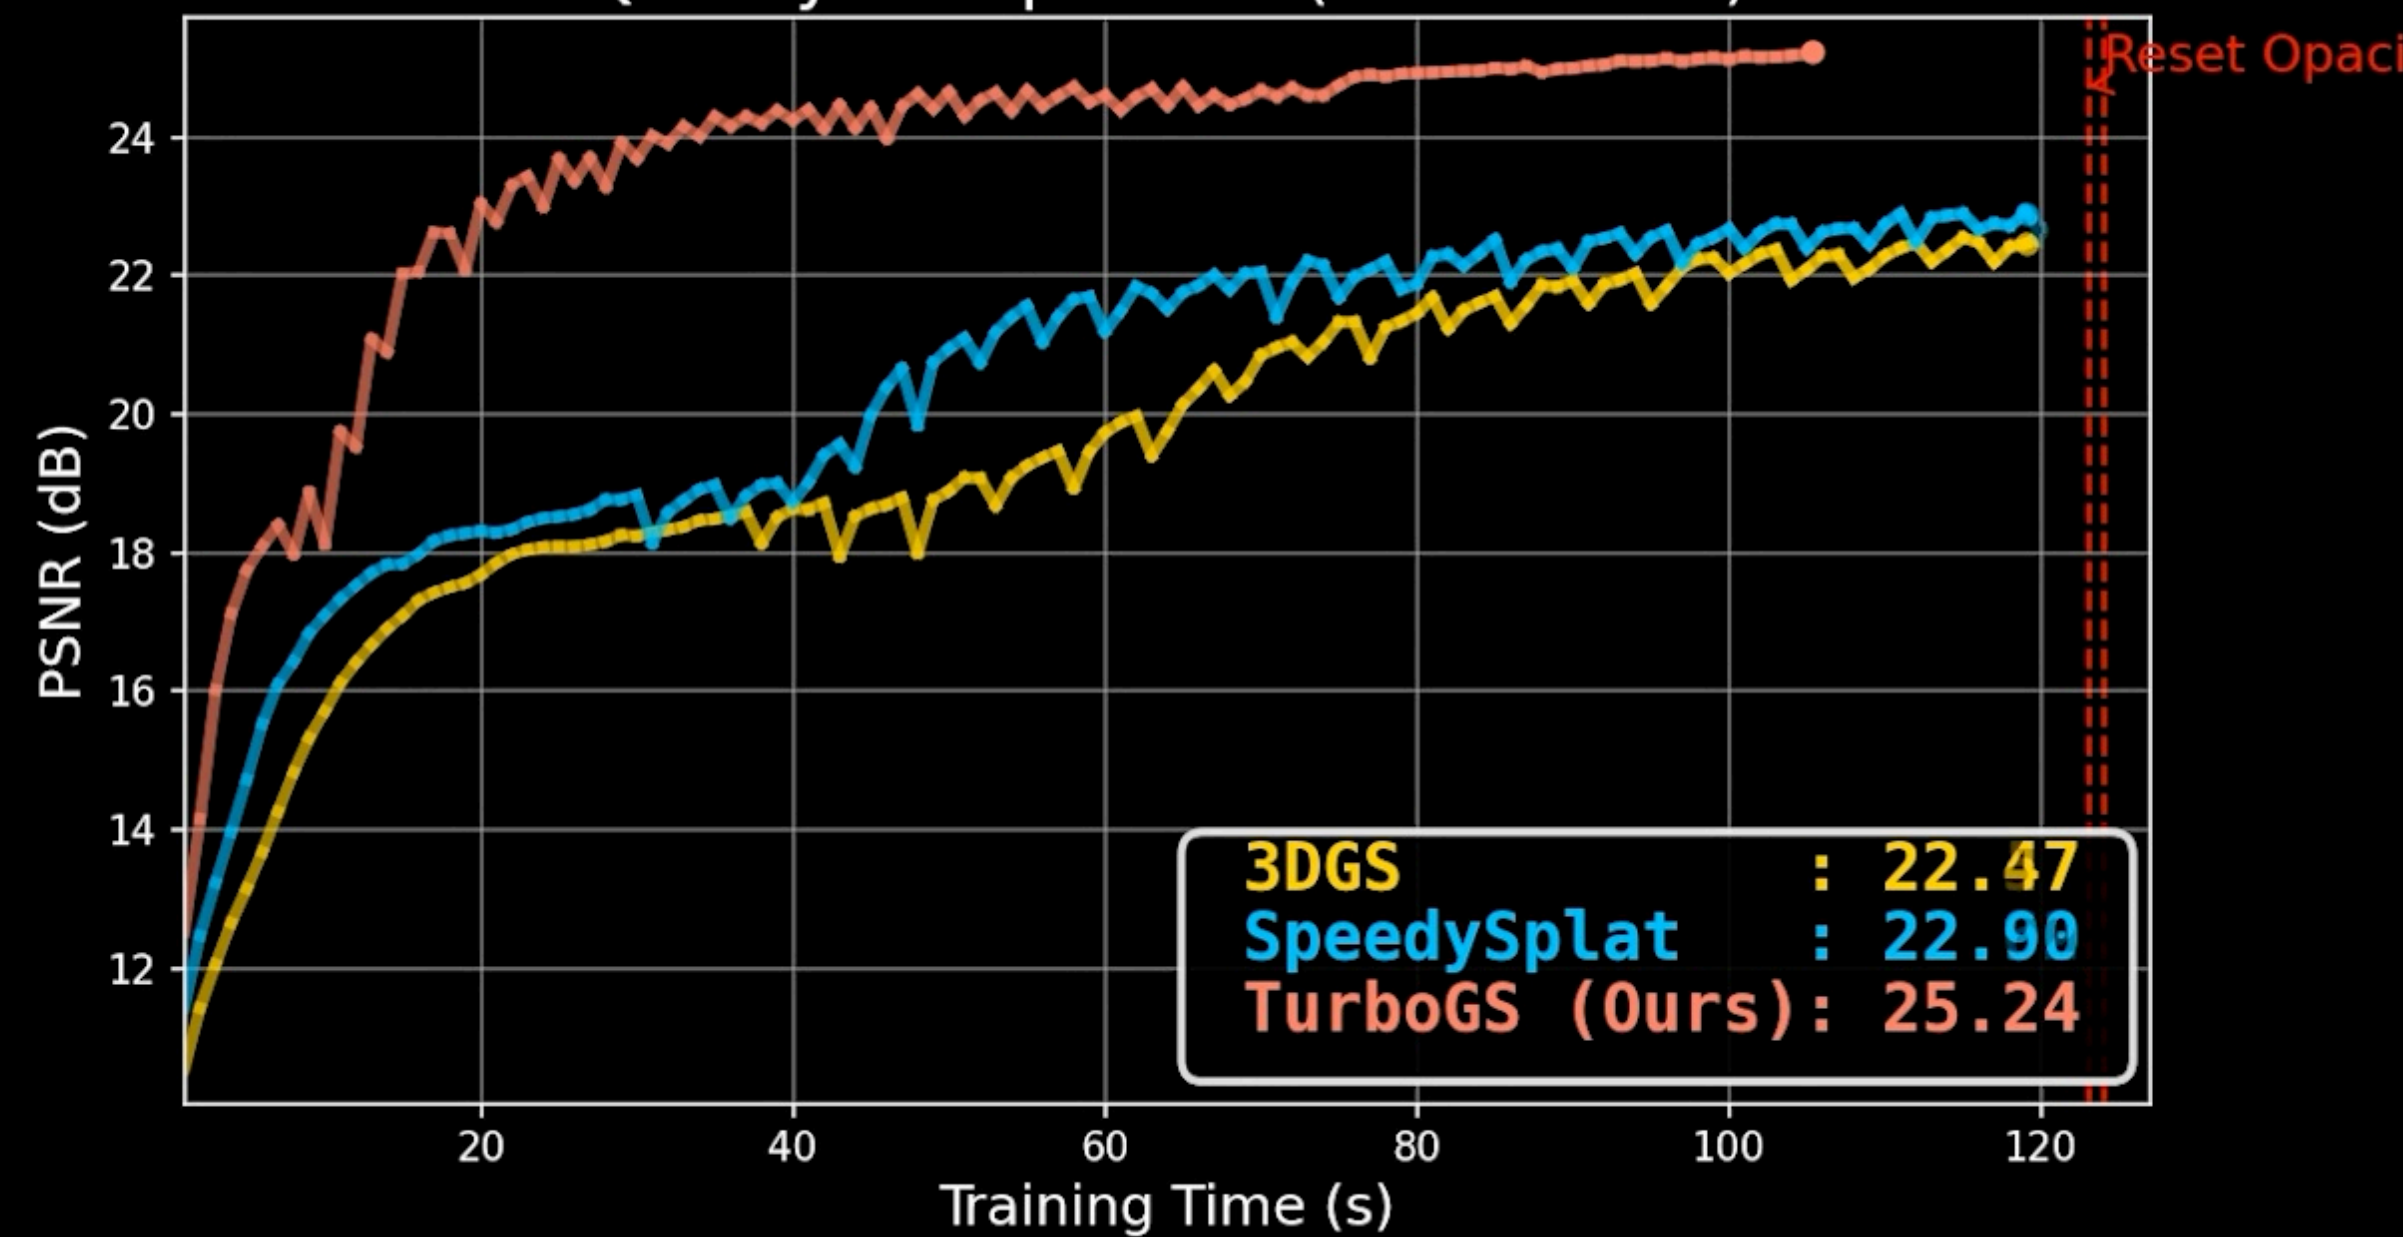
TIME 119.00s  
PSNR 22.47  
SSIM 0.698


## Speedy-Splat



TIME 119.00s  
PSNR 22.90  
SSIM 0.715

## Quality Comparison (TIME-PSNR)





# Key Takeaways

**Main Insight : Not All Pixels deserve equal computation.**

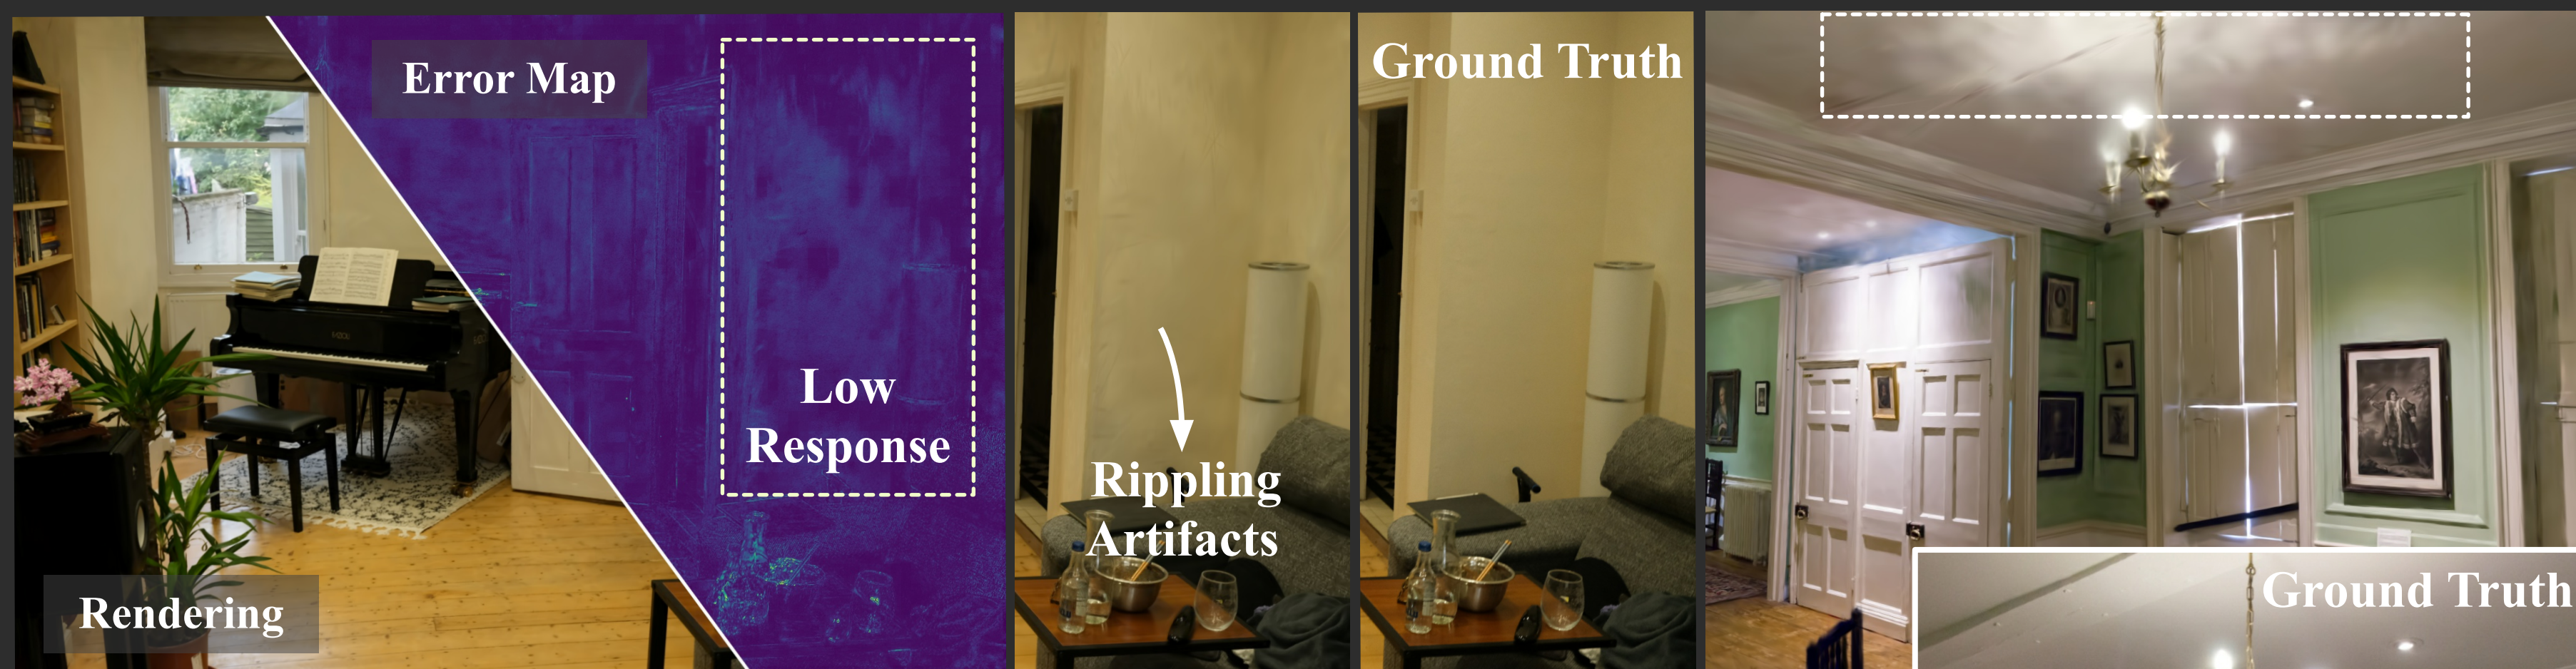
*Allocate computation adaptively across views, pixels and Gaussians.*

**Conclusion :**


- **Error-Guided Sparse Optimization achieves  $> 10x$  Speedup than Baseline 3DGS**
- **Compact Gaussian Representation with Error-Guided Density Control**

**Limitation :**

- **Low-Texture Regions : May Need Adaptive Dense Supervision;**
- **Extend to Dynamic Scene**



# THANKS YOU



**Project Page**

*<https://zhengdong.site/projects/TurboGS/>*