

# Quantifying and Alleviating Co-Adaptation in Sparse-View 3D Gaussian Splatting

## Supplementary Material

### 1 A Theoretical Analysis of Co-Adaptation Score

2 **Note.** In the main paper, CA is reported as the average pixel-wise variance over the commonly visible  
3 region of an entire image. Here, for clarity, we derive the formula for a single pixel  $u$ .

4 We begin by formalizing how random dropout perturbs the rendered color. Under a dropout mask  
5  $\mathbf{z} = (z_1, \dots, z_n)$ —where each  $z_i \in \{0, 1\}$  independently indicates whether Gaussian  $i$  is kept—the  
6 color at pixel  $u$  is

$$C^{(\mathbf{z})}(u) = \sum_{i \in \mathcal{N}(u)} z_i c_i \alpha_i \prod_{j < i} (1 - z_j \alpha_j). \quad (1)$$

7 Here  $\mathcal{N}(u)$  is the depth-sorted list of Gaussians projecting to pixel  $u$ ,  $c_i$  and  $\alpha_i$  are the color  
8 and opacity of Gaussian  $i$ , and the product term is the usual front-to-back transmittance. The  
9 Co-Adaptation Score measures how much  $C^{(\mathbf{z})}(u)$  fluctuates as we sample different masks:

$$\text{CA}(u) = \text{Var}_{\mathbf{z}}(C^{(\mathbf{z})}(u)). \quad (2)$$

10 A large variance indicates that dropping a subset of Gaussians greatly alters the pixel color, i.e. that  
11 Gaussians have become overly co-dependent. To see how CA depends on the individual parameters,  
12 we apply a first-order expansion of the transmittance:

$$\prod_{j < i} (1 - z_j \alpha_j) \approx 1 - \sum_{j < i} z_j \alpha_j. \quad (3)$$

13 Substituting back and discarding second-order (and higher) interactions yields

$$C^{(\mathbf{z})}(u) \approx \sum_i z_i c_i \alpha_i - \sum_i z_i c_i \alpha_i \sum_{j < i} z_j \alpha_j. \quad (4)$$

14 In this approximation, the dominant term is the simple sum of active contributions. Neglecting the  
15 smaller second term, the variance becomes

$$\text{CA}(u) \approx \text{Var}_{\mathbf{z}}\left(\sum_i z_i c_i \alpha_i\right). \quad (5)$$

16 Since each  $z_i$  is a Bernoulli random variable with success probability  $\Pr[z_i = 1] = 1 - p$ , where  $p$   
17 denotes the dropout probability, its first and second moments are given by

$$\mathbb{E}[z_i] = 1 - p, \quad \text{Var}(z_i) = p(1 - p). \quad (6)$$

18 As the dropout decisions  $z_i$  are independent across different Gaussians, we can derive the variance of  
19 the total weighted sum of Gaussian contributions to pixel  $u$  as follows:

$$\text{Var}\left(\sum_i z_i c_i \alpha_i\right) = \sum_i \text{Var}(z_i) (c_i \alpha_i)^2 = p(1 - p) \sum_i (c_i \alpha_i)^2. \quad (7)$$

20 Hence, we obtain a simplified expression for the Co-Adaptation Score (CA) at pixel  $u$ :

$$\text{CA}(u) \approx p(1 - p) \sum_{i \in \mathcal{N}(u)} (c_i \alpha_i)^2. \quad (8)$$

21 This final form shows that the CA score grows proportionally with the sum of squared weighted  
22 color-opacity terms across all Gaussians contributing to pixel  $u$ , scaled by the dropout-related variance  
23 factor  $p(1 - p)$ . It quantitatively captures how variations in both color and opacity among Gaussians  
24 manifest as co-adaptation under random dropout perturbations. Equation (8) highlights several key  
25 insights. **Dependence on opacity and color:** Gaussians with large  $c_i \alpha_i$  terms contribute more  
26 significantly to the CA score. When these contributions are similar or redundant, the resulting

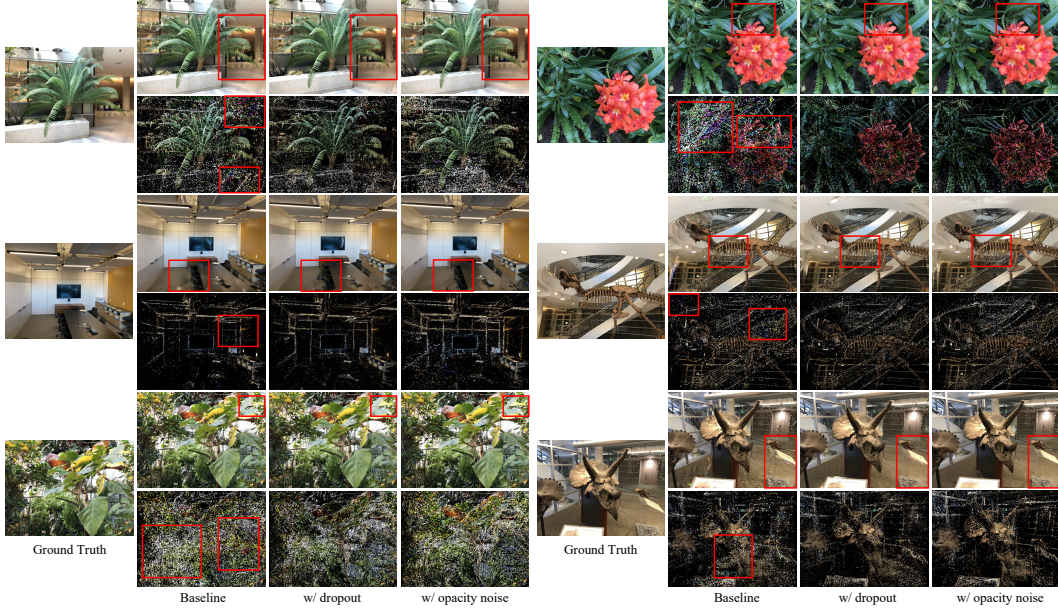


Figure 1: Visual comparison on the LLFF dataset based on 3DGS. Suppressing co-adaptation reduces color noise (see the changes in the GS point cloud) and improves scene geometry and detail quality.

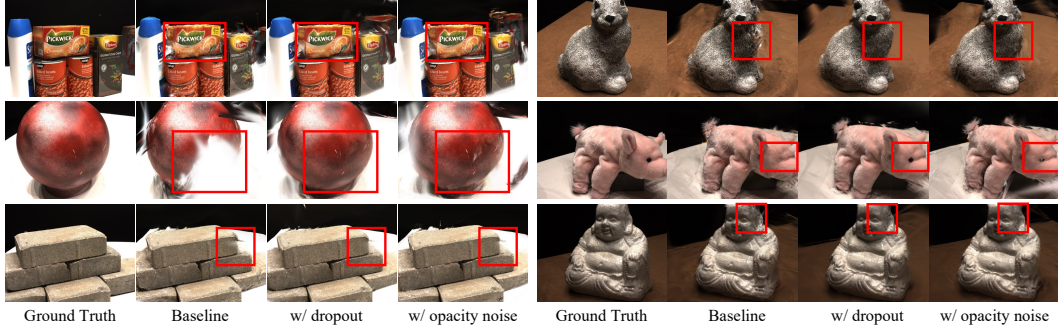


Figure 2: Visual comparison on the DTU dataset based on 3DGS. Suppressing co-adaptation leads to more consistent scene fusion across views, resulting in clearer structure and more accurate details.

27 variance—and thus co-adaptation—is low. **Effect of dropout rate:** The multiplicative factor  $p(1 - p)$   
 28 reaches its maximum at  $p = 0.5$ , justifying our choice to adopt this value in practice to enhance the  
 29 sensitivity of the CA metric. **Regularization implications:** Perturbations to  $c_i \alpha_i$  (e.g., via noise  
 30 injection or dropout) reduce the magnitudes of the summation terms, leading to lower CA scores and  
 31 implicitly mitigating excessive interdependence among Gaussians.

32 In summary, this theoretical analysis substantiates the CA score as a principled indicator of color-  
 33 opacity coupling strength under dropout perturbations, thereby providing a foundation for the design  
 34 of our co-adaptation suppression strategies.

## 35 B More Qualitative Results

36 We present additional qualitative results on the LLFF [1], DTU [2], and Blender [3] datasets to further  
 37 illustrate the visual improvements brought by our co-adaptation suppression strategies.

38 **LLFF.** We first examine the LLFF dataset, where improvements are clearly observed both in rendered  
 39 images and Gaussian point clouds. As shown in Figure 1, our methods yield point clouds that are  
 40 cleaner and more geometrically coherent. The baseline 3D Gaussian Splatting (3DGS) [4] often  
 41 produces scattered, colorful speckles, which are significantly suppressed after applying our co-  
 42 adaptation regularization. These visual improvements in point clouds directly lead to better rendering  
 43 in novel views. For instance, in the *trex* scene, fossil bones exhibit high-frequency flickering and

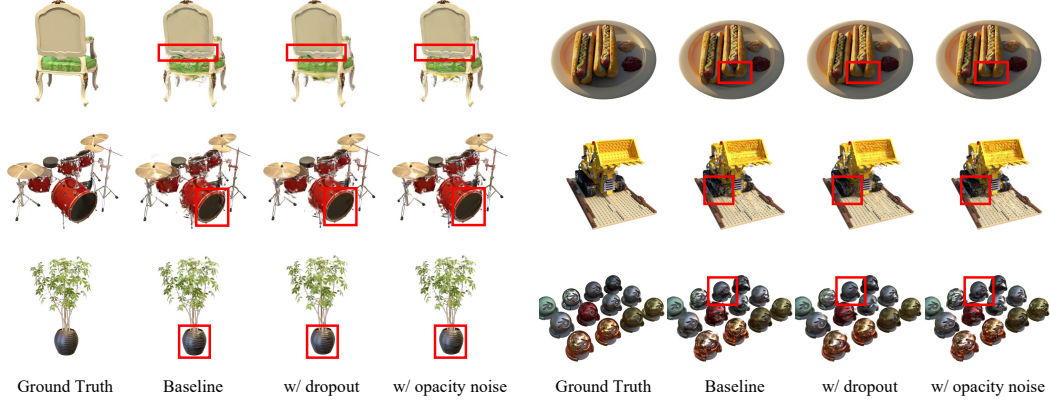


Figure 3: Visual comparison on the Blender dataset based on 3DGS. Suppressing co-adaptation reduces floating artifacts and enhances the completeness and clarity of fine structures.



Figure 4: Visual comparison on LLFF dataset across multiple methods, using Binocular3DGS as the baseline. Suppressing co-adaptation further reduces colorful speckle artifacts and floating noise, while enhancing the completeness and structural coherence of the reconstructed geometry.

44 speckling in the baseline, which are clearly alleviated with our methods. In the *flower* scene, petals  
 45 and leaves display incorrect colorization under novel views, while the regularized versions better  
 46 preserve the natural appearance. Similarly, in the *fern* scene, severe color inconsistencies appear  
 47 on the pole, ground, and leaf textures under the baseline, but these are significantly reduced when  
 48 co-adaptation is suppressed. Overall, the results suggest that reducing co-adaptation enhances both  
 49 structural consistency and appearance fidelity in sparse-view 3DGS.

50 **DTU.** On the DTU dataset, we observe that both co-adaptation suppression strategies significantly  
 51 enhance structural coherence and reduce visual artifacts in sparse-view 3DGS. As illustrated in  
 52 Figure 2, suppressing co-adaptation leads to improved fusion across views, enabling more accurate  
 53 and consistent detail reconstruction. In the scene containing multiple stacked boxes, our methods help  
 54 preserve fine-scale textual details on the box surfaces, with the characters on the packaging appearing  
 55 noticeably clearer compared to the baseline. In the scene with the red ball, our regularization  
 56 strategies substantially reduce the bright floating speckles that frequently appear in front of the  
 57 camera, especially near the lens center, leading to a cleaner and more complete object reconstruction.  
 58 Furthermore, in the scene featuring a small pink pig plush toy, the eyes of the pig—barely visible in



Table 1: Ablation study on the noise scale  $\sigma$  applied to **scale parameters** in Binocular3DGS on the LLFF dataset. We use multiplicative noise with different  $\sigma$  values and evaluate the impact on rendering quality and co-adaptation.

Scale Noise $\sigma$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Train CA $\downarrow$	Test CA $\downarrow$
0.0 (Baseline)	21.440	0.751	0.168	0.001845	0.001951
0.1	21.551	0.755	0.166	0.001520	0.001682
0.2	<b>21.760</b>	0.760	0.161	0.001326	0.001432
0.3	21.724	0.762	0.161	0.001148	0.001238
0.4	21.690	0.762	0.161	0.001075	0.001237
0.5	21.674	0.765	<b>0.159</b>	0.001009	0.001091
0.6	21.685	<b>0.766</b>	<b>0.159</b>	<b>0.000927</b>	0.001014
0.7	21.597	0.762	0.162	0.000945	<b>0.001012</b>

the baseline—are successfully reconstructed under novel views when using either of our suppression methods. These improvements confirm that controlling co-adaptation strengthens the model’s ability to render fine geometric and appearance details in complex real-world scenes.

**Blender.** On the Blender dataset, we observe consistent improvements in rendering quality and structural fidelity when applying co-adaptation suppression strategies. As shown in Figure 3, both dropout and opacity noise help reduce subtle artifacts and enhance reconstruction realism under sparse-view supervision. In the *chair* and *drums* scenes, novel view renderings generated with our methods exhibit fewer detail-related artifacts, particularly around fine structures such as chair legs and drum edges, resulting in more accurate geometry and cleaner appearance. In the *figus* scene, both strategies significantly improve the photorealism of the reconstructed flowerpot, leading to more consistent shading and geometry that better match the true scene layout. Additionally, in the *hotdog* scene, we observe a clear reduction in floating speckle artifacts outside the object boundary—especially near the sausage and tray—demonstrating the effectiveness of co-adaptation suppression in eliminating non-semantic visual noise. Together, these results highlight the generalizability of our techniques to synthetic datasets with complex textures and fine-scale geometry.

**Method Comparison.** Figure 4 presents a visual comparison across multiple 3DGS methods, using Binocular3DGS [5] as the baseline. Binocular3DGS is a strong baseline with geometry-aware initialization, yet our co-adaptation suppression strategies (dropout and opacity noise) still lead to noticeable improvements. In the *flower* scene, existing sparse-view approaches exhibit occasional color speckles and geometry gaps; applying our methods further reduces these artifacts and yields more complete, coherent novel view reconstructions.

## C More Experiments

### C.1 Exploration of Noise Injection on Different Gaussian Parameters

Beyond the opacity noise experiments presented in the main paper, we further explore noise injection on other Gaussian parameters, including the scale, position, and SH (spherical harmonics) coefficients. These experiments aim to analyze whether perturbing other parameters can also suppress co-adaptation and improve rendering quality under sparse-view settings.

**Scale Noise Injection.** We first examine the effect of injecting multiplicative noise into the scale parameters of Gaussians. Similar to opacity noise, we apply multiplicative perturbation as follows:

$$\tilde{s} = s \cdot (1 + \epsilon), \quad \epsilon \sim \mathcal{N}(0, \sigma^2), \quad (9)$$

where  $s$  denotes the original scale and  $\sigma$  is the noise strength. The results of this experiment are summarized in Table 1. Compared to the Binocular3DGS baseline, scale noise yields limited improvements in overall reconstruction metrics such as PSNR, SSIM [6], and LPIPS [7]. However, it reduces the co-adaptation strength (CA) on both training and novel views, indicating that scale perturbation helps decouple Gaussian interactions to some extent. This finding aligns with our visual observations in Figure 5, where scale noise improves scene-level geometric coherence in some cases, but introduces mild local blurring in object details. The observed trade-off can be explained by the functional role of scale in Gaussian rendering: perturbing scale slightly alters the spatial support of Gaussians, which may improve coverage and fusion in sparse-view settings, but risks oversmoothing



Table 2: Ablation study of **position** and **SH noise** injection on the LLFF dataset. Position noise is added as Gaussian noise scaled by nearest-neighbor distance; SH noise is applied multiplicatively.

Noise Setting	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Train CA $\downarrow$	Test CA $\downarrow$
Baseline	21.440	0.751	0.168	0.001845	0.001951
w/ position noise 0.1	21.093	0.742	0.171	0.001701	0.001812
w/ position noise 0.3	20.760	0.726	0.182	0.001594	0.001730
w/ position noise 0.5	20.540	0.715	0.189	0.001653	0.001733
w/ SH noise 0.1	21.475	0.749	0.171	0.001769	0.001876
w/ SH noise 0.3	21.391	0.750	0.170	0.001633	0.001753
w/ SH noise 0.5	21.390	0.751	0.168	0.001799	0.001906

Table 3: Ablation study on different **SH orders** used in 3DGS under sparse-view settings on the LLFF dataset. We evaluate rendering quality and co-adaptation across SH0 to SH3.

SH Order	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Train CA $\downarrow$	Test CA $\downarrow$
0	19.293	0.652	0.235	<b>0.006736</b>	<b>0.007259</b>
1	19.239	0.651	0.234	0.007087	0.007587
2	<b>19.377</b>	<b>0.657</b>	<b>0.229</b>	0.007167	0.007858
3 (Baseline)	19.357	0.651	0.232	0.007543	0.008206

fine structures. As a result, while scale noise contributes to better scene fusion in some examples, it reduces local detail sharpness, leading to modest overall gains.

**Position Noise Injection.** For position parameters, we introduce additive noise scaled by the nearest-neighbor distance of each Gaussian, which perturbs their spatial locations and disrupts alignment with the image plane. Even under small perturbations, this leads to degraded convergence and blurred geometry in sparse-view training. Although a coarse-to-fine training strategy can somewhat stabilize optimization, its effect on overall reconstruction quality remains limited.

**SH Noise Injection.** For SH parameters, we apply multiplicative noise in the form  $\tilde{c} = c \cdot (1 + \epsilon)$  with  $\epsilon \sim \mathcal{N}(0, \sigma^2)$  to perturb the directional color components. However, since SH parameters do not affect the spatial configuration of Gaussians along a ray, they do not influence the contribution weights of Gaussians during the alpha blending process. As a result, SH noise has limited effect on alleviating co-adaptation. While it may seem intuitive that perturbing SH and opacity could yield similar regularization effects—since both influence final pixel colors—their impact on rendering differs fundamentally. Opacity noise directly modifies the weight assigned to each Gaussian during alpha blending, thereby altering the effective contribution and even the number of Gaussians used to render a pixel. This enables opacity noise to disrupt the co-adapted structure more substantially. In contrast, SH perturbation only changes the emitted color of each Gaussian without modifying its blending weight, thus providing limited relief against co-adaptation.




Figure 5: Visual comparison of scale noise effects in Binocular3DGS.

## C.2 Comparison of Different SH Orders in Sparse-View 3DGS

To investigate whether the choice of SH (spherical harmonics) order plays a significant role in co-adaptation, we conduct a controlled ablation study using SH orders from 0 to 3 in vanilla 3DGS training under sparse-view settings. The quantitative results are shown in Table 3, and corresponding qualitative comparisons of Gaussian point clouds are visualized in Figure 6.

We observe that lower SH orders, especially SH0, tend to yield slightly lower co-adaptation scores and cleaner Gaussian point distributions. In contrast, SH1 to SH3 occasionally produce localized color speckles, but these artifacts appear stochastically—without a consistent correlation to SH order. For example, SH1 may exhibit more noticeable artifacts than SH2 in certain training runs. Overall, the differences in co-adaptation strength across SH0 to SH3 remain relatively minor, and the emergence of appearance artifacts is not strongly determined by the SH order. These findings suggest that while




Figure 6: Visualization of Gaussian point clouds under different SH orders on the LLFF *flower* scene. SH0 yields cleaner points, while higher SH orders occasionally introduce more color speckles.

SH complexity may influence directional appearance expressiveness, it is not a dominant factor in controlling co-adaptation or mitigating color artifacts in sparse-view 3DGS.

## D More Experimental Details

To ensure fair comparisons, we adopt unified parameter settings across all scenes within each dataset. For example, while the original DNGaussian [8] uses different training configurations for each Blender scene, we apply a consistent setup across all scenes. Binocular3DGS is retrained using white backgrounds to match other methods, as its original results were obtained using black backgrounds. For co-adaptation suppression, we use a fixed dropout probability of 0.2 across all methods and datasets, based on ablations conducted on Binocular3DGS. Because each method learns a distinct opacity distribution, we tune the opacity noise scale individually for each method within  $[0.05, 0.8]$  and fix it across all scenes in the same dataset. For CA score computation, we randomly drop 50% of the contributing Gaussians for each view and calculate the pixel-wise variance across multiple renderings. The number of training iterations for each baseline follows its official implementation.

## E Limitations

Although our proposed strategies—random Gaussian dropout and opacity noise injection—effectively reduce co-adaptation and enhance novel view synthesis under sparse supervision, we observe that overly suppressing co-adaptation does not always yield further improvements in rendering quality. As evidenced in our ablation studies, rendering performance tends to plateau or even degrade once co-adaptation scores fall below a certain threshold. This suggests that co-adaptation is not inherently harmful; rather, it constitutes an essential aspect of 3DGS, facilitating cooperation among Gaussians to model fine-grained appearance and geometry. The negative effects of co-adaptation primarily emerge under sparse-view training, where excessive entanglement can lead to artifacts or overfitting. Consequently, while moderate suppression is beneficial for improving generalization, excessive suppression may compromise the model’s capacity to accurately fit scene content, ultimately limiting expressiveness. Our goal is thus not to eliminate co-adaptation entirely, but to selectively mitigate its detrimental forms. Future work may consider adaptive suppression mechanisms that dynamically balance generalization and representational fidelity.

## References

- [1] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (ToG)*, 38(4):1–14, 2019.
- [2] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014.
- [3] 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. *Communications of the ACM*, 65(1):99–106, 2021.

- 169 [4] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian  
170 splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.
- 171 [5] Liang Han, Junsheng Zhou, Yu-Shen Liu, and Zhizhong Han. Binocular-guided 3d gaussian  
172 splatting with view consistency for sparse view synthesis. In *The Thirty-eighth Annual Conference*  
173 *on Neural Information Processing Systems*.
- 174 [6] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment:  
175 from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):  
176 600–612, 2004.
- 177 [7] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable  
178 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*  
179 *computer vision and pattern recognition*, pages 586–595, 2018.
- 180 [8] Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. Dngaussian:  
181 Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In  
182 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages  
183 20775–20785, 2024.