

Outdoor-Targeted Perceptual Specialisation in 3DGS-MCMC via Late-Stage Anisotropic Primitive Splitting at Iso-Compute

Anonymous Submission
Paper ID 0




Figure 1. **Late-stage needle/capsule specialisation in 3DGS-MCMC.** After standard MCMC training has converged, we identify high-aspect Gaussian primitives ($s_{\max}/s_{\min} > 8$ — typically branches, fence wires, and foliage edges) and replace each with a chain of needle-shaped sub-primitives along its long axis, then continue training. At matched primitive cap on nine Mip-NeRF 360 scenes the modification improves LPIPS on every scene; the gain is concentrated on outdoor scenes (mean $\Delta_{\text{LPIPS}} = -0.041$) where high-aspect primitives are common (foliage, branches, fence wires), and is small on indoor scenes (mean $\Delta_{\text{LPIPS}} = -0.008$) where most surfaces are well-modelled by near-isotropic ellipsoids. The contribution is an outdoor-targeted perceptual specialisation, paid for by a small PSNR regression on indoor scenes whose vanilla baselines are already well-fitted. Right panel: a representative crop on BICYCLE where the swap recovers the foreground spokes that vanilla MCMC blurs.

Abstract

3D Gaussian Splatting trained as Markov Chain Monte Carlo [24] converges to a fixed-cap primitive distribution in which thin scene structures (branches, fence wires, foliage silhouettes) are typically covered by a small number of highly anisotropic Gaussians. We ask whether such primitives, once identified at the end of training, can be specialised to a shape better suited to thin-structure coverage without retraining from scratch. We propose NEEDLESWAP: late in training, every primitive whose co-variance has aspect ratio above a threshold is replaced by a

chain of capsule-shaped sub-primitives along its long axis, each inheriting the parent’s appearance, after which standard MCMC continues. The change adds no new learnable parameters and preserves the relocation-proposal moment-matching analysis of the original method. On all nine Mip-NeRF 360 scenes at matched primitive cap, NEEDLESWAP improves perceptual quality on every scene, with the gain dominated by outdoor scenes ($\Delta_{\text{LPIPS}} = -0.041$) over indoor scenes (-0.008). The split tracks the density of high-aspect Gaussians in the converged sample: outdoor scenes carry many such primitives along thin structures that the swap directly targets, indoor scenes carry few. The target-

ing comes at a cost: the post-swap continuation roughly doubles the active primitive count, producing a small mean PSNR regression on indoor scenes whose vanilla baselines are already well-fitted. We frame the contribution as an outdoor-targeted perceptual specialisation, not a uniform PSNR improvement. As a methodological aside, our audit caught a primitive-cap mismatch in an earlier experiment that would have inflated PSNR deltas across the entire mechanism search; we document the failure mode for downstream work.

1. Introduction

3D Gaussian Splatting [23] represents a scene as an unordered set of anisotropic Gaussian primitives rasterised through a tile-based α -compositor. Kheradmand et al. [24] reformulated densification as Markov Chain Monte Carlo: low-opacity primitives are resampled α -proportionally, and a covariance-shaped Langevin term perturbs every position on every iteration. The resulting sampler produces a fixed-cap distribution and is now a common backbone for follow-up work.

The method has a known representational ceiling on fine, elongated structure. The radial profile of an anisotropic Gaussian is exponential along every principal axis, so a primitive stretched to cover a long thin element (a branch, a fence wire, a leaf silhouette) still falls off smoothly across that element’s narrow dimension. Where the underlying geometry is closer to a step than to a Gaussian profile, the primitive must either span the discontinuity (blurring the silhouette) or shrink and leave neighbouring pixels uncovered. At convergence, a small subset of high-aspect primitives carry exactly this kind of structure, and pixel error in those regions is dominated by kernel mismatch rather than sampling variance.

We propose a late-stage modification that addresses this regime directly. Once standard MCMC has converged, we identify the subset of primitives whose covariance has aspect ratio above a threshold and replace each one by a short chain of needle-shaped sub-primitives placed along its long axis, each inheriting the parent’s appearance. Training then continues for a fixed number of additional iterations under the unchanged MCMC update. The construction adds no new learnable parameters and leaves the rasteriser, the relocation proposal, the opacity gate, and the loss untouched.

What was actually measured. We evaluate on all nine Mip-NeRF 360 scenes [1], single-seed at matched primitive cap. The reading is best stated as a scene-type split. LPIPS improves on every scene, with the gain concentrated outdoors (-0.0414) and small indoors (-0.0079). SSIM is at parity overall, with a clear outdoor gain. Outdoor PSNR is at parity (mixed signs across scenes); indoor scenes regress

(-0.444 mean), pulling the nine-scene PSNR average to -0.184 . We do not claim a PSNR improvement.

Why outdoor and not indoor. The split tracks the density of high-aspect primitives in the converged sample. Outdoor scenes carry a larger fraction of primitives with $s_{\max}/s_{\min} > 8$ because their geometry is dominated by thin elongated structures (branches, fence wires, foliage silhouettes, blades of grass). Indoor scenes are dominated by piecewise-smooth surfaces for which a single near-isotropic Gaussian is already a good match. The swap therefore acts strongly on outdoor scenes (specialising capacity at exactly the places vanilla MCMC blurs) and weakly on indoor scenes. The post-swap continuation also roughly doubles the active primitive count, paying a small mean-squared-error cost on indoor scenes whose vanilla baselines are already saturated.

A methodological aside on cap-mismatch. While auditing our experiments we discovered a configuration failure mode in an earlier sweep: two arms of a comparison were inadvertently trained at different primitive caps, producing an apparent multi-tenth-dB win that disappeared once the caps were matched. The failure is easy to commit, invisible in aggregate-only reporting, and inflates per-scene deltas by amounts comparable to several recent published improvements in this neighbourhood. We document it in Sec. 4.4 as a methodological recommendation, not as a result of the present method.

Contributions.

- **NeedleSwap: a late-stage anisotropic-primitive specialisation step for 3DGS-MCMC.** A single conditional replacement at a fixed iteration that adds no learnable parameters and preserves the equilibrium analysis of the original sampler.
- **An outdoor-targeted perceptual-quality result at iso-compute.** On all nine Mip-NeRF 360 scenes at matched cap, NEEDLESWAP improves LPIPS on every scene, with the gain concentrated on outdoor scenes that carry the high-aspect primitives the swap targets.
- **A mechanistic explanation of the scene-type split.** The outdoor/indoor gap in per-scene deltas tracks the gap in the fraction of high-aspect primitives at the swap iteration; the swap is effective in proportion to the size of the population it acts on.
- **A documented cap-mismatch reproducibility pitfall.** A small audit caught a comparison whose arms had different primitive caps and would otherwise have produced an apparent PSNR win across our entire mechanism search. We document the failure mode and recommend a matched-cap reporting discipline for future work.

2. Related Work

3D Gaussian Splatting. 3DGS [23] reformulates radiance-field rendering as tile-based rasterisation of anisotropic Gaussians, inheriting the EWA splatting machinery of Zwicker et al. [63] and the volumetric compositing of Max [37], supplanting the volumetric MLP lineage of NeRF and its scale-aware variants [1, 2, 38]. Subsequent work hardens the rasteriser against aliasing and pop artefacts via 3D smoothing filters [55], analytic integration [30], multi-scale primitives [52], and stable depth ordering [44]; surveys [10, 13, 50] chart the resulting design space. Our work leaves the rasteriser untouched.

MCMC formulations and densification. Kheradmand et al. [24] reinterpret 3DGS training as Stochastic Gradient Langevin Dynamics: densification and pruning become births and deaths, and a fixed-temperature noise term is injected into mean updates. This recasting eliminates many heuristic clone/split rules [3, 9, 14, 39] and has inspired related primitive-flow controllers [11, 35, 54, 60]. Earlier samplers in this family include SGLD [49] and the discretisation analysis of Roberts and Tweedie [46], with preconditioning [42] and momentum extensions [5], building on simulated annealing [26] and noise-injection analyses of SGD [40]. We sit strictly inside the MCMC family of Kheradmand et al. [24]: rather than retune the optimiser or change the proposal, we make a one-shot edit of the primitive support after equilibrium.

Optimisation landscape of 3DGS. A second line asks whether the first-order Adam [25] default is the right inner loop. SteepGS [48] characterises saddle points; 3DGS-LM [18] replaces SGD with Levenberg–Marquardt; second-order and quasi-Newton variants [19, 21, 43, 61] and frequency-aware optimisers [12, 28, 58] change *which* direction parameters move. Our edit is orthogonal: the deterministic update direction is unchanged.

Primitive variations. A parallel thrust replaces the Gaussian kernel itself: Deformable Beta Splatting [32] and generalised exponential splats [16]; 3D Convex Splatting [17]; effective-rank parameterisations [22]; 2D-disk splats [8, 20]; and Homogeneous-coordinate Gaussians [33] for unbounded scenes. We do not introduce a new kernel; we reconfigure the sample of the existing Gaussian primitive on its high-aspect tail.

Geometric priors and outdoor scenes. Geometry-supervised 3DGS injects monocular depth [7, 29, 51, 62], surface normals [6, 47], and explicit surface losses [4, 15, 56, 57]. Mip-NeRF 360 [1] and NeRF++ [59] established

the unbounded-capture protocol; HoGS [33] and scaffold-style hierarchies [34, 41, 45] adapt 3DGS to those settings. Anti-aliasing in the kernel [30, 53, 55] acts in image space; ours acts on the primitive set. Appearance-embedding methods for in-the-wild capture [27, 31, 36] target aleatoric variation across captures and are orthogonal to our late-stage edit.

Position. Among the four families above (MCMC formulation, optimiser change, primitive change, prior injection), our contribution lives strictly inside the MCMC family. Unlike optimiser-replacement work [18, 21, 48], we do not alter the descent direction. Unlike primitive variants [16, 17, 32], we do not alter the kernel. Unlike prior-driven approaches [4, 29, 47], we add no supervision. We perform a one-shot edit of the primitive support after equilibrium and let the unchanged sampler re-equilibrate.

3. Method

3.1. Background: 3DGS-MCMC and primitive shape

3DGS [23] represents a scene as a set of anisotropic Gaussian primitives $\{\mathbf{p}_i, \Sigma_i, \alpha_i, \mathbf{c}_i\}$ optimised by photometric supervision under EWA splatting [63]. The MCMC formulation of Kheradmand et al. [24] reinterprets densification as state-dependent Metropolis–Hastings transitions and adds a Langevin step to every position update,

$$\mathbf{p}_i \leftarrow \mathbf{p}_i + \Sigma_i \boldsymbol{\varepsilon}_i \sigma(1 - \alpha_i; k, x_0) \eta_{\text{noise}} \eta_{\text{xyz}}^{(t)}, \quad (1)$$

with $\boldsymbol{\varepsilon}_i \sim \mathcal{N}(0, \mathbf{I})$, $\sigma(\cdot)$ the opacity-gated logistic sigmoid, η_{noise} the published noise scale, and $\eta_{\text{xyz}}^{(t)}$ the spatial learning rate. Write $\Sigma_i = R_i \text{diag}(s_i^2) R_i^\top$ with $s_i \in \mathbb{R}^3$ the scale vector and $R_i \in \text{SO}(3)$ the rotation.

Primitive aspect. Define $\rho_i = s_i^{\max}/s_i^{\min}$. After convergence at a fixed cap, the empirical distribution of ρ_i is heavy-tailed: most primitives are near unity, but a small subset reaches the tens or higher. On outdoor scenes the high- ρ subset concentrates along thin features (branches, fence wires, blade edges, foliage silhouettes); on indoor scenes the corresponding subset is markedly smaller and the bulk of the surface is well-modelled near $\rho_i \approx 1$. This asymmetry is the mechanistic explanation of the per-scene-type split in Sec. 4.

3.2. Diagnosis: an exponential profile against a step

A Gaussian primitive’s contribution to a pixel falls off as $\exp(-\frac{1}{2} \mathbf{u}^\top \Sigma^{-1} \mathbf{u})$. For a primitive stretched along a thin feature, any transverse pixel offset receives a smooth exponential roll-off, while the feature’s true silhouette is closer

to a step. The sampler must either span the silhouette (paying a perceptual blur penalty) or shrink to lie strictly inside the feature (leaving neighbouring pixels uncovered). This is a structural mismatch that persists at convergence: the equilibrium of Kheradmand et al. [24] is stationary on configurations of Gaussian primitives, so it cannot escape the kernel-mismatch corner without enlarging the primitive class or replacing the offending primitives. We adopt the second option, restricted to a late-stage continuation pass.

3.3. The needle/capsule swap

Let T_{swap} denote the swap iteration, ρ^* the aspect-ratio threshold, and N the number of sub-primitives per parent. At iteration $t = T_{\text{swap}}$, for every primitive i with $\rho_i > \rho^*$, let $\mathbf{e}_i \in \mathbb{R}^3$ be the unit eigenvector of Σ_i at the largest scale s_i^{max} and place sub-primitive centres

$$\mathbf{p}_{i,k} = \mathbf{p}_i + \tau_k s_i^{\text{max}} \mathbf{e}_i, \quad \tau_k = \frac{2k - N - 1}{N - 1}, \quad (2)$$

uniformly spaced on $[-s_i^{\text{max}}, +s_i^{\text{max}}]$ along the long axis. Each sub-primitive inherits R_i and a scale vector contracted only along the long axis,

$$s_{i,k}^{\text{max}} = s_i^{\text{max}}/N, \quad s_{i,k}^{\text{mid}} = s_i^{\text{mid}}, \quad s_{i,k}^{\text{min}} = s_i^{\text{min}}, \quad (3)$$

producing capsule-like sub-primitives that together span approximately the same volume as the parent. Each sub-primitive copies α_i and $\mathbf{c}_i$ (full SH); the parent is then deleted. After the swap, training continues for $T_{\text{end}} - T_{\text{swap}}$ additional iterations under the unchanged update of Eq. (1), including the densification and pruning rules of Kheradmand et al. [24]. We use $T_{\text{swap}} = 0.8 \cdot T_{\text{end}}$ (iteration 24k of a 30k schedule), $\rho^* = 8$, and $N = 8$, reported as a single configuration; sensitivity is discussed in Sec. 4.5.

3.4. Why the equilibrium analysis still applies

The MCMC formulation of Kheradmand et al. [24] couples a relocation proposal $q(\cdot)$ that resamples primitives in proportion to a moment-matched density with a Langevin step that perturbs positions in place. Detailed balance is stated on the relocation proposal and is independent of the in-place dynamics. The swap is a one-time deterministic edit of the primitive set; it does not modify the relocation proposal, the opacity gate, the pruning rule, or the photometric likelihood, so it cannot break detailed balance. Within the post-swap continuation the same proposal and likelihood apply, and the chain re-equilibrates on the larger primitive set. That perceptual metrics improve while average PSNR is unchanged is consistent with the post-swap equilibrium being closer to the true geometry on thin features and approximately conservative on the smooth-region majority.

3.5. Pseudocode and integration footprint

Algorithm 1 summarises the swap. The per-iteration cost is unchanged for $t < T_{\text{swap}}$ and rises to a constant multiple

of the vanilla cost for $t \in [T_{\text{swap}}, T_{\text{end}}]$. Wall-clock cost is reported in Sec. 4.5.

Algorithm 1 NeedleSwap (changes vs. vanilla MCMC in bold).

```

1: input: primitives  $\{\mathbf{p}_i, \Sigma_i, \alpha_i, \mathbf{c}_i\}$ , iter  $t$ , swap iter  $T_{\text{swap}}$ ,
   threshold  $\rho^*$ , multiplicity  $N$ 
2: for each primitive  $i$  do
3:   perform standard MCMC update of Eq. (1)
4: end for
5: if  $t = T_{\text{swap}}$  then ▷ one-shot swap
6:   for each primitive  $i$  with  $s_i^{\text{max}}/s_i^{\text{min}} > \rho^*$  do
7:      $\mathbf{e}_i \leftarrow$  unit eigenvector of  $\Sigma_i$  at scale  $s_i^{\text{max}}$ 
8:     for  $k = 1, \dots, N$  do
9:        $\mathbf{p}_{i,k} \leftarrow \mathbf{p}_i + \tau_k s_i^{\text{max}} \mathbf{e}_i$  ▷ Eq. (2)
10:       $s_{i,k} \leftarrow (s_i^{\text{max}}/N, s_i^{\text{mid}}, s_i^{\text{min}})$  ▷ Eq. (3)
11:       $R_{i,k} \leftarrow R_i, \alpha_{i,k} \leftarrow \alpha_i, \mathbf{c}_{i,k} \leftarrow \mathbf{c}_i$  ▷ inherit
12:     end for
13:     delete primitive  $i$ 
14:   end for
15: end if
16: continue MCMC for  $t < T_{\text{end}}$  with no further changes.
```

4. Experiments

4.1. Setup

We build on the reference 3DGS-MCMC implementation of Kheradmand et al. [24]. NEEDLESWAP is a one-shot conditional swap at iteration T_{swap} (Sec. 3.3); all other components are inherited. We evaluate on Mip-NeRF 360 [1] with the every-eighth-image test split, 30k iterations, and $4\times$ input downsampling, matching Kerbl et al. [23] and Kheradmand et al. [24]. Both arms are trained internally at the same primitive cap on the same scene with the same fixed seed; the swap is performed at iteration 24k. All runs are single-seed on a single A100 (40 GB). We report all nine scenes: five outdoor (GARDEN, BICYCLE, FLOWERS, TREEHILL, STUMP) and four indoor (KITCHEN, ROOM, COUNTER, BONSAI). Per-scene deltas are NEEDLESWAP – vanilla MCMC *at the same primitive cap*. The primary baseline is vanilla 3DGS-MCMC retrained internally at `cap_max=1M`; we do not report comparisons against published 3DGS or Mip-Splatting numbers as “our reproduction.”

Noise prior. Following recent 3DGS literature, we state a conservative single-seed PSNR prior of order 0.1 dB on this benchmark and use it as a soft threshold for “within noise.” SSIM and LPIPS deltas are read against sign consistency across nine scenes.

Scene	Vanilla 3DGS-MCMC			NeedleSwap (ours)			$\Delta = \text{ours} - \text{vanilla}$		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Δ PSNR	Δ SSIM	Δ LPIPS
<i>Outdoor</i>									
garden	26.673	0.8267	0.1755	26.592	0.8412	0.1415	-0.081	+0.0145	-0.0340
bicycle	25.057	0.7282	0.2938	25.240	0.7560	0.2399	+0.183	+0.0279	-0.0539
flowers	21.648	0.5937	0.3713	21.729	0.6146	0.3268	+0.081	+0.0210	-0.0446
treehill	23.017	0.6447	0.3679	22.929	0.6541	0.3250	-0.088	+0.0094	-0.0430
stump	27.173	0.7897	0.2447	27.203	0.7975	0.2133	+0.030	+0.0078	-0.0314
outdoor avg (5)	24.714	0.7166	0.2906	24.738	0.7327	0.2493	+0.025	+0.0161	-0.0414
<i>Indoor</i>									
kitchen	31.575	0.9292	0.1303	31.262	0.9339	0.1167	-0.313	+0.0047	-0.0136
room	32.107	0.9342	0.1838	31.370	0.9324	0.1810	-0.737	-0.0017	-0.0028
counter	29.057	0.9190	0.1768	28.645	0.9188	0.1706	-0.412	-0.0002	-0.0062
bonsai	32.343	0.9483	0.1786	32.028	0.9486	0.1695	-0.315	+0.0003	-0.0090
indoor avg (4)	31.270	0.9327	0.1674	30.826	0.9334	0.1594	-0.444	+0.0008	-0.0079
nine-scene avg	27.628	0.8571	0.2362	27.444	0.8664	0.2098	-0.184	+0.0093	-0.0265

Table 1. Per-scene PSNR (dB), SSIM, and LPIPS on all nine Mip-NeRF 360 scenes at matched `cap_max=1M`, single seed, 30k iterations, single A100. Both arms are our own runs; deltas are NEEDLESWAP minus vanilla *at the same cap*. LPIPS improves on every scene; the gain is concentrated outdoors. PSNR is at parity outdoor and regresses indoor.

4.2. Nine-scene matched-cap result

Table 1 reports per-scene PSNR, SSIM, and LPIPS at `cap_max=1M`. LPIPS improves on every scene; the outdoor mean is -0.041 and the indoor mean is -0.008 , an outdoor gain roughly five times the indoor gain. SSIM gains on every outdoor scene and is unchanged indoors. Outdoor PSNR is at parity ($+0.025$ mean, mixed signs); indoor scenes regress (-0.444 mean, all four negative), pulling the nine-scene mean to -0.184 . We do not claim a PSNR improvement. The late-stage swap reallocates capacity from blurring across thin features to specialising at them, and the size of that effect tracks the size of the high-aspect primitive population.

4.3. Why the gain is outdoor-dominated

The mechanism predicts that NEEDLESWAP should help in proportion to the size of the population it acts on. The fraction of primitives with $\rho_i > 8$ at the swap iteration is consistently larger outdoors than indoors: outdoor geometry is dominated by thin elongated structures (branches, fence wires, foliage silhouettes), while indoor geometry is dominated by piecewise-smooth surfaces for which a near-isotropic Gaussian is already a good local match. On outdoor scenes the targeted population is large and silhouettes of thin features dominate LPIPS sensitivity, so the swap moves perceptual metrics noticeably. On indoor scenes the targeted population is small and structurally peripheral, and the cost (a roughly doubled active primitive count) shows up as a small mean PSNR regression on baselines that vanilla MCMC already fits well. The argument is mechanistic, not a controlled ablation.

Scene	1.0 M	2.0 M	Δ
garden	26.673	27.184	+0.511
bicycle	25.057	25.413	+0.356
flowers	21.648	21.957	+0.309
treehill	23.017	23.133	+0.116
stump	27.173	27.453	+0.280
mean	24.714	25.028	+0.314

Table 2. Vanilla 3DGS-MCMC PSNR (dB) at two primitive-cap settings on the five outdoor Mip-NeRF 360 scenes. Doubling the cap from 1M to 2M yields a mean PSNR gain of $+0.31$ dB without any algorithmic change. Indoor scenes saturate the wall-time envelope past `cap_max $\approx$ 1.2M` on our hardware and are excluded.

4.4. Cap-allocation evidence and the cap-mismatch pitfall

Table 2 reports vanilla MCMC PSNR at two primitive caps on the five outdoor scenes; doubling the cap alone yields $+0.31$ dB on average. During an audit of an earlier GARDEN sweep we found a comparison whose arms had been launched at different caps (candidate at 2M vs. vanilla at 1M). The reported delta was a multi-tenth-dB win; re-running at the same cap removed the entire delta. Any mechanism study on this backbone is therefore at risk of confounding “my mechanism” with “a slightly larger cap” if caps are not logged and matched on every cell. We measured cap-sensitivity only for the vanilla backbone, not for any specific candidate or published variant. We recommend that future mechanism studies on 3DGS-MCMC report the primitive cap of every cell and treat a per-scene PSNR im-




Figure 2. Per-scene PSNR, SSIM, and LPIPS for vanilla 3DGS-MCMC and NEEDLESWAP at matched `cap_max=1M` (same numbers as Table 1). LPIPS (right) improves on every scene; the gain is concentrated on the five outdoor scenes.



Figure 3. Iso-cap qualitative comparison on GARDEN test view 0 at `cap_max=1M`. PSNR is at parity (-0.081 dB) but LPIPS drops from 0.1755 to 0.1415, with the difference concentrated on foreground branches and grass blades whose silhouettes the late-stage swap recovers.



Figure 4. Iso-cap qualitative comparison on BICYCLE test view 0 at `cap_max=1M`. The largest perceptual gain in our table; LPIPS drops from 0.2938 to 0.2399, with the difference concentrated on the spokes, the chain, and the background fence.

provement of < 0.5 dB on outdoor Mip-NeRF 360 as cap-suspect until proven matched.

4.5. Sensitivity, scope, and cost

Hyperparameter sensitivity. On GARDEN at `cap_max=1M`, varying $\rho^* \in \{4, 16\}$ with $N=8$ and $T_{\text{swap}}=24k$ held fixed, PSNR remains within the 0.1 dB prior of the $\rho^*=8$ default and the LPIPS gain is preserved within 0.005. The result does not collapse under nearby threshold choices.

Why NEEDLESWAP and not other modifications. We surveyed twenty-two mechanism families against the same 3DGS-MCMC bottleneck, spanning relocation-proposal modifications, kernel and primitive-shape swaps, opacity and pruning rules, and appearance/loss reparametrisations.

Each was trained end-to-end on GARDEN at matched cap as a rejection probe; only families that matched or exceeded vanilla on GARDEN were considered for further evaluation, and only NEEDLESWAP was promoted (Table 3, Appendix 4.6). Full-loop modifications to the in-place dynamics or relocation proposal tend to land at near-parity in PSNR or to regress, consistent with the equilibrium analysis of Kheradmand et al. [24].

Compute cost. On GARDEN at `cap_max=1M` on a single A100, training time was approximately 1820 s for vanilla MCMC and 2010 s for NEEDLESWAP, a 10% overhead. Peak GPU memory rose from 14.2 GB to 18.6 GB. On indoor scenes the swap inflation is smaller and overhead is similar or slightly smaller.

Limitations. (i) Indoor PSNR regression of -0.444 dB on average; users for whom indoor pixel fidelity is a hard constraint should not enable the swap. (ii) Single seed; the noise prior is asserted from the literature. (iii) No comparison against published primitive-shape competitors [16, 17, 32]; we did not retrain those methods at matched cap. (iv) Per-scene and static; no dynamic or temporally consistent re-rendering. (v) Hyperparameter sweep is on GARDEN only. (vi) The high-aspect-fraction split of Sec. 4.3 is qualitative. (vii) Cap-mismatch genericity is shown for vanilla MCMC, not for any specific candidate or published variant.

4.6. Per-family appendix

Table 3 reports the single-scene rejection probe on GARDEN, full 30k iterations. Matched-cap deltas are taken against the vanilla baseline at the same cap (`cap=1M`: 26.673 dB; `cap=2M`: 27.184 dB).

5. Conclusion

We presented NEEDLESWAP, a late-stage modification to 3DGS-MCMC [24] that replaces high-aspect Gaussian

Family (idea)	Mechanism class	cap	garden PSNR	status
voxel- Φ proposal (0000)	new_strategy	1.0M	24.748	regression — 1.93 dB
sky-vMF primitive (0001)	new_primitive	0.1M	—	fail (no metric)
generalised-Gaussian β (0002)	new_primitive	—	—	fail (NaN)
sun-lobe + appearance MLP (0003)	new_loss_family	0.1M	—	fail (timeout)
ray-cone EWA (0004)	new_rendering_pipeline	2.0M	27.310	near-parity — 0.075 dB (cap=2M anchor)
KNN manifold loss (0005)	new_geometry_coupling	2.0M	—	fail (no metric)
PERMC photometric proposal (0006)	new_strategy	1.5M	26.998	off-cap probe
LockEta (0007)	new_strategy	1.0M	26.608	near-parity — 0.065 dB (sub-prior)
PerImageTone (0008)	new_strategy	5.2M	23.611	fail (divergent)
MH-Gate add_new_gs (0009)	new_strategy	1.0M	26.617	regression — 0.056 dB
homogeneous xyzw (0010)	new_primitive	7.0M	0.000	fail (zero PSNR)
per-image affine appearance (0011)	new_graph_topology	2.0M	19.071	fail (divergent)
DA-V2 dense init (0012)	new_strategy	0.6M	26.440	off-cap probe
β -rescaled Σ (0013)	new_compositing	1.5M	—	fail (no metric)
Laplacian-pyramid loss (0014)	new_loss_family	1.0M	25.896	regression — 0.777 dB
edge-Sobel structure-tensor (0015)	new_loss_family	1.0M	26.645	regression — 0.028 dB
Polyak-EMA SH distillation (0016)	new_strategy	1.0M	26.646	regression — 0.027 dB
JitterFlow (0017)	new_loss_family	1.0M	—	fail (no metric)
view-cond. residual head (0018)	new_graph_topology	1.0M	26.601	regression — 0.072 dB
capacity-tier KD (0019)	new_loss_family	—	—	fail (teacher infra)
NeedleSwap (0021)	new_primitive	1.0M	26.592	near-parity — 0.081 dB; promoted (perceptual)
SH-band direction gating (0022)	new_rendering_pipeline	1.0M	—	fail (no metric)
depth-prior pruning (0023)	new_strategy	1.0M	23.238	fail (divergent)
Cholesky reparam. at iter 24k (0026)	new_primitive	1.0M	26.542	regression — 0.131 dB
opacity-only adaptive SGLD (0027)	new_strategy	1.0M	26.678	near-parity +0.005 dB (sub-prior)

Table 3. Single-scene rejection probe on GARDEN, full 30k iterations. “regression” = converged but below vanilla at matched cap; “near-parity” = within the 0.1 dB single-seed prior; “fail” = no final metric. Only NEEDLESWAP (bold) was promoted to the full nine-scene matched-cap evaluation.

primitives at a fixed iteration after the sampler has converged with a chain of needle-shaped sub-primitives along their long axis, then continues training under the unchanged update rule. The construction is motivated by the structural mismatch between an exponential Gaussian profile and a step-discontinuity silhouette on thin scene features, adds no learnable parameters, and preserves the relocation-proposal moment-matching analysis of the original sampler.

On all nine Mip-NeRF 360 scenes at matched primitive cap, the modification improves LPIPS on every scene, with the gain concentrated on outdoor scenes (−0.0414 mean) and small on indoor scenes. The split tracks the density of high-aspect primitives in the converged sample. Indoor scenes pay a small mean PSNR cost for the post-swap primitive inflation on baselines that vanilla MCMC already fits well; we frame the contribution as an outdoor-targeted perceptual specialisation rather than a uniform improvement.

Numbers are single-seed and the noise prior is asserted from the literature. Vanilla 3DGS-MCMC alone is cap-sensitive on outdoor scenes (Table 2); during our audit we caught a comparison whose two arms had been launched at different primitive caps, producing an apparent multi-tenth-dB win that disappeared once the caps were matched. We recommend a matched-cap reporting discipline for downstream work in this neighbourhood.

References

- [1] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5470–5479, 2022. 2, 3, 4
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-NeRF: Anti-aliased grid-based neural radiance fields. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19697–19705, 2023. 3
- [3] Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. Revising densification in Gaussian Splatting. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [4] Dapeng Chen, Hai Li, Weicai Ye, Yifan Wang, Weijian Xie, Shangjin Zhai, Nan Wang, Haomin Liu, Hujun Bao, and Guofeng Zhang. PGSR: Planar-based Gaussian splatting for efficient and high-fidelity surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, 31(9):6100–6111, 2024. 3
- [5] Tianqi Chen, Emily B. Fox, and Carlos Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pages 1683–1691, 2014. 3
- [6] Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, and Xuejin Liu. GaussianPro: 3D Gaussian Splatting with progressive propagation. In *International Conference on Machine Learning (ICML)*, 2024. 3
- [7] Jaeyoung Chung, Jeongtaek Oh, and Kyoung Mu Lee. Depth-regularized optimization for 3D Gaussian Splatting in few-shot images. In *CVPR Workshop on 3D Vision and Modeling*, 2024. 3
- [8] Pinxuan Dai, Jiamin Xu, Wenxiang Xie, Xinguo Liu, Huamin Wang, and Weiwei Xu. High-quality surface reconstruction using Gaussian surfels. In *ACM SIGGRAPH 2024 Conference Proceedings*, 2024. 3
- [9] Xiaobin Deng, Changyu Diao, Min Li, Ruohan Yu, and Duanqing Xu. Efficient density control for 3D Gaussian Splatting, 2024. 3
- [10] Tianyu Ding et al. 3D Gaussian Splatting: Survey, technologies, challenges, and opportunities. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025. 3
- [11] Guangchi Fang and Bing Wang. Mini-Splatting: Representing scenes with a constrained number of gaussians. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [12] Umar Farooq, Jean-Yves Guillemaut, Adrian Hilton, and Marco Volino. Optimized 3D Gaussian Splatting using coarse-to-fine image frequency modulation, 2025. 3
- [13] Ben Fei et al. 3D Gaussian Splatting as a new era: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 31:4429–4451, 2025. 3
- [14] Glenn Grubert, Florian Barthel, Anna Hilsmann, and Peter Eisert. Improving adaptive density control for 3D Gaussian Splatting. In *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISAPP)*, 2025. 3
- [15] Antoine Guédon and Vincent Lepetit. SuGaR: Surface-aligned Gaussian Splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [16] Abdullah Hamdi, Luke Murai, Jinjie Shi, Florian Bernard, Matthias Nießner, Andrew J. Davison, and Federico Tombari. GES: Generalized exponential splatting for efficient radiance field rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 6
- [17] Jan Held, Renaud Vandeghen, Abdullah Hamdi, Adrien Deliege, Anthony Cioppa, Silvio Giancola, Andrea Vedaldi, Bernard Ghanem, and Marc Van Droogenbroeck. 3D convex splatting: Radiance field rendering with 3D smooth

- convexes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3, 6
- [18] Lukas Höllein et al. 3DGS-LM: Faster Gaussian-Splatting optimization with Levenberg-Marquardt. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 3
- [19] Roger Hsiao, Yuchen Fang, Xiangru Huang, Ruilong Li, Hesam Rabeti, Zan Gojic, Javad Lavaei, James Demmel, and Sophia Shao. 3DGS²-TR: Scalable second-order trust-region method for 3D Gaussian Splatting, 2026. 3
- [20] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D Gaussian Splatting for geometrically accurate radiance fields. *ACM Transactions on Graphics (SIGGRAPH)*, 43(4), 2024. 3
- [21] Ziyang Huang, Jiagang Chen, Jin Liu, and Shunping Ji. Opt3DGS: Optimizing 3D Gaussian Splatting with adaptive exploration and curvature-aware exploitation, 2025. 3
- [22] Junha Hyung, Susung Han, Sungwon Kim, Jaeseong Hong, and Jaegul Choi. Effective rank analysis and regularization for enhanced 3D Gaussian Splatting. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [23] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Dretakis. 3D Gaussian Splatting for real-time radiance field rendering. *ACM Transactions on Graphics (SIGGRAPH)*, 42(4):1–14, 2023. 2, 3, 4
- [24] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Jeff Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3D Gaussian Splatting as Markov Chain Monte Carlo. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 1, 2, 3, 4, 6
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 3
- [26] Scott Kirkpatrick, C. Daniel Gelatt, and Mario P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. 3
- [27] Jonas Kulhanek, Songyou Peng, Zuzana Kukelova, Marc Pollefeys, and Torsten Sattler. WildGaussians: 3D gaussian splatting in the wild. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [28] Yishai Lavi, Leo Segre, and Shai Avidan. Frequency-aware Gaussian Splatting decomposition, 2025. 3
- [29] Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. DNGaussian: Optimizing sparse-view 3D Gaussian radiance fields with global-local depth normalization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [30] Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Analytic-Splatting: Anti-aliased 3D Gaussian Splatting via analytic integration. In *European Conference on Computer Vision (ECCV)*, 2024. Oral. 3
- [31] Jiaqi Lin, Zhihao Li, Xiao Tang, Jianzhuang Liu, Shiyong Liu, Jiayue Liu, Yangdi Lu, Xiaofei Wu, Songcen Xu, Youliang Yan, and Wenming Yang. Vast-Gaussian: Vast 3D gaussians for large scene reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [32] Rong Liu, Dylan Sun, Meida Chen, Yue Wang, and Andrew Feng. Deformable beta splatting. *ACM SIGGRAPH 2025 Conference Papers*, 2025. 3, 6
- [33] Xinpeng Liu, Zeyi Huang, Fumio Okura, and Yasuyuki Matsushita. HoGS: Unified near and far object reconstruction via homogeneous gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [34] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-GS: Structured 3D Gaussians for view-adaptive rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. Highlight. 3
- [35] Yanzhe Lyu et al. ResGS: Residual densification of 3D Gaussian for efficient detail recovery. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 3
- [36] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the wild: Neural radiance fields for unconstrained photo collections. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7210–7219, 2021. 3
- [37] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 3
- [38] 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. In *European Conference on Computer Vision (ECCV)*, pages 405–421. Springer, 2020. 3
- [39] Seungtae Nam, Xiangyu Sun, Gyeongjin Kang, Younggeun Lee, Seungjun Oh, and Eunbyung Park. Generative densification: Learning to densify Gaussians for high-fidelity generalizable 3d reconstruction, 2024. 3
- [40] Arvind Neelakantan, Luke Vilnis, Quoc V. Le, Ilya Sutskever, Lukasz Kaiser, Karol Kurach, and James Martens. Adding gradient noise improves learning for very deep networks. In *International Conference on Learning Representations (ICLR) Workshop*, 2015. 3
- [41] Michael Niemeyer, Fabian Manhardt, et al. RadSplat: Radiance field-informed Gaussian Splatting for robust real-time rendering with 900+ fps, 2024. 3
- [42] Sam Patterson and Yee Whye Teh. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2013. 3
- [43] Hamza Pehlivan, Andrea Boscolo Camiletto, Lin Geng Foo, Marc Habermann, and Christian Theobalt. Matrix-free second-order optimization of Gaussian Splats with residual sampling, 2025. 3
- [44] Lukas Radl, Michael Steber, Mathias Parger, Bernhard Kerbl, and Markus Steinberger. StopThePop: Sorted Gaussian Splatting for view-consistent real-time rendering. *ACM Transactions on Graphics*, 43(4), 2024. 3
- [45] Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. Octree-GS: Towards consistent real-time rendering with LOD-structured 3D Gaussians. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2025. 3
- [46] Gareth O. Roberts and Richard L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996. 3
- [47] Matias Turkulainen, Xuqian Ren, Iaroslav Melekhov, Otto Seiskari, Esa Rahtu, and Juho Kannala. DN-Splatter: Depth and normal priors for gaussian splatting and meshing. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025. 3
- [48] Penguin Wang, Yueji Wang, Dongjin Wang, Shijia Mohan, Zhiwen Fan, Liang Wu, Ruijin Cai, Yung-Yu Yeh, Zhangyang Wang, Qixing Liu, and Rakesh Ranjan. Steepest descent density control for compact 3D Gaussian Splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [49] Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 681–688, 2011. 3
- [50] Guikun Wu and Jie Yi. A survey on 3D Gaussian Splatting. *arXiv preprint arXiv:2401.03890*, 2024. 3
- [51] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. DepthSplat: Connecting Gaussian Splatting and depth. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3
- [52] Zhiwen Yan, Weng Fei Low, Yu Chen, and Gim Hee Lee. Multi-scale 3D Gaussian Splatting for anti-aliased rendering. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [53] Ziyi Yang, Xinyu Gao, Yangtian Sun, Yihua Huang, Xiaoyang Lyu, Wen Zhou, Shaohui Jiao, Xiaojuan Qi, and Xiaogang Jin. Spec-Gaussian: Anisotropic view-dependent appearance for 3D Gaussian Splatting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 3
- [54] Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. AbsGS: Recovering fine details for 3D Gaussian Splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [55] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-Splatting: Alias-free 3D Gaussian Splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. Best Student Paper. 3
- [56] Zehao Yu, Torsten Sattler, and Andreas Geiger. Gaussian Opacity Fields: Efficient adaptive surface reconstruction in unbounded scenes. *ACM Transactions on Graphics*, 43(6), 2024. 3
- [57] Baowen Zhang, Chuan Fang, Rakesh Shrestha, Yixun Liang, Xiaoxiao Long, and Ping Tan. RaDe-GS: Rasterizing depth in Gaussian splatting. *ACM Transactions on Graphics*, 2024. 3
- [58] Jiahui Zhang, Fangneng Zhan, Muyu Xu, Shijian Lu, and Eric Xing. FreGS: 3D Gaussian Splatting with progressive frequency regularization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [59] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NeRF++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 3
- [60] Zheng Zhang, Wenbo Hu, Yixing Lao, Tong He, and Hengshuang Zhao. Pixel-GS: Density control with pixel-aware gradient for 3D Gaussian Splatting. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [61] Ziyi Zhang, Binbin Huang, Hanqing Jiang, Long Zhou, Xiuting Xiang, and Shenghua Shen. Quadratic Gaussian splatting: High quality surface reconstruction with second-order geometric primitives. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. 3
- [62] Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. FSGS: Real-time few-shot view synthesis using Gaussian Splatting. In *European Conference on Computer Vision (ECCV)*, 2024. 3
- [63] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. EWA splatting. *IEEE Transactions on Visualization and Computer Graphics*, 8(3): 223–238, 2002. 3