AffineGS: Unlocking Accurate Geometry in 3DGS via Spatially Varying Affine Rectification

Tze Ho Elden Tse[1,2*] Jizong Peng[1*] Richard Chen[1] Angela Yao[2]

[1]dConstruct Robotics [2]National University of Singapore [*]Equal contribution

Monocular depth is distorted piecewise linearly. Rectify it locally, and 3DGS gets accurate geometry.

3DGS overfits appearance: photoreal renders, noisy geometry. Depth-guided fixes align a monocular prior with **one global scale and shift**, an assumption that breaks.

3DGS (Splatfacto): photoreal render, noisy mesh

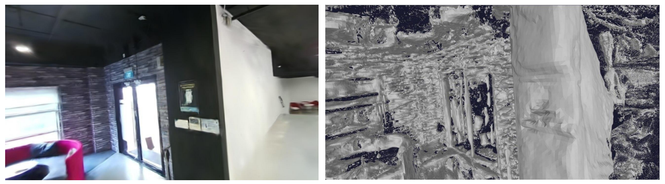

AffineGS (ours): same photorealism, clean mesh



Key observation

One global scale-and-shift cannot fit monocular disparity

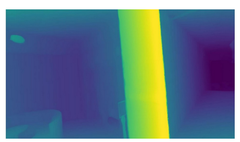

Ground-truth disparity

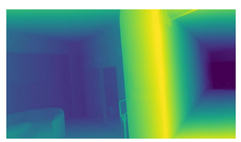

Monocular (Depth-Anything-V2)



Spatially coherent affine regions

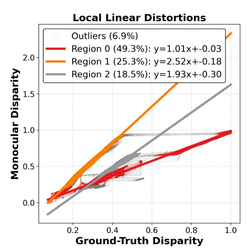

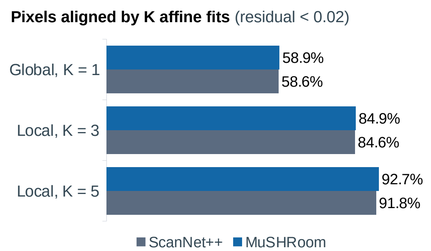

Breaks follow object boundaries; each region has its own scale and shift. MoGe and Metric3D-V2 behave the same. A single global fit either **averages the error** or **warps all but one region**.

Method

Spatially varying affine rectification of the monocular prior

Replace the global fit with scale and shift fields over the image:

$$\frac{1}{d(u,v)} = \alpha(u,v)\,\mathcal{D}(u,v) + \beta(u,v)$$

d: monocular disparity, $\mathcal{D}$: rendered 3DGS depth. Fields vary slowly within a surface, so gradients give α, β in closed form, **without sensor depth**; a guided filter keeps them sharp at edges.

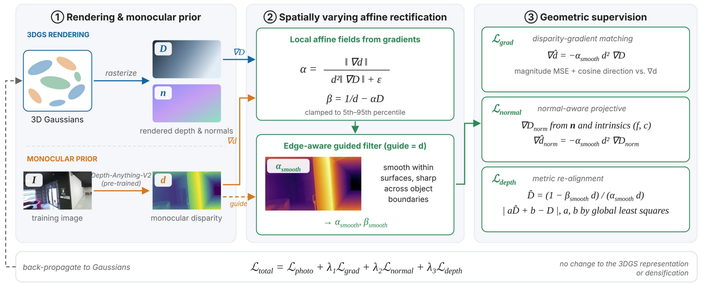

- $\mathcal{L}_{grad}$ matches predicted and measured disparity gradients, in magnitude and direction.
- $\mathcal{L}_{normal}$ derives depth gradients from rendered normals (pinhole model), orienting the Gaussians.
- $\mathcal{L}_{depth}$ applies L1 to the rectified, globally re-aligned depth, preventing scale drift.
- **No change to the Gaussian representation or densification.**

Sparse views (NeRF-LLFF)

2–5 training views · mean of 4 scenes × 10 random view sets

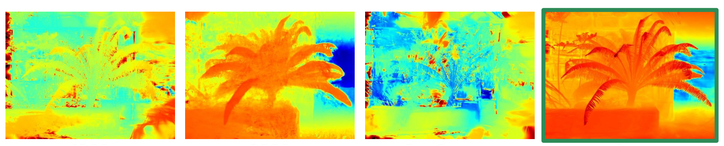

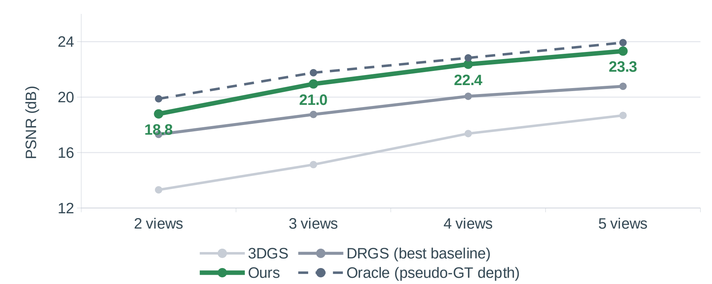

Depth & novel views (MuSHRoom)

Same Depth-Anything-V2 prior; DN-Splatter (mono) aligns it globally



DN-Splatter (mono)
global alignment

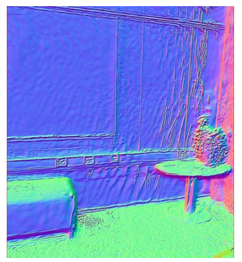

DN-Splatter
uses sensor depth

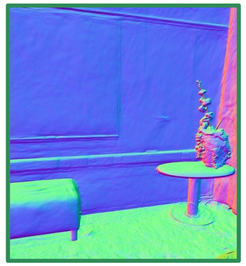

Ours
rectified prior

Depth error, Abs Rel ↓ (rendered normals shown above)

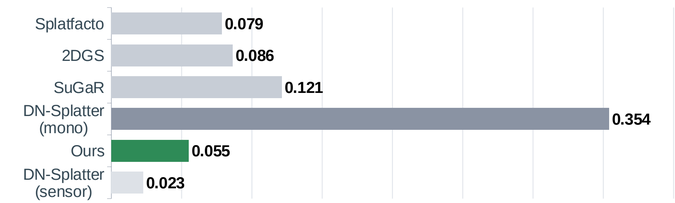

6.4× lower depth error than DN-Splatter (mono), with matching PSNR: **24.49** vs. 24.58 for DN-Splatter with sensor depth.

Meshes

Extracted meshes on ScanNet++



DN-Splatter



Ours

Mesh F-score ↑ (MuSHRoom, 5 cm)

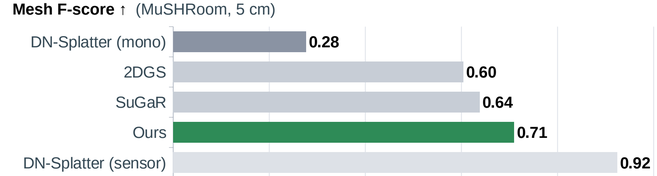

Ablation: all three losses lift F-score 0.728 → **0.855** for −0.35 dB PSNR.

Project page & code

eldentse.github.io/AffineGS · github.com/eldentse/AffineGS

Kerbl et al. 3D Gaussian Splatting. ACM ToG 2023. Turkulainen et al. DN-Splatter. WACV 2025. Chung et al. DRGS. CVPRW 2024. Yang et al. Depth Anything V2. NeurIPS 2024. Ren et al. MuSHRoom. WACV 2024.