Post-Training Depth-of-Field Control for 3D Gaussian Splatting Scenes

Controllable depth of field can be added to already trained 3DGS scenes at render time, but large blur requires alpha-tail suppression to avoid artifacts. — ***Adam Sztano*** (sztano@student.tudelft.nl), ***Michael Weinmann*** (Supervisor)

1. 3DGS: Photorealistic, but Always in Focus

3D Gaussian Splatting (3DGS) [2] reconstructs a scene from many input photographs by representing it as millions of small 3D Gaussian primitives. Each Gaussian stores position, shape, opacity, and color, allowing the scene to be rendered from new viewpoints in real time.

However, standard 3DGS rendering usually behaves like a pinhole camera: the whole scene appears sharp. Real cameras have finite apertures, which create depth of field: objects near the focal plane appear sharp, while objects in front of or behind it become blurred.

Goal: add interactive focus distance and aperture-like blur control to already trained 3DGS scenes, without retraining the scene.

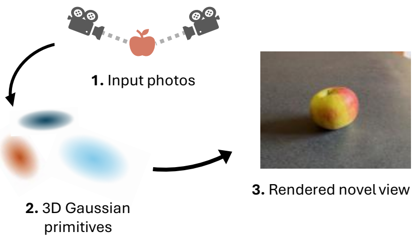

2. Research Question & Contributions

Research question:
Can an already trained 3D Gaussian Splatting scene be extended with interactive depth-of-field control at render time?

Contributions:

1. **Post-training depth of field:** defocus blur is added during rendering, without retraining the Gaussian scene.
2. **Interactive controls:** the user can change focal-plane depth and aperture-like blur strength in the viewer.
3. **Artifact mitigation:** per-pixel alpha thresholding reduces color bleeding and halo artifacts caused by enlarged Gaussian tails.

3. Method: Defocus by Expanding Gaussian Footprints

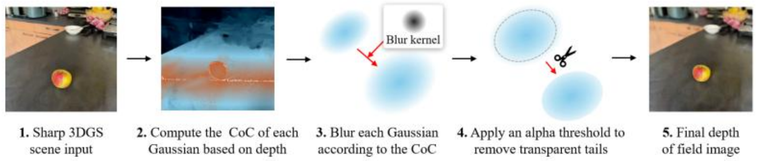

This implementation builds on the depth-of-field formulation used in DOF-GS [3]. Depth of field is controlled using a circle-of-confusion estimate: the circle of confusion is the blur radius of a point that is not on the focal plane. Points near the focal plane have a small blur radius, while points farther away have a larger one.

$$R_{\mathrm{CoC}} = \frac{1}{2} A \left| \frac{1}{z} - \frac{1}{f} \right|$$

where:
R_{CoC}: circle-of-confusion-like blur radius
z: camera-space depth of the Gaussian center
f: selected focal-plane depth
A: aperture-like blur strength

The computed blur radius is added to the projected 2D Gaussian footprint in screen space. Gaussians near the focal plane stay sharp, while Gaussians farther from it become larger and more blurred. Unlike full blur-aware training methods, this is applied only at render time, so the original trained 3DGS scene remains unchanged.

4. Artifact Control: Cutting Gaussian Tails

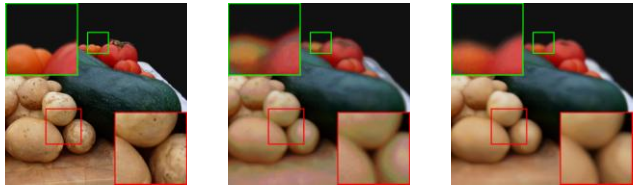

***τ* = 0, *A* = 0** | ***τ* = 0, *A* = 100** | ***τ* = 2/255, *A* = 100** [1]

Large aperture values enlarge the projected Gaussian footprints. This creates defocus blur, but also makes weak low-opacity Gaussian tails visible near object boundaries, causing color bleeding and halo-like artifacts. The minimum alpha threshold **τ** removes very weak per-pixel Gaussian contributions before they accumulate into visible artifacts.

Observed trade-off:
(**τ** = 0): strongest artifacts at high aperture
(**τ** = 1/255): reduces artifacts, but not always enough
(**τ** = 2/255): best visual trade-off in tested scenes
(**τ** = 3/255): can help extreme blur, but removes faint detail

5. Result 1: Interactive Refocusing Works

All sharp | ***Focus near*** | ***Focus far*** [1]

Changing the focal-plane depth changes which part of the scene remains sharp.
Experimental setup: Same scene, camera pose, and aperture-like blur strength. Only the focal-plane depth (f) changes, plus a reference sharp image.
Observation: Objects near the selected focal plane stay sharp, while objects in front of or behind it become blurred. The method provides predictable post-training focus control without retraining the 3DGS scene.

6. Result 2: Alpha Thresholding Improves Quality and Speed

Frame times (lower is better)	*A = 0*	*A = 50*	*A = 100*
τ = 0	Baseline (1x)	3.33x	9.13x
τ = 1/255	0.93x	1.17x	1.30x
τ = 2/255	0.93x	1.03x	1.10x

Increasing aperture-like blur strength ***A*** creates stronger depth-of-field blur, but also exposes low-opacity Gaussian tails as halos and color bleeding. Minimum alpha thresholding suppresses these weak tail contributions.
Visual result: ***τ*** *= 2/255* gave the best trade-off in the tested scenes: it reduced artifacts while preserving most visible detail. For an example, see the Artifact Control comparison.
Performance result: At ***A*** *= 100*, rendering without thresholding required 9.13x the baseline frame time. With **τ** *= 2/255,* the same setting was only (1.10x) the baseline.
Takeaway: Per-pixel alpha thresholding makes large-aperture rendering both cleaner and much faster.

7. Limitations + Generalization

The method was also tested on additional public 3DGS scenes, where focus control behaved consistently beyond the main test scene. These experiments show that the render-time approach is not specific to one reconstruction. However, the limitations also became clearer: alpha thresholding reduces halos and color bleeding, but it can remove real faint details when strong blur makes their opacity contribution very small. In some cases, subtle colors or faint background structures can partially disappear. Removing whole low-opacity Gaussians was less effective, because it still left halo artifacts near object boundaries. Overall, the method is practical for interactive post-training depth of field, but remains an approximation rather than a physically complete lens simulation.

References

[1] Simon Bethke (https://superspl.at/user?id=simonbethke). Vegetables HQ - Lichtfeld Studio. https://superspl.at/scene/f592397a, 2026. Creative Commons Attribution 4.0 License, **[2]** Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuhler, and George Drettakis. 3D Gaussian Splatting for real-time radiance field rendering. ACM Transactions on Graphics, 42(4):1–14, July 2023. **[3]** Yujie Wang, Praneeth Chakravarthula, and Baoquan Chen. DOF-GS: Adjustable depth-of-field 3D Gaussian Splatting for post-capture refocusing, defocus rendering and blur removal. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 21297–21306, 2025.