BIGUS: Beam Integrated Gaussian Scattering for Continuous Ultrasound View Synthesis

Youssif Abdelaziz*[1,2], Malak Ismail*[1,2], Nada Mourad*[1,2], Mohamed Badran[1,2], and Marwan Torki[1,2]

[1] Applied Innovation Center, Alexandria, Egypt
[2] Faculty of Engineering, Alexandria University, Egypt
*These authors contributed equally to this work.
{y.khaled,m.ismail,n.hossam,m.badran,mtorki}@aic.gov.eg

Abstract. Ultrasound imaging is widely used clinically due to its real-time capability, non-invasive nature, and low cost. Novel view synthesis in ultrasound allows better anatomical understanding, consistent evaluation across standard planes, and review of regions poorly captured during scanning. However, generating realistic novel views is challenging due to the prevalence of speckle patterns arising from complex wave–tissue interactions. These fine-grained textures often carry diagnostic significance, yet existing rendering methods frequently oversmooth or distort speckle, limiting performance in downstream tasks that rely on accurate speckle preservation. In this paper, we propose BIGUS, a 3D Gaussian Splatting–based framework for ultrasound novel view synthesis. We redesign the Gaussian rendering pipeline with a physics-aware scattering formulation that better captures the characteristics of ultrasound physics. In addition, we introduce a structure-aware initialization strategy that adaptively samples Gaussian primitives from tissue regions, combined with an error-guided densification mechanism to improve reconstruction fidelity. We evaluate BIGUS on three diverse ultrasound datasets, including fetal brain, forearm, and brain imaging, each with distinct speckle patterns and probe geometries. To better assess speckle preservation, we propose the Backscatter Fidelity Score, a speckle-sensitive evaluation metric that complements conventional image quality measures. Our experiments show that BIGUS achieves **67.26 dB** PSNR on the fetal brain dataset, establishing state-of-the-art performance with only **140k** Gaussians on a single GPU and training in under **15 minutes**. Moreover, it outperforms existing methods by up to **9 dB** PSNR on speckle-dominant datasets. The source code is available at https://github.com/YoussifKhaled/BIGUS.

Keywords: Novel View Synthesis · Gaussian Splatting · Ultrasound Imaging.

1 Introduction

Ultrasound is the most widely used medical imaging modality, valued for its real-time capability, portability, absence of ionizing radiation, and low cost [24].

Clinically, it operates in two regimes: conventional 2D imaging acquires cross-sectional slices in real time but provides only partial anatomical views, requiring the operator to mentally reconstruct 3D structures, while dedicated 3D probes enable multiplanar review at the cost of higher cost and limited availability [11].

Both limitations can be overcome through novel view synthesis (NVS), in which a volumetric representation is learned from tracked 2D acquisitions and queried at arbitrary poses to produce new slices [8, 9]. Such views enable downstream tasks, including lesion tracking across standardized planes [4], scan simulation for operator training [1], and synthetic data generation [18].

Learned volumetric representations offer a principled path toward this goal. Neural Radiance Fields (NeRF) [20] optimize an implicit mapping from coordinates to radiance via differentiable volume rendering, achieving high visual fidelity at a significant computational cost. 3D Gaussian Splatting (3DGS) [14] replaces the implicit field with an explicit set of anisotropic Gaussian primitives rasterized in real time, attaining comparable quality orders of magnitude faster.

Adapting these representations to the acoustic domain, Ultra-NeRF [27] integrates tissue-specific scattering and attenuation into the NeRF pipeline for B-mode synthesis, while UltraGauss [6] introduces probe-plane intersection of Gaussian primitives for real-time rendering. Both demonstrate that learned volumetric representations can produce realistic ultrasound views, yet neither explicitly models the speckle pattern arising from sub-resolution wave interference.

Speckle importance comes from its statistics, which carry tissue-specific diagnostic information used in characterization [3, 5], its decorrelation encodes inter-frame displacement for sensorless pose estimation [16], and its texture informs segmentation boundaries [23]. Preserving speckle fidelity is therefore a prerequisite for clinically useful reconstruction.

In this work, we present **BIGUS** (Beam Integrated Gaussian Scattering), a framework that reconstructs ultrasound imagery while faithfully preserving speckle, illustrated in Figure 1. We summarize our main contributions as follows:

1. **Signed beam-integrated Gaussian rendering:** A differentiable rasterizer with signed reflectivities, exact elevation-beam integration, implemented via custom CUDA kernels.
2. **Gaussian Field Construction and Refinement:** Using slice-based Harris initialization, and error-attributed densification that focuses on high-uncertainty regions.
3. **Multi-anatomy evaluation with speckle-aware assessment:** Experiments on three anatomically distinct datasets, in addition to a novel metric that captures speckle-distribution similarity.

2 Method

2.1 Preliminaries

B-mode image formation. Three physical phenomena underlie B-mode ultrasound image formation. First, pixels integrate backscatters across a volumetric

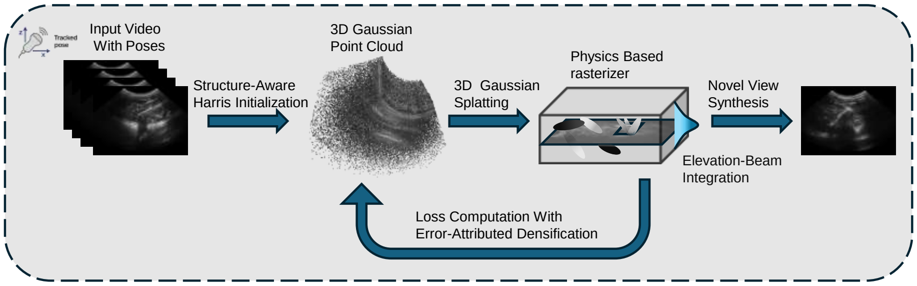

Fig. 1. BIGUS pipeline overview. Structure-aware initialization builds a 3D Gaussian representation from US frames, then it's optimized with error-guided densification and rendered via physics-based beam integration for NVS.

slab rather than sampling on a single plane, as the imaging beam has a finite elevation thickness [24, 12]. Second, scattered transmitted pulses within each resolution cell echo as signed pressure waves, producing the constructive and destructive interference that manifests as speckles [24]. Third, the magnitude of the resulting RF signal is extracted, compressed and mapped to a greyscale image [24].

3D Gaussian Splatting (3DGS) [14] represents a scene with N anisotropic Gaussian primitives, each carrying a mean $\boldsymbol{\mu}_i \in \mathbb{R}^3$, a unit quaternion $\mathbf{q}_i$ encoding a rotation $\mathbf{R}_i$, and a scale matrix $\mathbf{S}_i$. The covariance factorizes as

$$\boldsymbol{\Sigma}_i = \mathbf{R}_i\, \mathbf{S}_i \mathbf{S}_i^\top \mathbf{R}_i^\top \tag{1}$$

UltraGauss [6] adapts this representation for ultrasound by replacing camera projection with probe-plane intersection. Given a pose, each Gaussian is transformed to the probe frame with mean $\boldsymbol{\mu}_i^{(p)}$ and covariance $\boldsymbol{\Sigma}_i^{(p)}$, and each pixel $\mathbf{x}$ is lifted to $\mathbf{x}_{|0} = [x_1, x_2, 0]^\top$. The contribution of Gaussian i is evaluated as

$$G_i(\mathbf{x}) = \exp\Big(-\tfrac{1}{2}\left(\mathbf{x}_{|0} - \boldsymbol{\mu}_i^{(p)}\right)^\top \left(\boldsymbol{\Sigma}_i^{(p)}\right)^{-1} \left(\mathbf{x}_{|0} - \boldsymbol{\mu}_i^{(p)}\right)\Big) \tag{2}$$

and weighted contributions are accumulated into a normalized pixel intensity.

2.2 BIGUS Renderer Formulation

Previous works [6] fall short in two key aspects of US image formation: modeling the scan plane as infinitely thin Eq. (2) at z=0, ignoring the elevation beam thickness, and using non-negative I_i that cannot capture destructive interference, thus failing to reproduce the full speckle pattern. We base our rasterizer formulation depicted as a component in Fig. 1 on addressing these limitations, through three key design choices: a Gaussian parameterization that facilitates beam integration, signed reflectivities that capture interference, and an exact closed-form solution to the elevation beam integral.

Gaussian parameterization We retain a formulation of parameters similar to [14] as mentioned in Section 2.1. Building upon Eq. (1), we focus on the precision $\Lambda_i = \Sigma_i^{-1}$, which is more convenient for rendering and optimization.

$$\Lambda_i = R_i \, S_i^{-2} \, R_i^\top . \tag{3}$$

Signed reflectivities Within each resolution cell, sub-wavelength echoes combine as complex phasors $\sum_k a_k \, e^{j\phi_k}$, interfering constructively where phases align and destructively where they oppose [24]. A system restricted to non-negative contributions drives Gaussians toward pixel-sized primitives and cannot fully reproduce speckle contrast. We collapse each phasor to a signed scalar $r_i \in \mathbb{R}$ that absorbs both magnitude and phase sign, enabling constructive and destructive accumulation without the optimization difficulties a periodic phase variable would introduce.

Elevation beam integration Each 2D observation is inherently a volumetric measurement, due to the finite thickness of the ultrasound beam in the elevation direction. Modeling this integration strengthens cross-plane coherence and improves novel-view fidelity beyond what a thin-plane evaluation can achieve.

For a pixel at $\mathbf{p} = (u, v)$, and Gaussian μ_i let $\mathbf{d} = [\mathbf{d}_{xy}; d_z]$ denote the displacement, with $\mathbf{d}_{xy} = (u - \mu_x, \; v - \mu_y)^\top$ and $d_z = z - \mu_z$. And with elevation beam modeled as a Gaussian profile with precision $\lambda_E = 1/\sigma_E^2$.

We found the elevation beam integral, which sums each Gaussian's response over the full elevation range, weighted by the Gaussian beam profile.

$$w_i(\mathbf{p}) \;=\; \int_{-\infty}^{+\infty} \underbrace{\exp\!\Big(-\tfrac{1}{2}\,\mathbf{d}^\top \Lambda_i^{(p)}\,\mathbf{d}\Big)}_{\text{3D Gaussian}} \; \underbrace{\exp\!\Big(-\tfrac{\lambda_E}{2}\,z^2\Big)}_{\text{beam profile}} \, dz \tag{4}$$

We follow by partitionning the precision $\Lambda_i^{(p)}$ into an in-plane block $A \in \mathbb{R}^{2\times 2}$, a cross-coupling vector $\mathbf{b} \in \mathbb{R}^2$, and an elevation precision scalar c, and Expanding the quadratic form via this block structure gives

$$\mathbf{d}^\top \Lambda_i^{(p)}\,\mathbf{d} \;=\; \big(\mathbf{d}_{xy}^\top \; d_z\big) \begin{pmatrix} A & \mathbf{b} \\ \mathbf{b}^\top & c \end{pmatrix} \begin{pmatrix} \mathbf{d}_{xy} \\ d_z \end{pmatrix} \;=\; \mathbf{d}_{xy}^\top A\,\mathbf{d}_{xy} + 2\,\mathbf{b}^\top \mathbf{d}_{xy}\,d_z + c\,d_z^2 \tag{5}$$

Substituting $z = d_z + \mu_z$ in the beam term and merging both exponents of Eq. (4), the combined exponent E collects as a quadratic in d_z:

$$E \;=\; -\tfrac{1}{2}\Big[\underbrace{(c + \lambda_E)}_{\tilde{c}}\,d_z^2 \;+\; 2\,\underbrace{(\mathbf{b}^\top \mathbf{d}_{xy} + \lambda_E \mu_z)}_{\beta}\,d_z \;+\; \mathbf{d}_{xy}^\top A\,\mathbf{d}_{xy} + \lambda_E \mu_z^2\Big] \tag{6}$$

By completing the square, we isolate a shifted Gaussian in d_z from a residual that depends only on $\mathbf{d}_{xy}$. The integral therefore factors as

$$w_i(\mathbf{p}) \;=\; \exp\!\Big(\tfrac{\beta^2}{2\tilde{c}} - \tfrac{1}{2}\,\mathbf{d}_{xy}^\top A\,\mathbf{d}_{xy} - \tfrac{\lambda_E \mu_z^2}{2}\Big) \;\cdot\; \int_{-\infty}^{+\infty} \exp\!\Big(-\tfrac{\tilde{c}}{2}\Big(d_z + \tfrac{\beta}{\tilde{c}}\Big)^2\Big)\,dz \tag{7}$$

Expanding $\beta^2/\tilde{c}$ by powers of $\mathbf{d}_{xy}$ and using the fact that standard Gaussian integrals evaluate to $\int e^{-\alpha x^2/2}dx = \sqrt{2\pi/\alpha}$, we arrive at a closed-form expression for the beam-integrated weight:

$$w_i(\mathbf{p}) \;=\; \kappa_{0,i} \;\cdot\; \exp\!\Big(-\tfrac{1}{2}\,\mathbf{d}_{xy}^\top \tilde{A}_i\,\mathbf{d}_{xy} \;+\; \boldsymbol{\ell}_i^\top \mathbf{d}_{xy}\Big) \tag{8}$$

with three quantities produced by the marginalization:

$$\tilde{A} = A - \frac{1}{\tilde{c}}\,\mathbf{b}\mathbf{b}^\top \qquad \text{(Schur complement)} \tag{9}$$

$$\kappa_0 = \sqrt{\frac{2\pi}{\tilde{c}}}\;\exp\!\left(-\frac{c\,\lambda_E}{2\,\tilde{c}}\,\mu_z^2\right) \qquad \text{(elevation decay)} \tag{10}$$

$$\boldsymbol{\ell} = \frac{\lambda_E\,\mu_z}{\tilde{c}}\,\mathbf{b} \qquad \text{(parallax shift)} \tag{11}$$

The Schur complement $\tilde{A}$ is the 2D precision after marginalizing out elevation. The scalar κ_0 penalizes Gaussians centered far from the scan plane, with joint decay governed by both the primitive and beam precisions. The parallax vector $\boldsymbol{\ell}$ shifts the apparent 2D peak when a tilted Gaussian sits off-plane. No approximation is involved: Eq. (8) with (9)–(11) are an exact closed-form solution to the elevation integral in Eq. (4), and were efficiently implemented into two-phase custom CUDA kernels.

Rendering equation The beam-integrated weights are assembled into a signed scattering field and mapped to pixel intensity through a two-stage transfer:

$$S(\mathbf{p}) = \textstyle\sum_{i=1}^{M} r_i\; w_i(\mathbf{p}) \tag{12}$$

$$\hat{I}(\mathbf{p}) = 1 - \exp\!\big(-\gamma \cdot \max\!\big(S(\mathbf{p}),\,0\big)\big) \tag{13}$$

The summation is order-independent due to absence of geometric occlusion in US imaging, the depth sort required by 3DGS can be removed [14]. Signed reflectivities allow S to be negative where destructive interference dominates; the rectifier zeroes these regions, mirroring envelope detection, while the saturating map $1{-}e^{-x}$ compresses strong echoes in analogy with B-mode log-compression. Because S is linear in $\{r_i\}$, the image is invariant under $(r_i,\gamma)\to(\alpha r_i,\gamma/\alpha)$; we fix $\gamma = -\ln(1{-}\bar{I}_{\mathrm{GT}})$ so that at convergence the mean scattering field across slices is unity, placing reflectivities in a stable numerical range, and maintaining a more stable early optimization. The full pipeline from Eq. (8) to (13) is differentiable, enabling end-to-end optimization from 2D US frames.

2.3 Gaussian Field Construction and Refinement

Structure-Aware Harris Initialization In practice, suboptimal mean initialization may render primitives effectively inactive before they can migrate toward

Table 1. Reconstruction fidelity on the fetal brain dataset under two input densities.

		5 mins		**Convergence**	
Training Set	**Method**	**PSNR ↑**	**SSIM ↑**	**PSNR ↑**	**SSIM ↑**
160 Linearly Spaced Slices	ImplicitVol [28]	26.44	0.793	39.36	0.96
	RapidVol [7]	15.96	0.240	42.01	0.982
	UltraNeRF [27]	19.39	0.494	22.78	0.679
	UltraGauss-100k [6]	43.07	0.986	45.75	0.992
	UltraGauss-2M [6]	41.17	0.982	54.26	1.000
	BIGUS-140k (ours)	**58.08**	**0.998**	**67.26**	**1.000**
80 Linearly Spaced Slices	ImplicitVol [28]	24.72	0.752	26.13	0.745
	RapidVol [7]	15.88	0.248	32.13	0.924
	UltraNeRF [27]	18.49	0.428	23.42	0.688
	UltraGauss-100k [6]	39.99	0.980	38.59	0.974
	UltraGauss-2M [6]	39.63	0.981	44.40	0.991
	BIGUS-140k (ours)	**51.65**	**0.999**	**60.72**	**1.000**

anatomically meaningful regions. Prior ultrasound methods initialize means either by distributing them uniformly across slices [15], or within a predefined 3D volume [6] this delays convergence. We instead propose *Structure-Aware Harris Initialization*. For each slice, we compute the Harris corner response [10] to identify high-gradient and speckle-rich anatomical structures. Seed locations sampled from this response distribution define the initial Gaussian mean configuration. By concentrating primitives at structurally informative regions, Structure-Aware Harris Initialization provides early convergence.

Error-Attributed Adaptive Densification(EA-D) In light-based 3DGS, adaptive densification is commonly driven by image-space residuals or primitive gradient statistics [14]. However, positional-gradient-based strategies are sensitive to threshold tuning and may fail to activate in persistently underfit regions, especially when gradients vanish despite remaining reconstruction error [22]. This issue is amplified in our setting: under signed, order-independent linear accumulation, $\|\nabla_{\mathbf{x}}\mathcal{L}\|$ rapidly diminishes once primitives stabilize spatially, even in speckle-rich or boundary regions that remain imperfectly reconstructed. To address this limitation, instead of relying on instantaneous position gradients, EA-D selects primitives based on their attributed responsibility for residual error. Using pixel-wise gradients with respect to the signed scattering field, each Gaussian i receives an error score E_i weighted by its beam-integrated footprint:

$$E_i = \sum_p \left| \frac{\partial \mathcal{L}}{\partial S(p)} \right| \Phi_i(p) \tag{14}$$

Primitives in the top percentile of normalized E_i are selected for densification. To maintain stability, EA-D applies controlled structural updates: small primitives are cloned with scaled reflectivity, large ones are split along a descent direction, and inactive or negligible primitives are pruned. This reallocates capacity to persistently underfit regions while preserving stable optimization.

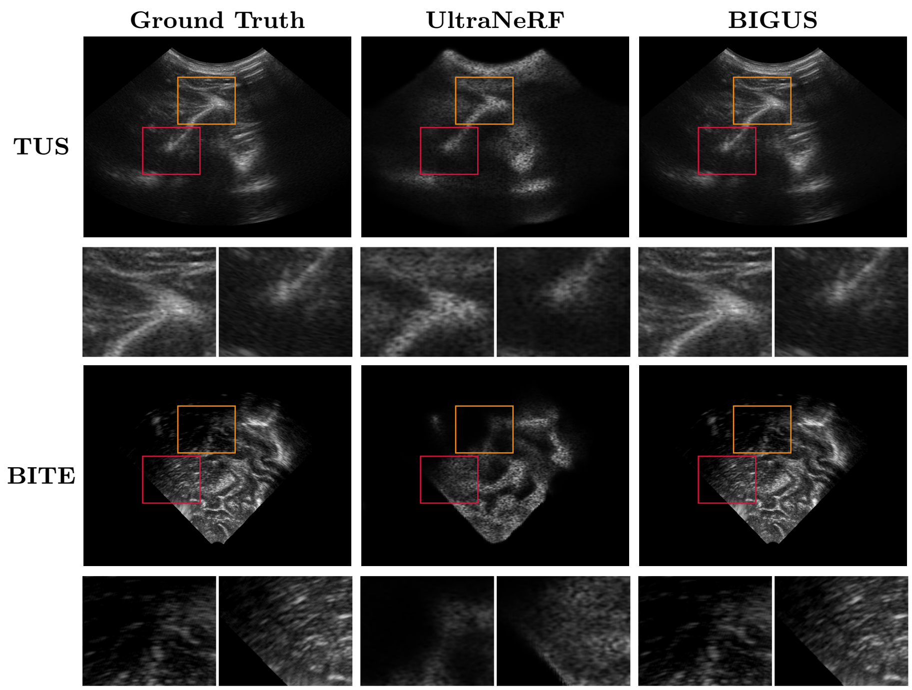

Fig. 2. Visual comparison of zoomed crops highlighting speckle fidelity. BIGUS matches ground truth speckle while UltraNeRF over-smooths diagnostic texture.

3 Experiments

3.1 Datasets

We evaluate on three datasets spanning different anatomies and speckle regimes. **Fetal Brain** from the INTERGROWTH-21st study [21] serves as our primary benchmark for direct comparison with prior work. Both **TUS** from TUS-REC2024 [17], a freehand forearm sweep benchmark, and **BITE** [19], containing brain tumour volumes, exhibit stronger speckles, allowing us to test robustness.

3.2 Evaluation Metrics

We report standard image-quality metrics: PSNR and SSIM. These metrics don't capture speckle statistics, so we introduce the Backscattering Fidelity Score (BFS) to compare the speckle distributions between the reconstructed and ground-truth images. BFS is computed by extracting speckle-contrast values over small overlapping patches, forming two distributions (GT versus reconstruction).

$$\mathrm{BFS} = 0.4\, s_\mu + 0.4\, \exp(-W_1/\sigma_W) + 0.2\, s_{\mathrm{Ray}} \tag{15}$$

where s_μ measures the similarity of the mean speckle contrast [2], W_1 measures distributional mismatch, and s_{Ray} rewards matching the expected speckle

Table 2. Reconstruction and speckle fidelity on TUS and BITE. **Bold** indicates best.

Dataset	Model	PSNR↑	SSIM↑	BFS↑
TUS	UltraNeRF	24.03±1.73	0.711±0.068	0.686±0.045
	BIGUS	**34.40**±0.06	**0.917**±0.003	**0.805**±0.001
BITE	UltraNeRF	19.41±1.16	0.770±0.022	0.630 ±0.065
	BIGUS	**34.7**±1.51	**0.965**±0.084	**0.871**±0.019

Table 3. Ablation study on BITE and TUS. **Bold** indicates best.

Init	DSSIM +FFL	Dens	#GS	BITE			TUS		
				PSNR	SSIM	BFS	PSNR	SSIM	BFS
random3d	–	–	100K	26.12	0.8663	0.7779	29.98	0.8249	0.7959
harris	–	–	100K	30.27	0.9285	0.8744	31.14	0.8478	0.8156
harris	–	–	2M	33.82	0.9571	0.9084	34.16	0.9020	0.9089
harris	✓	–	2M	34.28	0.9626	0.9166	34.85	0.9118	0.9320
harris	–	✓	2M	34.39	0.9606	0.9122	35.21	0.9136	0.9439
harris	✓	✓	2M	**35.02**	**0.9661**	**0.9202**	**36.38**	**0.9244**	**0.9722**

regime [25]. The score lies between 0 and 1, with higher values indicating better speckle fidelity.

3.3 Results Discussion

Slices were rendered via custom CUDA kernels on a single NVIDIA V100 GPU. Table 1 reports results on the fetal brain benchmark. Using only 140k Gaussians, BIGUS achieves state-of-the-art performance in two experimental settings: one stopped at 5 minutes, and one run to full convergence which took under 15 minutes for BIGUS. In both cases, it surpasses prior work. To assess speckle fidelity, we evaluate the more challenging TUS and BITE datasets (Table 2) and compare directly with UltraNeRF [27]. BIGUS outperforms UltraNeRF across all three metrics on both datasets, and the qualitative comparison in Fig. 2 shows that it preserves the granular appearance of the ground truth.

3.4 Ablation Study

We conduct ablations on representative samples from TUS and BITE to isolate the impact of different design choices under controlled conditions, as shown in Table 3. We first compare uniform 3D initialization against our slice-based Harris initialization, which performs better on both datasets. The base configuration uses Harris initialization and L1 loss. Increasing the number of Gaussians from 100K to 2M substantially improves performance on both datasets. Adding DSSIM and FFL losses [26, 13] further enhances structural fidelity, reflected in consistent SSIM gains. EA-D provides additional improvements, particularly on TUS. The final configuration achieves the best overall results on both datasets.

Table 4. Ablation study on Signed r (positive-only r when off) and elevation-beam integration ($\sigma_E \to 0$ thin-plane when off) on BITE and TUS. **Bold** indicates best.

Signed r	**Elev**	**TUS**			**BITE**		
		PSNR ↑	SSIM ↑	BFS ↑	PSNR ↑	SSIM ↑	BFS ↑
–	–	29.47	0.8138	0.6539	25.75	0.8693	0.7266
✓	–	29.91	0.8248	0.6754	26.46	0.8861	0.7519
–	✓	28.67	0.7999	0.6704	31.47	0.9412	0.8996
✓	✓	**34.14**	**0.9057**	**0.8501**	**33.68**	**0.9600**	**0.9196**

Additionally, Table 4 presents an ablation study of the proposed rasterizer components, further demonstrating the importance of each component.

4 Conclusion

In this paper, we presented BIGUS, a Gaussian splatting framework that approximates ultrasound physics for ultrasound NVS. We designed a differentiable rendering pipeline with signed reflectivities and exact elevation beam integration. Additionally, we complemented the rasterizer with Harris Based Initialization and EA-Adaptive Densification to concentrate representational capacity at more informative regions. Experiments on three anatomically diverse datasets demonstrate that BIGUS consistently outperforms prior baselines in reconstruction quality and speckle fidelity, as measured by our proposed BFS metric.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Blum, T., Rieger, A., Navab, N., Friess, H., Martignoni, M.: A review of computer-based simulators for ultrasound training. Simulation in Healthcare **8**(2), 98–108 (2013)
2. Boas, D.A., Dunn, A.K.: Laser speckle contrast imaging in biomedical optics. Journal of biomedical optics **15**(1), 011109–011109 (2010)
3. Burckhardt, C.B.: Speckle in ultrasound b-mode scans. IEEE Transactions on Sonics and ultrasonics **25**(1), 1–6 (2005)
4. Che, H., Ramanathan, S., Foran, D.J., Nosher, J.L., Patel, V.M., Hacihaliloglu, I.: Realistic ultrasound image synthesis for improved classification of liver disease. In: International Workshop on Advances in Simplifying Medical Ultrasound. pp. 179–188. Springer (2021)
5. Damerjian, V., Tankyevych, O., Souag, N., Petit, E.: Speckle characterization methods in ultrasound images–a review. Irbm **35**(4), 202–213 (2014)
6. Eid, M.C., Namburete, A., Henriques, J.F.: Ultragauss: Ultrafast gaussian reconstruction of 3d ultrasound volumes. In: The Fourteenth International Conference on Learning Representations (2026), https://openreview.net/forum?id=0imrI7UXdu

7. Eid, M.C., Yeung, P.H., Wyburd, M.K., Henriques, J.F., Namburete, A.I.: Rapidvol: rapid reconstruction of 3d ultrasound volumes from sensorless 2d scans. In: 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2025)
8. Gaits, F., Mellado, N., Basarab, A.: Ultrasound volume reconstruction from 2d freehand acquisitions using neural implicit representations. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024)
9. Gaits, F., Vidal, F., Basarab, A., Mellado, N.: Efficient neural reconstruction for freehand 3-d ultrasound imaging and visualization in augmented reality. IEEE Access **12**, 166253–166267 (2024)
10. Harris, C., Stephens, M., et al.: A combined corner and edge detector. In: Alvey vision conference. vol. 15, pp. 10–5244. Manchester, UK (1988)
11. Huang, Q., Zeng, Z.: A review on real-time 3d ultrasound imaging technology. BioMed research international **2017**(1), 6027029 (2017)
12. Jensen, J.A., Svendsen, N.B.: Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers. IEEE transactions on ultrasonics, ferroelectrics, and frequency control **39**(2), 262–267 (1992)
13. Jiang, L., Dai, B., Wu, W., Loy, C.C.: Focal frequency loss for image reconstruction and synthesis. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 13919–13929 (2021)
14. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., et al.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. **42**(4), 139–1 (2023)
15. Krähenmann, M., Tascon-Morales, S., Laumer, F., Vogt, J.E., Ozkan, E.: From slices to structures: Unsupervised 3d reconstruction of female pelvic anatomy from freehand transvaginal ultrasound. arXiv preprint arXiv:2508.14552 (2025)
16. Laporte, C., Arbel, T.: Probabilistic speckle decorrelation for 3d ultrasound. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 925–932. Springer (2007)
17. Li, Q., Saeed, S.U., Huang, Y., Luo, M., Yan, Z., Chen, J., Yang, X., Ni, D., Winter, N., Nguyen, P., et al.: Tus-rec2024: A challenge to reconstruct 3d freehand ultrasound without external tracker. arXiv preprint arXiv:2506.21765 (2025)
18. Liang, J., Yang, X., Huang, Y., Li, H., He, S., Hu, X., Chen, Z., Xue, W., Cheng, J., Ni, D.: Sketch guided and progressive growing gan for realistic and editable ultrasound image synthesis. Medical image analysis **79**, 102461 (2022)
19. Mercier, L., Del Maestro, R.F., Petrecca, K., Araujo, D., Haegelen, C., Collins, D.L.: Online database of clinical mr and ultrasound images of brain tumors. Medical physics **39**(6Part1), 3253–3261 (2012)
20. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021)
21. Papageorghiou, A.T., Kennedy, S.H., Salomon, L.J., Altman, D.G., Ohuma, E.O., Stones, W., Gravett, M.G., Barros, F.C., Victora, C., Purwar, M., et al.: The intergrowth-21st fetal growth standards: toward the global integration of pregnancy and pediatric care. American journal of obstetrics and gynecology **218**(2), S630–S640 (2018)
22. Rota Bulò, S., Porzi, L., Kontschieder, P.: Revising densification in gaussian splatting. In: European Conference on Computer Vision. pp. 347–362. Springer (2024)
23. Slabaugh, G., Unal, G., Wels, M., Fang, T., Rao, B.: Statistical region-based segmentation of ultrasound images. Ultrasound in medicine & biology **35**(5), 781–795 (2009)

24. Szabo, T.L.: Diagnostic ultrasound imaging: inside out. Academic press (2013)
25. Wagner, R.F., Smith, S.W., Sandrik, J.M., Lopez, H.: Statistics of speckle in ultrasound b-scans. IEEE Transactions on sonics and ultrasonics **30**(3), 156–163 (2005)
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing **13**(4), 600–612 (2004)
27. Wysocki, M., Azampour, M.F., Eilers, C., Busam, B., Salehi, M., Navab, N.: Ultranerf: Neural radiance fields for ultrasound imaging. In: Medical Imaging with Deep Learning. pp. 382–401. PMLR (2024)
28. Yeung, P.H., Hesse, L.S., Aliasi, M., Haak, M.C., Xie, W., Namburete, A.I., 21st Consortium, I., et al.: Sensorless volumetric reconstruction of fetal brain freehand ultrasound scans with deep implicit representation. Medical Image Analysis **94**, 103147 (2024)