Horseshoe Splatting

Handling Structural Sparsity for Uncertainty-Aware Gaussian-Splatting Radiance Field Rendering

Feng Wu[1*], **Tsai Hor Chan**[2*], **Yihang Chen**[1], **Lingting Zhu**[1],
Guosheng Yin[1], **Lequan Yu**[1†]

[1] University of Hong Kong [2] University of Pennsylvania

[*] Equal contribution [†] Corresponding author

Core idea. We make 3D Gaussian Splatting **uncertainty-aware** by using a **global-local Horseshoe prior** on covariance scales, so noisy directions shrink while useful anisotropy remains.

Problem: fast 3DGS still lacks reliable confidence

Three practical gaps

- Existing 3DGS pipelines are mostly **deterministic**, so they cannot say where rendering is reliable.
- Per-splat covariance scales are not regularized for **structural sparsity**.
- Under **sparse views**, occlusion, or thin structures, noisy directions hurt both fidelity and trustworthiness.

Goal. Keep explicit Gaussian splatting fast while making it **sparsity-aware** and **uncertainty-aware**.

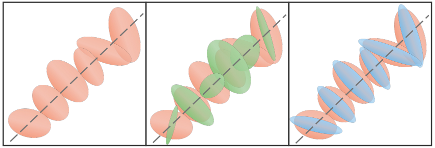
Optimal Representation Unstructured Noise Structured Noise

The key issue is **selective shrinkage**: suppress irrelevant covariance directions without destroying scene geometry.

Method overview: Horseshoe Splatting

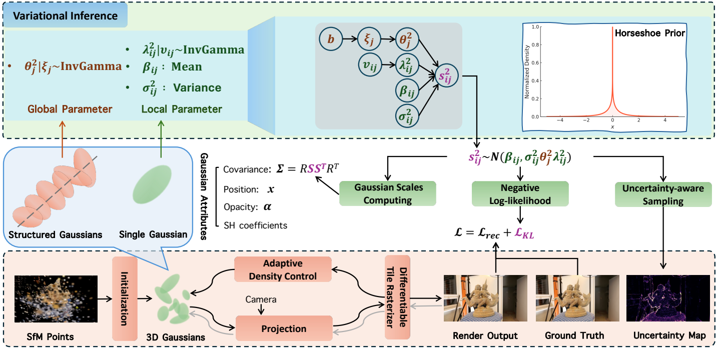

1) Structured prior. Put **global** and **local** shrinkage on each Gaussian scale.

2) Variational inference. Learn a factorized posterior with reconstruction loss plus KL regularization.

3) MC rendering. Sample from the posterior to obtain predictive mean and **pixel-wise uncertainty maps**.

Result. One training run yields a regularized splatting representation plus calibrated predictive uncertainty at test time.

Why Horseshoe instead of generic shrinkage

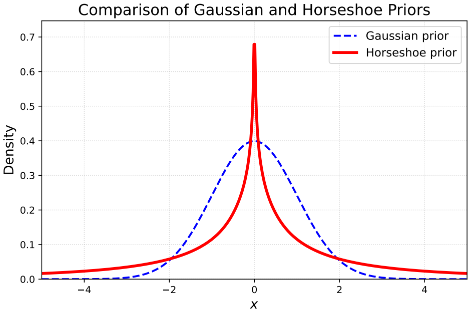

Horseshoe = **spike near zero** + **heavy tails**.

What this buys us

- **Spike near zero**: aggressively removes noise-dominated covariance scales.
- **Heavy tails**: keeps large, signal-bearing anisotropy instead of over-shrinking it.
- **Better calibration**: predictive variance tracks true reconstruction difficulty more faithfully.

Theory. More observations lead to lower estimation error and lower predictive uncertainty in image space.

Quantitative results: better fidelity and calibration together

LF dataset

Method	PSNR ↑	AUSE ↓	NLL ↓
FisherRF	29.13	0.54	7.02
Variational 3DGS	27.39	0.26	-0.30
Ours	**30.05**	**0.25**	**-0.74**

LLFF dataset

Method	PSNR ↑	AUSE ↓	NLL ↓
FisherRF	25.34	0.51	7.05
Variational 3DGS	23.97	0.32	0.23
Ours	**25.86**	**0.31**	**0.14**

LF depth uncertainty. We also reach the best average depth AUSE of **0.18**.

Takeaway. The uncertainty gain does **not** come at the cost of rendering quality. We improve calibration while also achieving the strongest or near-strongest novel-view fidelity on both benchmarks.

Qualitative comparison I: sharper novel-view rendering

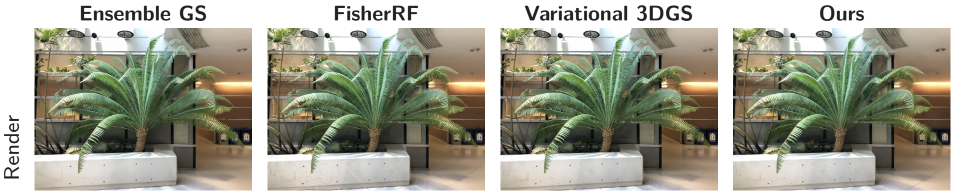

Visual evidence. On thin leaves and scene boundaries, our result stays **sharper and cleaner**, showing that selective shrinkage preserves useful structure instead of washing it out.

Qualitative comparison II: uncertainty aligns with actual error

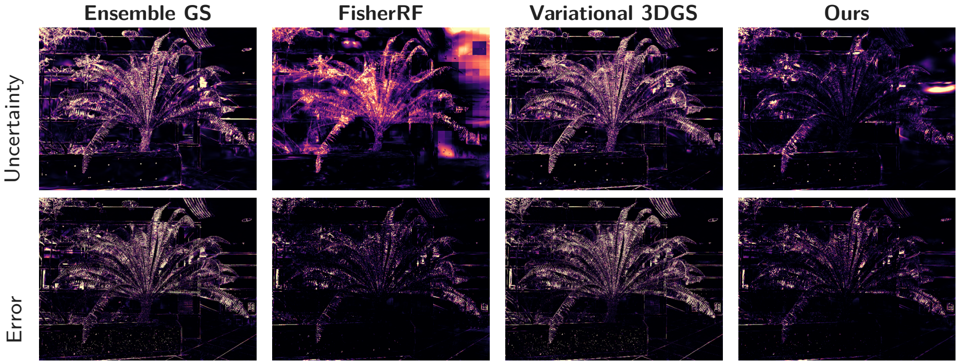

Uncertainty quality. Our uncertainty is **concentrated on genuinely hard regions** and visually matches the error pattern better, rather than spreading noisy confidence estimates across the image.

Practical impact and take-home message

Active view selection

26.23 PSNR

Best LLFF performance under the same acquisition budget.

Inference speed

0.03 s / view

Fastest among the uncertainty-aware 3DGS baselines we compare.

Prior ablation

-0.74 NLL

Much better than Laplace **10.58** and Gaussian **9.15** on LF.

Take-home message

- **Bayesian 3DGS**: a principled global-local Horseshoe hierarchy for covariance regularization.
- **Reliable uncertainty**: pixel-wise confidence from variational inference and Monte Carlo rendering.
- **Practical benefit**: better fidelity, calibration, view selection, and efficiency in one framework.

Thank you!

GitHub: `https://github.com/HKU-MedAI/Horseshoe-Splatting`