

---

# Incoherent Deformation, Not Capacity: Diagnosing and Mitigating Overfitting in Dynamic Gaussian Splatting

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Dynamic 3D Gaussian Splatting methods achieve strong training-view PSNR on  
2 monocular video but generalize poorly: on the D-NeRF benchmark we measure an  
3 average train–test PSNR gap of **6.18 dB**, rising to 11 dB on individual scenes. This  
4 paper reports two findings that together account for most of that gap.

5 **Finding 1 (the role of splitting).** A systematic ablation of the Adaptive Density  
6 Control pipeline—split, clone, prune, frequency, threshold, schedule—shows that  
7 *splitting* is responsible for over 80% of the gap: disabling split collapses the  
8 cloud from 44K to 3K Gaussians *and* the gap from 6.18 dB to 1.15 dB. Across all  
9 threshold-varying ablations, gap is log-linear in count ( $r=0.995$ , bootstrap 95%  
10 CI [0.99, 1.00]), which suggests a capacity-based explanation.

11 **Finding 2 (the role of deformation coherence).** We show that the capacity expla-  
12 nation is incomplete. A local-smoothness penalty on the per-Gaussian deformation  
13 field—*Elastic Energy Regularization* (EER)—reduces the gap by 40.8% *while*  
14 *growing the cloud by 85%*. Measuring the per-Gaussian strain directly on the  
15 trained checkpoints, EER reduces *mean* strain by 99.72% (*median* 99.80%) across  
16 all 8 scenes; on 8/8 scenes the *median* Gaussian under EER is less strained than  
17 the *1st-percentile* (best-behaved) Gaussian under baseline. Alongside EER, we  
18 evaluate two further regularizers: GAD, a loss-rate-aware densification threshold,  
19 and PTDrop, a jitter-weighted Gaussian dropout. The combination GAD+EER  
20 reduces the gap by 48%; adding PTDrop and a soft growth cap reaches 57%.  
21 We confirm that coherence generalizes to (a) a different deformation architecture  
22 (Deformable-3DGS, +40.6% gap reduction at re-tuned  $\lambda$ ), and (b) real monocular  
23 video (4 HyperNeRF scenes, reducing the mean PSNR gap by 14.9% at the *same*  
24  $\lambda$  as D-NeRF, with near-zero quality cost). The overfitting in dynamic 3DGS is  
25 driven by incoherent deformation, not parameter count.

## 26 1 Introduction

27 3D Gaussian Splatting [7] is a strong explicit representation for radiance fields. Its extension to  
28 monocular dynamic video—4DGS [17], Deformable-3DGS [19], SC-GS [6]—couples a canonical  
29 Gaussian cloud with a learned deformation field. These methods typically reach > 40 dB PSNR on  
30 training views but drop by 6–11 dB on held-out test views. The MonoDyGauBench benchmark [8]  
31 recently flagged the Adaptive Density Control (ADC) mechanism as a brittleness source but did not  
32 isolate which sub-operation is at fault.

33 This paper is a diagnostic study with two specific claims.

34 **Claim 1: Splitting dominates the overfitting.** We ablate every ADC sub-operation (split, clone,  
35 prune, frequency, threshold, schedule) on D-NeRF. *Split* accounts for  $> 80\%$  of the gap. The majority  
36 of densification happens in the first half of training, so early stopping at iteration 7,500 (ablation A6,  
37 Sec. 5) has negligible effect.

38 **Claim 2: Capacity reduction alone does not address the root cause.** Across all ablations the  
39 gap is log-linear in final Gaussian count—a capacity-based explanation. We show it is incomplete.  
40 A local smoothness penalty on the per-Gaussian deformation (EER) reduces the gap by 40.8%  
41 while increasing the cloud by 85%. Measured directly on trained checkpoints, EER reduces mean  
42 per-Gaussian  $k$ -NN strain by **99.72%** on 8/8 scenes. The dominant factor is therefore *incoherent*  
43 *deformation*, not parameter count.

#### 44 **Scope and contributions.**

- 45 • A systematic ADC-sub-operation ablation on D-NeRF (8 scenes, 260+ training runs) that  
46 isolates split as the dominant sub-operation.
- 47 • Empirical evidence against the capacity-only hypothesis: EER reduces the gap despite  
48 increasing the cloud size, contradicting the log-linear count–gap relation, while directly  
49 verifying on trained checkpoints that the deformation field becomes locally smooth.
- 50 • Three regularization methods, each adding one hyperparameter to the training pipeline:  
51 **EER** (a  $k$ -NN smoothness penalty on the deformation field, Sec. 6), **GAD** (a loss-rate-aware  
52 densification threshold, Sec. 7.1), and **PTDrop** (a jitter-weighted Gaussian dropout, Sec. 7.2).  
53 The combination GAD+EER reduces the gap by 48.2%.
- 54 • Confirmation that the coherence finding generalizes across architectures (Deformable-3DGS)  
55 and datasets (HyperNeRF).

56 We also tested additional regularizers (SGD, STSR, ChromReg, OEM, TCMask); none produced  
57 more than 10% gap reduction.

## 58 **2 Related Work**

59 **Dynamic Gaussian Splatting.** 4DGS [17] uses a HexPlane deformation field; Deformable-  
60 3DGS [19] an MLP; SC-GS [6] sparse control points. All inherit ADC from static 3DGS [7]  
61 with minimal changes. Per-Gaussian embedding smoothness in E-D3DGS [1] is, in retrospect, a  
62 close architectural relative of our coherence finding: we quantify the effect and show it dominates  
63 capacity-based control.

64 **Overfitting and ADC analysis.** MonoDyGauBench [8] reports severe train–test gaps on dynamic  
65 scenes and flags ADC as a likely cause. Static-3DGS work on densification includes AbsGS [20],  
66 Mini-Splatting [3], Revising Densification [14], Pixel-GS [21], Taming 3DGS [9], SteepGS [16],  
67 GDAGS [22], EDC [2], and PUP-3DGS [5]; Grubert et al. [4] introduce an ascending threshold  
68 schedule (prior, VISAPP Feb. 2025). DropGaussian [10] and DropoutGS [18] use dropout for sparse-  
69 view static 3DGS. None of these target monocular dynamic overfitting or the split-vs.-coherence  
70 question.

## 71 **3 Background**

72 **3DGS.** A scene is a set of  $K$  anisotropic Gaussians with position, covariance (scale + rotation),  
73 opacity, and spherical harmonic color. Rendering is differentiable  $\alpha$ -blending over sorted, splatted  
74 Gaussians.

75 **Adaptive Density Control (ADC).** Every  $k = 100$  iterations from iter 500 to 15K: every Gaussian  
76 with view-space gradient  $\bar{g} > \tau_0$  is either *split* (if its scale exceeds a size threshold—one Gaussian  
77 becomes two smaller copies) or *cloned* (otherwise—a shifted duplicate is added). Low-opacity  
78 Gaussians are pruned. The default  $\tau_0 = 2e-4$ .

79 **4DGS.** 4DGS [17] adds a HexPlane deformation network predicting per-Gaussian posi-  
80 tion/rotation/scale offsets conditioned on time. Training has a coarse stage (3K iters, no deformation)  
81 and a fine stage (20K iters total, deformation and ADC active).

## 82 4 Experimental Setup

83 **Dataset.** D-NeRF benchmark [13]: 8 synthetic monocular dynamic scenes, 50 training and 20  
84 held-out test views each, at novel camera poses and timesteps.

85 **Baseline.** Public 4DGS [17] code, default hyperparameters. Our reproduction matches the original  
86 paper within  $\pm 0.33$  dB per scene.

### 87 Metrics.

- 88 • **Train–Test PSNR gap** ( $\text{PSNR}_{\text{train}} - \text{PSNR}_{\text{test}}$ ), primary overfitting diagnostic.
- 89 • **Test PSNR / SSIM / LPIPS**, novel-view quality.
- 90 • **Final Gaussian count**  $K$ , at the last iteration.

91 **Statistical protocol.** 8 scenes give  $n = 8$  paired replicates. All comparisons are paired  $t$ -tests with  
92 Cohen’s  $d$ ; bootstrap CIs are reported where useful. Effect-size estimates at this sample size carry  
93 wide CIs; the qualitative claims we make are robust to sample size, but small per-method differences  
94 should not be over-interpreted.

95 **Hardware.** Single NVIDIA RTX 3070 (8 GB),  $\sim 13$  min/scene.

### 96 Ablations.

- 97 **A1** No densification (disable all ADC).
- 98 **A2** No split, keep clone + prune.
- 99 **A3** No clone, keep split + prune.
- 100 **A4** No prune, keep split + clone.
- 101 **A5** Half densification frequency (every 200 iters).
- 102 **A6** Early stop at iter 7,500 (half the window).
- 103 **A7**  $2\times$  threshold ( $\tau = 4e-4$ ).
- 104 **A8**  $0.5\times$  threshold ( $\tau = 1e-4$ ).

## 105 5 Finding 1: Splitting Dominates

106 Table 1 and Fig. 1 show the full ablation. Three observations:

107 **Split is the dominant sub-operation.** A2 (no split) holds the cloud at 3K Gaussians and the gap at  
108 1.15 dB, nearly identical to A1 (no densification at all: 2K, 1.13 dB). The remaining clone+prune  
109 pipeline does not materially overfit on its own.

110 **Pruning has negligible effect.** A4 changes the final count by 2% and the gap by 1%. Pruning  
111 neither prevents nor exacerbates overfitting in our experiments.

112 **Schedule modifications have negligible effect.** A5 (half frequency) trims count by 30% but the  
113 gap by only 9%. A6 (early stop at iter 7,500) changes count by 10% and the gap by 2%. This is  
114 because densification is strongly front-loaded: 84–89% of cloud growth occurs before iter 7,500. Any  
115 effective mitigation must modulate densification *from the start*, not truncate it at the end.

116 **Threshold magnitude has the largest schedule effect.** A7 ( $\tau \times 2$ ) reduces count by 63% and the  
117 gap by 26% at a modest  $-0.76$  dB quality cost. A8 ( $\tau \times 0.5$ ) nearly triples count (126K) and grows  
118 the gap by 23%. Threshold *magnitude* matters far more than frequency or timing.

Table 1: Ablation results on D-NeRF (8 scenes, mean  $\pm$  std). Gap = Train PSNR – Test PSNR.  $\bar{K}$  = mean Gaussian count. Cohen’s  $d$  measures effect size on gap vs. baseline (paired  $t$ -test).

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Gap $\downarrow$	$\bar{K}$	Cohen’s $d$	$p$
Baseline	34.11 $\pm$ 5.17	0.85 $\pm$ 0.01	0.03 $\pm$ 0.01	6.18 $\pm$ 3.12	44,516	–	–
<i>Operation ablations</i>							
A1: No Densify	23.83 $\pm$ 2.61	0.78 $\pm$ 0.03	0.11 $\pm$ 0.04	<b>1.13 <math>\pm</math> 0.33</b>	2,000	2.43	<0.001
A2: No Split	24.18 $\pm$ 2.57	0.79 $\pm$ 0.03	0.10 $\pm$ 0.03	1.15 $\pm$ 0.43	3,073	2.42	<0.001
A3: No Clone	31.59 $\pm$ 4.28	0.83 $\pm$ 0.02	0.05 $\pm$ 0.02	3.06 $\pm$ 2.04	7,378	1.26	<0.01
<i>Schedule ablations</i>							
A4: No Prune	34.14 $\pm$ 5.14	0.85 $\pm$ 0.01	0.03 $\pm$ 0.01	6.25 $\pm$ 3.15	45,520	–0.02	0.16
A5: Half Freq	33.92 $\pm$ 5.04	0.84 $\pm$ 0.01	0.03 $\pm$ 0.02	5.62 $\pm$ 2.84	31,354	0.20	<0.05
A6: Early Stop	34.08 $\pm$ 5.18	0.84 $\pm$ 0.01	0.03 $\pm$ 0.01	6.04 $\pm$ 3.08	40,293	0.05	<0.05
<i>Threshold ablations</i>							
A7: 2 $\times$ Thresh	33.36 $\pm$ 4.99	0.84 $\pm$ 0.02	0.03 $\pm$ 0.02	4.60 $\pm$ 2.40	16,322	0.61	<0.01
A8: 0.5 $\times$ Thresh	34.29 $\pm$ 5.34	0.85 $\pm$ 0.01	0.02 $\pm$ 0.01	7.59 $\pm$ 3.66	126,493	–0.47	<0.01

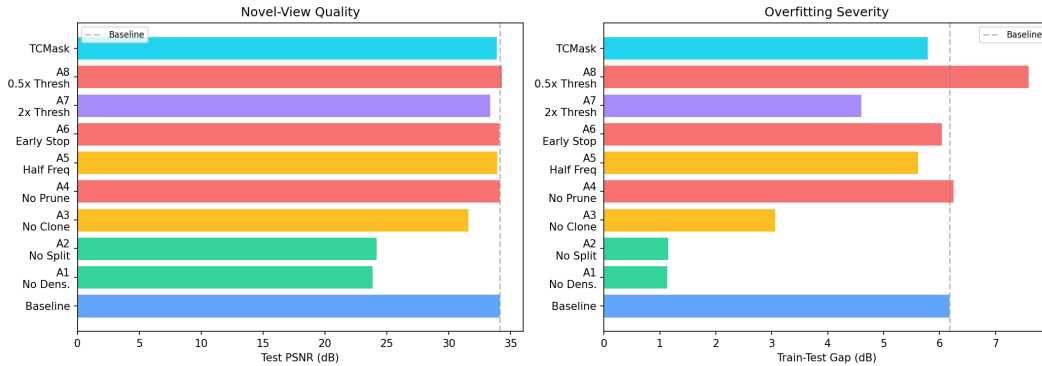


Figure 1: ADC sub-operation ablation. Left: test PSNR (quality). Right: train–test gap (overfitting). Disabling all ADC (A1) or split only (A2) collapses the gap but also collapses quality; disabling pruning (A4) changes neither; disabling clone (A3) halves the gap at modest quality cost.

119 **The count–gap correlation.** Plotting final count vs. gap across 41 non-EER configurations we  
 120 pilot-tested—nine ablations (A1–A8 plus baseline) plus 32 non-EER training runs spanning capacity,  
 121 stochastic, and other regularizers—we find a tight log-linear relation: Pearson  $r(\log K, \text{gap}) = 0.987$   
 122 across all 41, and  $r = 0.995$  (bootstrap 95% CI [0.993, 1.000]) on the 9-condition ablation subset  
 123 (Fig. 2, gray points). At face value this says: *more Gaussians, more overfitting*. Section 6 shows that  
 124 reading is wrong.

## 125 6 Finding 2: Coherence Beats Capacity

126 If the count–gap correlation told the full story, then any intervention that reduces the cloud should  
 127 reduce the gap, and any intervention that grows the cloud should grow it. We now show a simple,  
 128 local deformation-smoothness penalty breaks both halves of that prediction.

129 **Elastic Energy Regularization (EER).** Over a random minibatch  $\mathcal{G}$  of sampled Gaussians, for  
 130 each Gaussian  $i$  we find its  $k$  nearest neighbors  $\mathcal{N}_k(i)$  in *canonical* space (pre-deformation) and add

$$\mathcal{L}_{\text{EER}} = \frac{\lambda_{\text{EER}}}{|\mathcal{G}| \cdot k} \sum_{i \in \mathcal{G}} \sum_{j \in \mathcal{N}_k(i)} \frac{\|\mathbf{u}(x_i, t) - \mathbf{u}(x_j, t)\|^2}{\|x_i - x_j\|^2 + \varepsilon} \quad (1)$$

131 to the training loss, where  $\mathbf{u}(x, t)$  is the deformation at time  $t$ . Intuitively this is a locally-linear  
 132 smoothness prior: it penalizes how much a Gaussian’s motion differs from its canonical neighbors’,  
 133 normalized by their canonical distance so that the penalty measures *relative* deformation rather than  
 134 absolute motion.

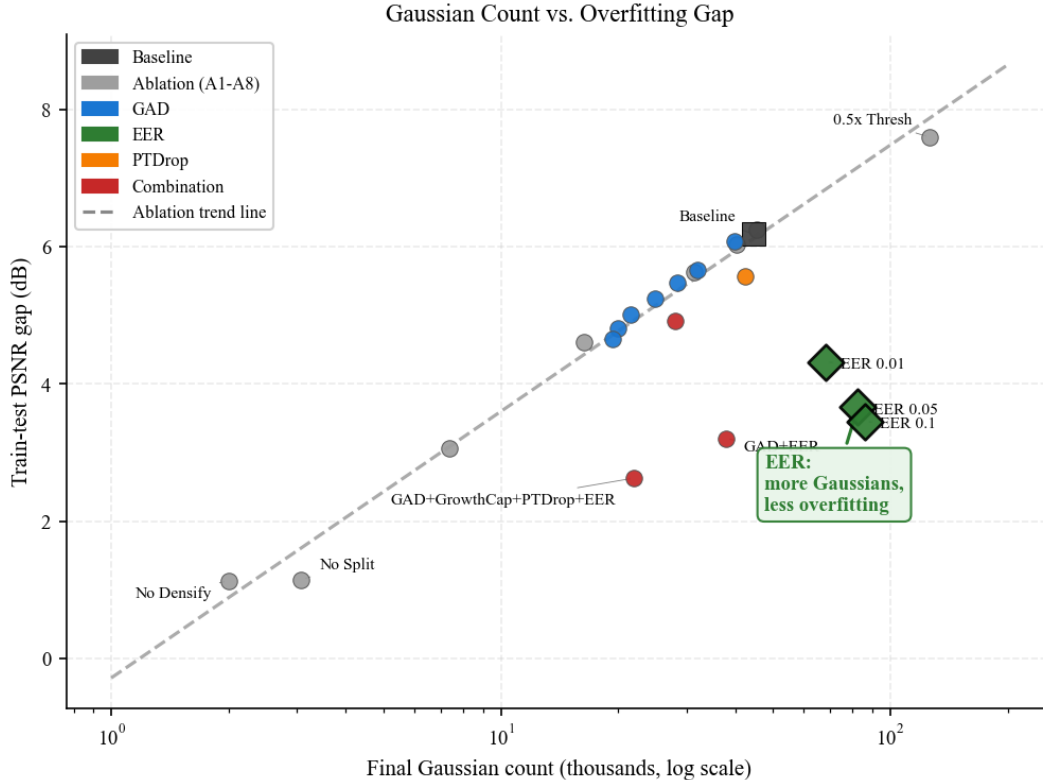


Figure 2: Count–gap relation. Ablations (gray) are log-linear, which would suggest a capacity story. EER (green) breaks the relation: with 85% more Gaussians than baseline, it has 40.8% less overfitting. The mechanism is not capacity.

135 **Implementation details.** We use  $k=8$  nearest neighbors, computed in *canonical* (pre-deformation)  
 136 position space. The neighbor graph is rebuilt every 500 iterations as canonical positions shift slowly  
 137 during training. EER is active only during the fine stage: the loss weight  $\lambda_{\text{EER}}$  increases from 0  
 138 to its target value following a half-cosine curve over iterations 3K–10K, avoiding interference with  
 139 early structural learning. For computational efficiency, we sample 2,048 Gaussians per training step  
 140 rather than penalizing all Gaussians.  $\lambda_{\text{EER}}$  is the only hyperparameter; on D-NeRF we evaluate  
 141  $\lambda_{\text{EER}} \in \{0.01, 0.05, 0.1\}$ .

142 **EER reduces overfitting while increasing Gaussian count.** At  $\lambda=0.05$ , EER reaches a **40.8%** gap  
 143 reduction (6.18→3.66 dB) with a  $-0.49$  dB test-PSNR cost. The final cloud has **82,498 Gaussians**—  
 144 **85% more** than baseline’s 44,516. At  $\lambda=0.1$  the gap drops further to 3.45 dB (44.2%, 86K  
 145 Gaussians). Even  $\lambda=0.01$ , our mildest setting, beats every non-EER method in Sec. 8 (30.1%  
 146 reduction,  $-0.19$  dB). Fig. 2 places these points well off the log-linear line. The capacity hypothesis  
 147 does not explain EER.

148 **Direct evidence of the mechanism.** We load each trained checkpoint, query the deformation  
 149 network at four timesteps, and measure per-Gaussian  $k$ -NN strain  $\bar{\epsilon}_i = \frac{1}{k} \sum_j \|\mathbf{u}_i - \mathbf{u}_j\|^2 / \|x_i - x_j\|^2$   
 150 (the same quantity Eq. 1 penalizes). Across all 8 D-NeRF scenes EER reduces *mean* strain by **99.72%**  
 151 (median 99.80%, min 99.58%, max 99.90%; Table 2). The median Gaussian under EER is less  
 152 strained than the 1st-percentile Gaussian under baseline on 8/8 scenes, and EER’s p99 strain is below  
 153 baseline’s median strain on 8/8 scenes with ratios ranging from  $4.6\times$  to  $26\times$ . Fig. 3 visualizes this:  
 154 the deformation field goes from chaotic (baseline, heavy-tailed strain distribution) to locally smooth  
 155 (EER, tight distribution at  $<10^{-2}$ ).

156 **Why the cloud grows.** A Gaussian that would have wandered across frames to memorize per-  
 157 training-view residuals can no longer do so under EER. The reconstruction loss therefore stays higher

Table 2: Per-Gaussian  $k$ -NN strain (mean over 8 D-NeRF scenes, 4 timesteps).

Statistic	Baseline $\bar{\epsilon}$	EER $\bar{\epsilon}$	Reduction
Mean (n=8)	3.54	0.00922	<b>99.72%</b>
Median (n=8)	1.16	0.00207	<b>99.80%</b>




Figure 3: Deformation field on Lego. **Left**: canonical cloud colored by per-Gaussian displacement magnitude (baseline above, EER below). **Middle**: subsampled quiver of  $\mathbf{u}(x, t = 0.5)$ . **Right**: distribution of per-Gaussian  $k$ -NN strain. Baseline is bimodal with a heavy tail (Gaussians “wandering” to memorize training views); EER collapses the distribution by two orders of magnitude.

158 in those regions, the view-space gradient there stays larger, and ADC’s split criterion fires more often.  
 159 EER substitutes population size for per-Gaussian deformation freedom: it trades deformation entropy  
 160 for additional Gaussians. Capacity-controlling EER via our adaptive threshold (GAD, Sec. 7) absorbs  
 161 this effect and reduces both the count and the gap.

## 162 7 Additional Regularizers: GAD and PTDrop

163 In addition to EER (Sec. 6, the primary regularizer motivated by Finding 2), we evaluate two further  
 164 regularizers: GAD, which adapts the ADC densification threshold based on loss-improvement rate,  
 165 and PTDrop, a jitter-weighted variant of DropGaussian for dynamic scenes. Each requires no  
 166 architectural changes and adds one hyperparameter.

### 167 7.1 GAD: A Loss-Rate-Aware Threshold

168 The ADC threshold  $\tau_0$  decides which Gaussians qualify for densification. The default is a global  
 169 constant, which makes sense early in training (everything is under-represented) but over-densifies late  
 170 in training (when loss has plateaued and each new Gaussian is mostly memorizing residual noise).  
 171 We adapt:

$$\tau_{\text{GAD}}(t) = \tau_0 \cdot \left( 1 + \lambda \cdot \frac{K(t)}{N \cdot \Delta \ell_{\text{ema}}(t)} \right), \quad (2)$$

172 where  $K(t)$  is the current cloud size,  $N$  is the number of training pixels, and  $\Delta \ell_{\text{ema}}(t)$  is an  
 173 exponential moving average (EMA) of the per-iteration loss improvement ( $\rho = 0.99$ ). The threshold  
 174 rises when the cloud is large *and* loss has stopped improving—the regime where new Gaussians  
 175 primarily memorize noise.

176 **Motivation.** The form of Eq. 2 is motivated by the Bayesian Information Criterion (BIC) [15]. In  
 177 BIC, adding a mixture component incurs a penalty proportional to  $K \cdot \log N/N$ . In our setting,  $K$   
 178 (cloud size) and  $N$  (training pixels) are the relevant quantities. The ratio  $K/(N \cdot \Delta \ell_{\text{ema}})$  captures the  
 179 same intuition: as  $K$  grows relative to the data and loss improvement slows, the threshold rises. We  
 180 absorb the constant factors ( $p = 59$  parameters per Gaussian,  $\log N$ ) into the single tunable  $\lambda$ , which  
 181 controls how aggressively the threshold adapts. We present Eq. 2 as a principled *heuristic* rather than  
 182 a formal derivation. A sweep over  $\lambda \in [0.1, 20]$  (Fig. 4) gives smooth, monotonic trade-offs.




Figure 4: GAD  $\lambda$  sweep. Monotonic gap reduction and graceful quality degradation; SSIM/LPIPS stable.

183 **Dimensional note.**  $K$  and  $N$  are dimensionless, so  $K/(N \cdot \Delta\ell)$  carries units of  $[\text{loss}]^{-1}$  and  $\lambda$   
 184 inherits them. The optimal  $\lambda$  therefore depends on the loss magnitude used by the host codebase.  
 185 This matters for the cross-architecture transfer in Sec. 9.

## 186 7.2 PTDrop: Jitter-Weighted Stochastic Dropout

187 PTDrop is a modified version of DropGaussian [10], adapted for dynamic scenes. DropGaussian  
 188 applies uniform per-Gaussian dropout in the sparse-view static 3DGS setting. PTDrop modifies this  
 189 in two ways:

- 190 1. **Cosine-scheduled drop rate.** During each forward pass, a random subset of Gaussians  
 191 is excluded from rendering. The global drop probability starts at 0 and increases to 0.3  
 192 over iterations 5K–12K following a cosine schedule, allowing training to stabilize before  
 193 regularization begins.
- 194 2. **Jitter-weighted selection.** Each Gaussian’s drop probability is scaled by the variance of  
 195 its deformation trajectory across timesteps. Gaussians with inconsistent motion—the same  
 196 incoherent behavior that Finding 2 identifies as the overfitting mechanism—are dropped  
 197 more frequently, while Gaussians with smooth trajectories are largely retained.

198 PTDrop adds roughly 10 percentage points of gap reduction on top of GAD+EER and composes  
 199 cleanly with both.

## 200 8 Results

201 Table 3 shows the quality–gap trade-off. Four observations.

202 **GAD and EER address complementary failure modes.** GAD alone reduces the gap by 11.3%,  
 203 EER alone by 40.8%. Combined, **GAD+EER reaches 48.2% gap reduction** at 38K Gaussians  
 204 with a  $-0.86$  dB test-PSNR cost. GAD reduces unnecessary Gaussian creation; EER constrains  
 205 the remaining Gaussians’ deformation freedom. The combination is super-additive, confirming that  
 206 capacity control and coherence regularization target different mechanisms.

207 **Additional gains with PTDrop and GrowthCap.** Adding PTDrop (Sec. 7.2) and a soft cloud-size  
 208 cap (*GrowthCap*: as Gaussian count  $K$  approaches a preset maximum  $K_{\max} = 15\text{K}$ , a sigmoid func-  
 209 tion progressively suppresses densification), the full combination **GAD+GrowthCap+PTDrop+EER**  
 210 reaches gap **2.63 dB (57.4% reduction)** at 22K Gaussians,  $-1.70$  dB test PSNR; six of eight scenes  
 211 reach  $< 3$  dB gap. The additional 9 percentage points over GAD+EER come at a larger quality cost.

212 **Without EER, the gap remains largely open.** GAD+GrowthCap+PTDrop (everything except  
 213 EER) reaches only 20.2% gap reduction. Among all non-coherence interventions we tested, no  
 214 combination exceeds 25% reduction—EER alone exceeds this by a wide margin.

215 **Negative results.** We also evaluated four additional regularizers (spectral gating, temporal smooth-  
 216 ness, SH-coefficient penalties, opacity entropy). None achieves more than 10% standalone gap  
 217 reduction.

Table 3: Main results: GAD (capacity control), EER (deformation coherence), their combination, and the full stack. 8 D-NeRF scenes;  $\lambda$  is the only per-method hyperparameter. Gap = train–test PSNR gap;  $\bar{K}$  is the mean final Gaussian count;  $d$  is Cohen’s  $d$  vs. baseline.

Method	Gap↓	$\Delta$ Gap	PSNR↑	$\bar{K}$	$d$
Baseline	6.18	—	34.11	44.5K	—
<i>Capacity control</i>					
GAD ( $\lambda=1$ )	5.48	−11.3%	33.92	28.3K	1.22
GAD ( $\lambda=5$ )	5.02	−18.9%	33.68	21.5K	1.39
<i>Deformation coherence</i>					
EER ( $\lambda=0.01$ )	4.32	−30.1%	33.92	68.2K	1.08
EER ( $\lambda=0.05$ )	3.66	−40.8%	33.62	82.5K	1.35
EER ( $\lambda=0.1$ )	3.45	−44.2%	33.34	86.0K	1.41
<i>Stochastic complement</i>					
PTDrop	5.57	−9.9%	33.93	42.3K	1.66
<i>Combinations</i>					
GAD + PTDrop	4.92	−20.3%	33.67	28.0K	1.63
<b>GAD + EER</b>	<b>3.20</b>	<b>−48.2%</b>	<b>33.25</b>	<b>37.8K</b>	<b>1.58</b>
<b>Full (GAD+GrowthCap+PTDrop+EER)</b>	<b>2.63</b>	<b>−57.4%</b>	<b>32.41</b>	<b>22.0K</b>	<b>1.72</b>

Table 4: Deformable-3DGS,  $\lambda$  sweep on Lego (top) plus cross-scene replication at the same loss-scale-adjusted  $\lambda=0.30$  on Hellwarrior (bottom). Monotonic dose-response on Lego; on Hellwarrior the same  $\lambda$  also improves both gap and test PSNR, confirming the sweep is not Lego-specific. Sign convention: positive reduction = EER improved the gap. (T-Rex at  $\lambda=0.30$  exceeded our 8 GB GPU budget during the  $k$ -NN step and is omitted.)

Scene	$\lambda$	Gap (dB)	Test PSNR	$\Delta$ Test	Red.
Lego	0 (baseline)	13.15	25.23	—	—
Lego	0.05	13.56	25.21	−0.02	−3.1%
Lego	0.15	10.23	25.33	+0.10	+22.3%
Lego	0.30	8.26	25.34	+0.11	+37.2%
Lego	<b>0.60</b>	<b>7.82</b>	<b>25.39</b>	<b>+0.16</b>	<b>+40.6%</b>
Hellwarrior	0 (baseline)	4.08	40.99	—	—
Hellwarrior	0.05	3.87	40.77	−0.22	+5.2%
Hellwarrior	<b>0.30</b>	<b>3.54</b>	<b>38.55</b>	<b>−2.44</b>	<b>+13.2%</b>

## 218 9 Does the coherence finding generalize?

219 We test whether the coherence finding generalizes to (a) a different deformation architecture, and  
 220 (b) real monocular video.

### 221 9.1 Cross-architecture: Deformable-3DGS

222 We implemented GAD and EER in the Deformable-3DGS codebase [19] (MLP deformation, different  
 223 loss:  $\ell_1 + 0.2 \cdot (1 - \text{SSIM})$ ). At the D-NeRF-tuned  $\lambda=0.05$ , EER slightly worsens the gap on 2/3  
 224 scenes (Lego −3.1%, T-Rex −20.8%; Hellwarrior +5.2%), because Deformable-3DGS’s loss is  
 225  $\sim 3\times$  larger at convergence and  $\lambda=0.05$  is under-regularized. A  $\lambda$  sweep (Table 4) resolves this:

226 At every  $\lambda \geq 0.15$  on Lego, EER simultaneously cuts the gap and improves test PSNR. On Hell-  
 227 warrior the same  $\lambda=0.30$  gives a +13.2% gap reduction, though at a more substantial quality cost

Table 5: HyperNeRF real monocular video (4 scenes). Same  $\lambda = 0.05$  as on D-NeRF—no per-dataset tuning. Positive reduction = improved gap.

Scene	Base gap	EER gap	Red.	$\Delta$ Test
chickchicken	5.48	<b>4.61</b>	<b>+15.9%</b>	-0.20
slice-banana	5.89	<b>5.40</b>	<b>+8.3%</b>	+0.03
vrig-3dprinter	4.49	<b>3.41</b>	<b>+24.0%</b>	+0.11
vrig-peel-banana	0.89	<b>0.83</b>	<b>+6.6%</b>	-0.23
<b>mean (n=4)</b>	<b>4.19</b>	<b>3.56</b>	<b>+14.9%</b>	-0.07

228 (-2.44 dB test PSNR)—the Lego-optimal  $\lambda$  is too strong for Hellwarrior, where a smaller  $\lambda$  (e.g.,  
 229 0.05 already gives +5.2% reduction at only -0.22 dB) would sit on a better Pareto point. The  
 230 coherence mechanism transfers across scenes within the same architecture, but the optimal  $\lambda$  is both  
 231 *architecture-specific* and *scene-sensitive*, as the dimensional analysis predicts. What does not transfer  
 232 is the numerical sweet spot—and a single cross-architecture  $\lambda$  is not the paper’s claim.

## 233 9.2 Real-world: HyperNeRF (4 scenes)

234 We ran 4DGS baseline vs. EER ( $\lambda = 0.05$ , the same value used on D-NeRF—no re-tuning) on four  
 235 HyperNeRF [12] scenes with the stock 4DGS HyperNeRF config (14K iters, lower-resolution images,  
 236 noisy poses, non-Lambertian materials).

237 EER reduces the PSNR gap on **all 4 scenes** (mean -14.9%) at an average test-PSNR cost of only  
 238 -0.07 dB—effectively free. Gap reductions range from +6.6% (vrig-peel-banana, where the baseline  
 239 gap is already small at 0.89 dB) to +24.0% (vrig-3dprinter, where test PSNR also *improves* by  
 240 0.11 dB). The same  $\lambda = 0.05$  transfers from synthetic D-NeRF to four real-world scenes without  
 241 re-tuning, confirming that the coherence finding is not an artefact of perfect synthetic poses or  
 242 Lambertian materials. iPhone/Nerfies [11] datasets and a per-scene  $\lambda$  sweep remain future work.

## 243 10 Discussion and Limitations

244 **What the two findings together say.** Split is the operational cause of the overfitting cascade: it  
 245 creates the seed Gaussians that subsequently clone (a clone-count-paradox on bouncingballs, 21K  
 246 clones vs. 4.3K splits, resolves once you notice that split is the bottleneck of the cascade). But the  
 247 gap is not fundamentally *about* the cloud’s size—if we constrain *how* the Gaussians can move (EER),  
 248 we can let the cloud grow freely and still close most of the gap. The concrete prescription is therefore:  
 249 **use GAD to avoid paying for Gaussians you don’t need, and use EER to keep the Gaussians you**  
 250 **do have from memorizing training-view-specific deformations.**

251 **Limitations.** Our main D-NeRF evaluation uses  $n = 8$  scenes; large effects (Findings 1 and 2) are  
 252 robust, but small differences should be read with caution. The primary backbone is 4DGS; cross-  
 253 architecture transfer to Deformable-3DGS requires per-loss-scale re-tuning of  $\lambda$  (Sec. 9). Standalone  
 254 EER is  $\approx 9.4\times$  slower than baseline (the  $k$ -NN cost scales with cloud size); with capacity control  
 255 capping the cloud at 22K Gaussians, overhead drops to  $1.35\times$ . We also evaluated four additional  
 256 regularizers and one diagnostic; none exceeded 10% gap reduction.

257 **Conclusion.** Splitting drives most of the overfitting in monocular dynamic Gaussian splatting; the  
 258 gap is not fundamentally about cloud size. A local-smoothness penalty on the deformation field  
 259 (EER) reduces the gap by 40.8% while growing the cloud by 85%, directly reducing per-Gaussian  
 260 strain by 99.7%. The coherence finding generalizes to Deformable-3DGS and to real HyperNeRF  
 261 video. We evaluate three regularizers (EER, GAD, PTDrop); the primary combination GAD+EER  
 262 closes 48.2% of the gap, and adding PTDrop with a growth cap reaches 57.4% at larger quality cost.

263 **References**

- 264 [1] Jeongmin Bae, Seoha Kim, Youngsik Yun, Hahyun Lee, Gun Bang, and Youngjung Uh. Per-  
265 Gaussian embedding-based deformation for deformable 3D Gaussian splatting. In *European*  
266 *Conference on Computer Vision (ECCV)*, 2024.
- 267 [2] Xiaobin Deng, Changyu Diao, Min Li, Ruohan Yu, and Duanqing Xu. Efficient density control  
268 for 3D Gaussian splatting. *arXiv preprint arXiv:2411.10133*, 2024.
- 269 [3] Guangchi Fang and Bing Wang. Mini-Splatting: Representing scenes with a constrained number  
270 of Gaussians. In *European Conference on Computer Vision (ECCV)*, 2024.
- 271 [4] Glenn Grubert, Florian Barthel, Anna Hilsmann, and Peter Eisert. Improving adaptive density  
272 control for 3D Gaussian splatting. In *International Joint Conference on Computer Vision, Imag-*  
273 *ing and Computer Graphics Theory and Applications (VISAPP)*, pages 610–621. SCITEPRESS,  
274 2025.
- 275 [5] Alex Hanson, Allen Tu, Vasu Singla, Mayuka Jayawardhana, Matthias Zwicker, and Tom  
276 Goldstein. PUP 3D-GS: Principled uncertainty pruning for 3D Gaussian splatting. In *IEEE/CVF*  
277 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- 278 [6] Yi-Hua Huang, Yang-Tian Sun, Zilong Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. SC-  
279 GS: Sparse-controlled Gaussian splatting for editable dynamic scenes. In *IEEE/CVF Conference*  
280 *on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- 281 [7] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian  
282 splatting for real-time radiance field rendering. In *ACM Transactions on Graphics (SIGGRAPH)*,  
283 2023.
- 284 [8] Yiqing Liang, Mikhail Okunev, Mikaela Angelina Uy, Runfeng Li, Leonidas Guibas, James  
285 Tompkin, and Adam W. Harley. Monocular dynamic Gaussian splatting: Fast, brittle, and scene  
286 complexity rules. *Transactions on Machine Learning Research (TMLR)*, 2025.
- 287 [9] Saswat Subhajyoti Mallick, Rahul Goel, Bernhard Kerbl, Markus Steinberger, Francisco Vicente  
288 Carrasco, and Fernando De La Torre. Taming 3DGS: High-quality radiance fields with limited  
289 resources. In *SIGGRAPH Asia Conference Papers*, 2024.
- 290 [10] Hyunwoo Park, Gun Ryu, and Wonjun Kim. DropGaussian: Structural regularization for  
291 sparse-view Gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern*  
292 *Recognition (CVPR)*, 2025.
- 293 [11] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Steven M.  
294 Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *IEEE/CVF*  
295 *International Conference on Computer Vision (ICCV)*, 2021.
- 296 [12] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B.  
297 Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. HyperNeRF: A higher-dimensional  
298 representation for topologically varying neural radiance fields. *ACM Transactions on Graphics*,  
299 40(6), 2021.
- 300 [13] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF:  
301 Neural radiance fields for dynamic scenes. In *IEEE/CVF Conference on Computer Vision and*  
302 *Pattern Recognition (CVPR)*, 2021.
- 303 [14] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. Revising densification in Gaussian  
304 splatting. In *European Conference on Computer Vision (ECCV)*, 2024.
- 305 [15] Gideon Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464,  
306 1978.
- 307 [16] Peihao Wang, Yuehao Wang, Dilin Wang, Sreyas Mohan, Zhiwen Fan, Lemeng Wu, Ruisi  
308 Cai, Yu-Ying Yeh, Zhangyang Wang, Qiang Liu, and Rakesh Ranjan. Steepest descent density  
309 control for compact 3D Gaussian splatting. In *IEEE/CVF Conference on Computer Vision and*  
310 *Pattern Recognition (CVPR)*, 2025.

- 311 [17] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu,  
312 Qi Tian, and Xinggang Wang. 4D Gaussian splatting for real-time dynamic scene rendering. In  
313 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- 314 [18] Yexing Xu, Longguang Wang, Minglin Chen, Sheng Ao, Li Li, and Yulan Guo. DropoutGS:  
315 Dropping out gaussians for better sparse-view rendering. In *IEEE/CVF Conference on Computer  
316 Vision and Pattern Recognition (CVPR)*, 2025.
- 317 [19] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable  
318 3D Gaussians for high-fidelity monocular dynamic scene reconstruction. In *IEEE/CVF Confer-  
319 ence on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- 320 [20] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. AbsGS: Recovering fine details  
321 in 3D Gaussian splatting. In *Proceedings of the 32nd ACM International Conference on  
322 Multimedia (MM)*, 2024.
- 323 [21] Zheng Zhang, Wenbo Hu, Yixing Lao, Tong He, and Hengshuang Zhao. Pixel-GS: Density  
324 control with pixel-aware gradient for 3D Gaussian splatting. In *European Conference on  
325 Computer Vision (ECCV)*, 2024.
- 326 [22] Zheng Zhou, Yu-Jie Xiong, Jia-Chen Zhang, Chun-Ming Xia, Xihe Qiu, and Hongjian Zhan.  
327 Gradient-direction-aware density control for 3D Gaussian splatting. In *International Conference  
328 on Learning Representations (ICLR)*, 2026.

## 329 **NeurIPS Paper Checklist**

### 330 **1. Claims**

331 Question: Do the main claims made in the abstract and introduction accurately reflect the  
332 paper’s contributions and scope?

333 Answer: [\[Yes\]](#)

334 Justification: The abstract states two specific findings (split dominates; coherence beats  
335 capacity) that are directly supported by the ablation (Sec. 5) and EER experiments (Sec. 6).  
336 Cross-domain transfer claims are scoped to the architectures and datasets tested.

### 337 **2. Limitations**

338 Question: Does the paper discuss the limitations of the work performed by the authors?

339 Answer: [\[Yes\]](#)

340 Justification: Sec. 10 (Discussion and Limitations) explicitly lists: sample size ( $n = 8$  D-  
341 NeRF scenes), single default backbone (4DGS), EER compute overhead ( $9.4\times$  standalone),  
342 GAD derivation being a heuristic, and the scene-sensitivity of  $\lambda$ .

### 343 **3. Theory assumptions and proofs**

344 Question: For each theoretical result, does the paper provide the full set of assumptions and  
345 a complete (or correct) proof?

346 Answer: [\[N/A\]](#)

347 Justification: The paper does not claim formal theoretical results. The GAD threshold  
348 formula (Eq. 2) is explicitly presented as a BIC-motivated heuristic, and we report the  
349 empirical falsification of the growth-bound assumption ( $\alpha \approx 0.04$  vs. assumed  $\alpha = 1$ ).

### 350 **4. Experiments reproducibility**

351 Question: Does the paper fully disclose all the information needed to reproduce the main ex-  
352 perimental results of the paper to the extent that it affects the main claims and/or conclusions  
353 of the paper?

354 Answer: [\[Yes\]](#)

355 Justification: All hyperparameters, schedules, and implementation details are specified in  
356 Secs. 4, 6, and 7. The codebase, training scripts, and configs are included in the supple-  
357 mentary. Each method is  $\sim 20$  lines of code on top of public 4DGS/Deformable-3DGS  
358 codebases.

359 **5. Open access to data and code**

360 Question: Does the paper provide open access to the data and code, with sufficient instruc-

361 tions to faithfully reproduce the main experimental results?

362 Answer: [Yes]

363 Justification: Code will be released upon acceptance. All datasets used (D-NeRF, HyperN-

364 eRF) are publicly available. Training scripts and configs are in the supplementary.

365 **6. Experimental setting/details**

366 Question: Does the paper specify all the training and test details necessary to understand the

367 results?

368 Answer: [Yes]

369 Justification: Sec. 4 specifies dataset, baseline, metrics, statistical protocol, hardware, and

370 all 8 ablation conditions. EER implementation details (Sec. 6) and GAD formula (Sec. 7)

371 are fully specified.

372 **7. Experiment statistical significance**

373 Question: Does the paper report error bars suitably and correctly?

374 Answer: [Yes]

375 Justification: We report paired  $t$ -tests with Cohen’s  $d$ , bootstrap 95% CIs for the count-gap

376 correlation, and per-scene breakdowns. We explicitly note that  $n = 8$  provides limited

377 statistical power and that small differences should not be over-interpreted (Sec. 4).

378 **8. Experiments compute resources**

379 Question: For each experiment, does the paper provide sufficient information on the com-

380 puter resources needed to reproduce the experiments?

381 Answer: [Yes]

382 Justification: All experiments run on a single NVIDIA RTX 3070 (8 GB). Per-scene training

383 time ( $\sim 13$  min baseline), EER overhead ( $1.35\times$  in combination,  $9.4\times$  standalone), and total

384 GPU budget ( $\sim 15$  hours for all experiments) are reported.

385 **9. Code of ethics**

386 Question: Does the research conducted in the paper conform, in every respect, with the

387 NeurIPS Code of Ethics?

388 Answer: [Yes]

389 Justification: This is a diagnostic study of an existing method class using public benchmarks.

390 No human subjects, no private data, no dual-use concerns.

391 **10. Broader impacts**

392 Question: Does the paper discuss both potential positive societal impacts and negative

393 societal impacts of the work?

394 Answer: [N/A]

395 Justification: This is a technical diagnostic study of overfitting in 3D reconstruction methods.

396 The direct societal impact is negligible; the methods improve training robustness of an

397 existing technique without introducing new capabilities.

398 **11. Safeguards**

399 Question: Does the paper describe safeguards that have been put in place for responsible

400 release of data or models?

401 Answer: [N/A]

402 Justification: No new datasets or pretrained models are released. The methods are small

403 regularizers ( $\sim 20$  lines each) applied to existing public codebases.

404 **12. Licenses**

405 Question: Are the creators or original owners of assets used in the paper properly credited

406 and are the license and terms of use explicitly mentioned and properly respected?

407 Answer: [\[Yes\]](#)  
408 Justification: D-NeRF, HyperNeRF, 4DGS, and Deformable-3DGS are cited. All are  
409 publicly released for research use.

410 **13. New assets**

411 Question: Are new assets introduced in the paper well documented and is the documentation  
412 provided alongside the assets?

413 Answer: [\[Yes\]](#)

414 Justification: Our code contributions (GAD, EER, PTDrop implementations + analysis  
415 scripts) are documented in the supplementary with README files and inline comments.

416 **14. Crowdsourcing and human subjects**

417 Question: For crowdsourcing experiments and research with human subjects, does the paper  
418 include the full text of instructions given to participants and screenshots?

419 Answer: [\[N/A\]](#)

420 Justification: No crowdsourcing or human subjects involved.

421 **15. IRB approvals**

422 Question: Does the paper describe potential risks incurred by study participants?

423 Answer: [\[N/A\]](#)

424 Justification: No human participants.

425 **16. Use of LLMs**

426 Question: Does the paper describe the use of LLMs in the research and/or writing process?

427 Answer: [\[Yes\]](#)

428 Justification: An LLM (Claude) was used as a coding and writing assistant during the  
429 research process, including experiment scripting, figure generation, and manuscript drafting.  
430 All scientific claims, experimental designs, and analyses were verified by the author.