This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

G^{3}R: Gaussian-based Geometry-Guided Reconstruction for Fetal Brain MRI

Zhibao Cai[1], Yao Lv[1], Xin Zhang[1(✉)], and Chaoxiang Yang[2]

[1] School of Electronic and Information Engineering, South China University of Technology, Guangzhou, Guangdong, China
eexinzhang@scut.edu.cn

[2] Department of Radiology, Guangdong Women and Children Hospital, Guangzhou, Guangdong, China

Abstract. Reconstructing 3D fetal brain MRI volume from 2D slices is crucial for prenatal assessment but challenged by unpredictable fetal motion and large inter-slice gap. While 3D Gaussian Splatting (3DGS) has recently enabled rapid slice-to-volume reconstruction (SVR), the optimization of independent Gaussian primitives lacks explicit geometric constraints. This unconstrained flexibility often leads to structural distortion and artifacts, limiting reconstruction fidelity. We propose G^{3}R, Gaussian-based Geometry-Guided Reconstruction, a geometry-guided Gaussian reconstruction framework that introduces a straightforward yet effective structural prior to stabilize Gaussian optimization. Specifically, we develop the Density-Gated Prior Network (DGPN) to yield a prior volume from physics-derived density map. This prior volume is then incorporated in a two-stage Gaussian reconstruction process: (1) prior-based primitive initialization to establish stable geometry, and (2) slice-driven refinement with an annealed distillation constraint to preserve structural consistency while recovering fine details. We demonstrate that our approach matches the reconstruction speed of existing Gaussian frameworks while delivering state-of-the-art quality with 24.56 dB PSNR in simulated dataset, specifically achieving up to +1.87 dB PSNR in a single-stack stress test. These promising results highlights the significant potential of application into clinical routine of real-time fetal 3D MRI.

Keywords: Slice-to-Volume Reconstruction · 3D Gaussian Splatting · Super-Resolution · Fetal MRI · Single-stack Reconstruction

1 Introduction

Non-invasive fetal and pediatric brain magnetic resonance imaging (MRI) is crucial for prenatal diagnosis and neuro developmental assessment [21, 2]. Moreover, an isotropic high resolution volume MRI is increasingly critical for assessing cortical development and diagnosing congenital neurological disorders [7, 18]. Unlike conventional brain imaging, fetal scans are affected by unpredictable inter-slice motion, thick-slice acquisition with inter-slice gaps, and limited orientation coverage. Hence, fetal brain MRI reconstruction aims to recover high-resolution 3D volumetric structures from motion-corrupted 2D slices.

Classical slice-to-volume reconstruction (SVR) methods rely on iterative slice registration [20], evolving with many variants [10, 6, 14, 4]. Recently, Xu et al. [26] incorporated attention mechanisms to capture long-range slice correlations, while Wu et al. [24] adopted Mamba-based sequence modeling to precisely model spatial dependencies. Although progressively enhanced with robust estimators and patch-based strategies, these approaches remain sensitive to motion complexity and noise. In parallel, Implicit Neural Representations (INRs) [25] have emerged as an efficient reconstruction tool to reduce memory usage and enable arbitrary resolution reconstruction. INRs model the volume as a continuous function and utilize Point Spread Function (PSF) [19, 5, 11] to simulate the slice acquisition process. However, PSF within INRs requires expensive stochastic Monte Carlo sampling [25, 17], resulting in prohibitive inference latency.

3D Gaussian Splatting (3DGS) is an advanced rasterization technique that demonstrates significant utility in generating photorealistic scenes [12] and medical imaging [16]. Recently, GSVR [3] attempts to address inference fetal MRI reconstruction latency from the perspective of 3DGS. It eliminates the need for expensive stochastic Monte Carlo sampling by explicitly modeling the high-resolution volume as a sparse cloud of anisotropic Gaussian primitives, unlock a remarkable reconstruction speed. However, as analyzed by Bao et al. [1], 3DGS inherently suffers from structural instability in under-determined scenarios. This limitation is particularly prominent in clinical fetal MRI, where thick slices and motion render the reconstruction highly ill-posed. Without explicit regularization, optimizing massive weakly coupled Gaussian primitives leads to severe overfitting. Instead of forming a coherent 3D structure, the primitives merely minimize 2D projection errors by flattening or stretching, ultimately causing severe structural distortion. To alleviate this instability, introducing an additional prior serves as a straightforward and effective method to guide the optimization process [1]. While conventional prior-injection methods rely on generative models (e.g., Diffusion Models [13, 22, 15]), their prohibitive computational costs neutralize 3DGS's rapid inference advantage.

To provide robust prior guidance without sacrificing efficiency, we propose G^{3}R: Gaussian-based Geometry-Guided Reconstruction. Unlike costly models above, G^{3}R uses a fast, pre-trained Density-Gated Prior Network (DGPN) to provide robust structural guidance with minimal computational overhead. We introduce a Density-Gated Refinement (DGR) module within DGPN. DGR constructs a physical density map by accumulating the PSF of the input slices. This density map quantifies observation confidence and distinguishes well-scanned regions from acquisition voids. Further, according to this map, a refiner explicitly inpaints missing anatomical structures, producing a structurally continuous and robust geometric prior even in extreme clinical scenarios (e.g., single-stack).

To effectively assimilate this prior, G^{3}R employs a tailored two-stage Gaussian reconstruction strategy. The Gaussian primitives are first initialized via a rapid volume-fitting to the geometric prior, establishing a reliable structural foundation. During the subsequent slice-fitting optimization against the observed 2D slices, this prior acts as a continuous constraint on the global structure. As

a result, G^{3}R effectively addresses the structural instability of unconstrained 3DGS and yields significant improvements in recovering coherent 3D anatomy, demonstrating superior robustness especially when clinical observations are limited or insufficient. Our main contributions are as follows:

1. We propose G^{3}R, a novel framework that explicitly enforces continuous geometric constraints to elevate the reconstruction fidelity of 3DGS-based SVR. By integrating a fast Density-Gated Prior Network (DGPN) into a two-stage Gaussian process, our approach ensures structural coherence and superior anatomical detail recovery across varying acquisition conditions.

2. We introduce a novel Density-Gated Refinement (DGR) mechanism within the DGPN. Guided by physical density maps, DGR effectively inpaints missing anatomical structures to ensure a continuous and reliable volume prior, even from highly sparse inputs.

3. Experiments confirm G^{3}R's superiority in robustness-critical regimes. It effectively prevents structural distortion while maintaining fast inference speeds.

2 Methodology

2.1 Preliminary

Mathematically, reconstructing the volume V from K motion-corrupted slices $\mathcal{Y} = \{y_k\}_{k=1}^K$ is an inverse problem defined by:

$$V^* = \arg\min_V \mathcal{M}(\mathcal{Y}, \mathcal{A}(V)) + \lambda \mathcal{R}(V), \tag{1}$$

where $\mathcal{A}$ models the forward acquisition physics [6], $\mathcal{M}$ is the metric that measures the similarities and $\mathcal{R}(V)$ represents a regularization term. Physically, $\mathcal{A}(V)$ models the observed intensity $\hat{y}_k(\mathbf{u})$ by integrating the continuous volume V over a slice defined by the PSF, while accounting for motion:

$$\hat{y}_k(\mathbf{u}) = \int_{\mathbb{R}^3} V(\mathbf{x}) \cdot \mathrm{PSF}(\mathbf{x} - \mathcal{T}_k(\mathbf{u}))\, d\mathbf{x}, \tag{2}$$

where $\mathcal{T}_k(\mathbf{u}) = \mathbf{R}_k\mathbf{u} + \mathbf{t}_k$ denotes a rigid transformation with motion rotation $\mathbf{R}_k$ and translation $\mathbf{t}_k$. This continuous formulation is typically approximated via Monte Carlo sampling, leading to high computational cost.

To circumvent this limitation, following [3], we parameterize the volume V as a continuous mixture of N 3D Gaussian primitives $\Theta = \{\mu_j, \Sigma_j, c_j\}_{j=1}^N$, where $\mu_j \in \mathbb{R}^3$ denotes the spatial mean position, $c_j \in \mathbb{R}$ represents the intensity coefficient, and covariance $\Sigma_j \in \mathbb{R}^{3\times3}$ is factorized into rotation $\mathbf{R}_j$ and scaling $\mathbf{S}_j$ matrices ($\Sigma_j = \mathbf{R}_j\mathbf{S}_j\mathbf{S}_j^T\mathbf{R}_j^T$):

$$V(\mathbf{x}) = \sum_{j=1}^{N} c_j \cdot \exp\left(-\frac{1}{2}(\mathbf{x} - \mu_j)^T \Sigma_j^{-1}(\mathbf{x} - \mu_j)\right) \tag{3}$$

Crucially, by modeling the PSF as an anisotropic Gaussian kernel Σ_{PSF}, we leverage the property that Gaussians are closed under convolution. This explicit

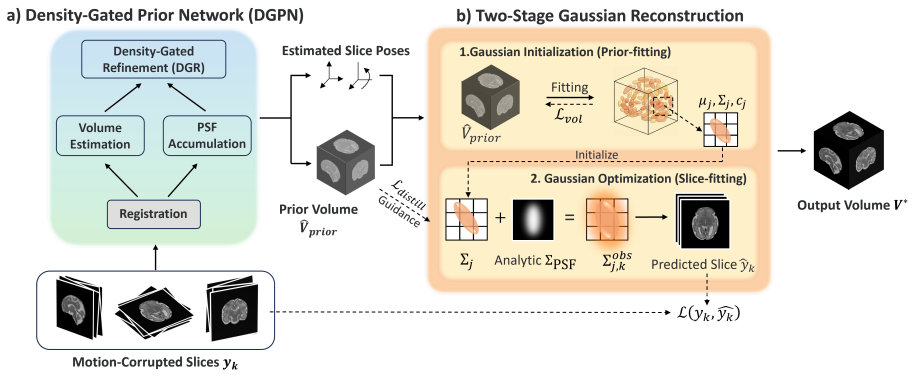

Fig. 1. Overview of the G^{3}R framework. (a) Density-Gated Prior Network (DGPN): Estimates slice poses from motion-corrupted slices and synthesizes a robust geometric prior $\hat{V}_{prior}$ using Density-Gated Refinement (DGR) to fill acquisition voids. (b) Two-Stage Gaussian Reconstruction: Gaussian primitives are initialized by volume-fitting to $\hat{V}_{prior}$. During subsequent slice-fitting process via exact covariance addition ($\Sigma_j + \Sigma_{\text{PSF}}$), primitives are refined against observed slices under the guidance of $\hat{V}_{prior}$.

formulation transforms the intractable integral in Eq. (2) into a closed-form analytical solution, reducing the slice acquisition process to an exact covariance addition. For a primitive j observed in slice k, its effective observed covariance $\Sigma^{\text{obs}}_{j,k}$ becomes:

$$\Sigma^{\text{obs}}_{j,k} = \Sigma_j + \mathbf{R}_k \Sigma_{\text{PSF}} \mathbf{R}_k^T \tag{4}$$

The final intensity $\hat{y}_k(\mathbf{u})$ can be directly and rapidly evaluated at the motion-corrected coordinate $\mathbf{x} = \mathcal{T}_k(\mathbf{u})$:

$$\hat{y}_k(\mathbf{u}) = \sum_{j \in \mathcal{N}(\mathbf{u})} c_j \exp\left(-\frac{1}{2} (\mathcal{T}_k(\mathbf{u}) - \mu_j)^T (\Sigma^{\text{obs}}_{j,k})^{-1} (\mathcal{T}_k(\mathbf{u}) - \mu_j) \right), \tag{5}$$

where $\mathcal{N}(\mathbf{u})$ denotes the primitives contributing to the queried coordinate.

2.2 Overview of G^{3}R

For the first time, we present a framework that leverages a Density-Gated Prior Network (DGPN) to explicitly guide volumetric Gaussian optimization. As illustrated in Fig. 1, G^{3}R operates in two synergistic stages:

1. Density-Gated Prior Network (Sec. 2.3): From the motion-corrupted slices, the pre-trained Density-Gated Prior Network (DGPN) employs Density-Gated Refinement (DGR) to explicitly inpaint acquisition voids. This synthesizes a robust, structurally continuous geometry prior, $\hat{V}_{prior}$.

2. Two-Stage Gaussian Reconstruction (Sec. 2.4): The Gaussian primitives Θ are first initialized by volume-fitting to $\hat{V}_{prior}$. During the subsequent slice-fitting optimization, this prior constrains the global geometry, enabling the

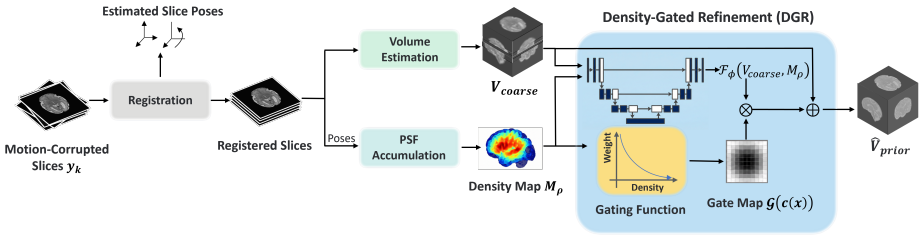

Fig. 2. Architecture of the DGPN. Registered slices yield a coarse volume V_{coarse} and a physical density map M_ρ via PSF accumulation. DGR module derives confidence $c(\mathbf{x})$ from M_ρ to generate a gating map $\mathcal{G}(c(\mathbf{x}))$. This gate modulates the network to preserve dense regions while inpainting voids, yielding the robust prior $\hat{V}_{prior}$.

primitives to accurately recover high-frequency features from the observed slices without structural distortion.

2.3 Density-Gated Prior Network (DGPN)

To robustly initialize Gaussian primitives, we propose DGPN. Unlike standard learning-based SVR methods [8, 26] where volume reconstruction relies on optimization based solvers (e.g., Conjugate Gradient) that struggle with motion artifacts in low-quality data regimes, we introduce Density-Gated Refinement (DGR) (Fig. 2). By leveraging MRI acquisition physics, DGR explicitly inpaints voids to generate a structurally continuous, robust prior $\hat{V}_{prior}$.

In the context of SVR problem [11], Eq. (2) shows that the observational support at any 3D voxel $\mathbf{x}$ is determined by the cumulative PSF contributions from all acquired 2D pixels. From this perspective, the observational support at location $\mathbf{x}$ can be explicitly extracted as a Physical Density Map $M_\rho(\mathbf{x})$:

$$M_\rho(\mathbf{x}) = \sum_{k=1}^{K} \int_{\Omega_k} \mathrm{PSF}\left(\mathbf{x} - \mathcal{T}_k(\mathbf{u})\right) d\mathbf{u}, \tag{6}$$

where $M_\rho(\mathbf{x})$ represents the spatial coverage of the forward operator and naturally approaching zero in acquisition voids [11, 14].

Given a coarse volume V_{coarse} derived from registered slices, we synthesize the anchor $\hat{V}_{prior}$ by modulating a residual refinement network $\mathcal{F}_\phi$ with a density-dependent gating function $\mathcal{G}(\cdot)$. We define the observation confidence as $c(\mathbf{x}) = 1 - \exp(-M_\rho(\mathbf{x}))$. As shown in Fig. 2, the gating function blends the network prediction based on this confidence:

$$\hat{V}_{prior}(\mathbf{x}) = V_{coarse}(\mathbf{x}) + \mathcal{G}(c(\mathbf{x})) \cdot \mathcal{F}_\phi(V_{coarse}, M_\rho), \tag{7}$$

where $\mathcal{G}(c(\mathbf{x})) = \alpha \cdot c(\mathbf{x}) + \beta \cdot (1 - c(\mathbf{x}))$. We set $\alpha = 0.05$ to preserve high-fidelity features in dense regions ($c \to 1$), and $\beta = 1.0$ to fully utilize the neural prior in acquisition voids ($c \to 0$). The entire prior learning network is pre-trained in a

supervised manner on a synthetic dataset [26] generated from high-quality fetal brain atlases. This pre-training minimizes the reconstruction L_1 loss against ground-truth volumes, enabling the network to generalize across subjects and provide a reliable initialization before Gaussian reconstruction.

2.4 Two-Stage Gaussian Reconstruction

As illustrated in Fig. 1(b), our framework proceeds in two stages to ensure structural stability and high-frequency features recovery.

Prior-Fitting. We first initialize the Gaussian primitives Θ from $\hat{V}_{prior}$ by minimizing the volumetric difference within the foreground region Ω_{fg}:

$$\mathcal{L}_{vol} = \frac{1}{|\Omega_{fg}|} \sum_{\mathbf{x} \in \Omega_{fg}} \|V(\mathbf{x}) - \hat{V}_{prior}(\mathbf{x})\|_2^2 + \lambda_{scale} \mathcal{L}_{scale}, \tag{8}$$

where $\mathcal{L}_{scale} = \sum_j \|\mathbf{s}_j - \mathbf{s}_{tgt}\|_2^2$ regularizes the covariance scaling factors $\mathbf{s}_j$ (where $\mathbf{S}_j = \mathrm{diag}(\mathbf{s}_j)$) towards a target mean $\mathbf{s}_{tgt}$ to prevent overfitting to noise.

Slice-Fitting. Subsequently, we refine these primitives against the observed 2D slices to recover high-frequency features. The loss is defined as:

$$\mathcal{L}_{slice} = \sum_k \|\hat{y}_k - y_k\|_1 + \lambda_{scale} \mathcal{L}_{scale} + \omega(i) \cdot \mathcal{L}_{distill}, \tag{9}$$

where $\mathcal{L}_{distill} = \|V - \hat{V}_{prior}\|_1$ anchors the global geometry, and $\mathcal{L}_{scale}$ stays the same as Eq. (8). To smoothly shift the optimization focus from robust global structural constraint to fine detail recovery, the guidance weight $\omega(i)$ follows a linear annealing schedule, decaying from an initial strong penalty ω_{start} to a minimal bound ω_{end} between iterations τ_1 and τ_2.

3 Experiments and Results

For quantitative comparison of the reconstructed images, we measure some metrics such as the peak-signal-to-noise ratio (PSNR), structural similarity index measure (SSIM) [23], Normalized Cross-Correlation (NCC), root mean square error (RMSE), Learned Perceptual Image Patch Similarity (LPIPS) [28].

Datasets and baselines. We performed extensive experiments to evaluate the proposed model in two datasets. **1**) Simulated Data: Motion-corrupted stacks were synthesized from high-quality CRL [7] and Feta 2.2 [18] atlases (GA: 21-38w) with rigid motion and slice-wise intensity inhomogeneities [9]. Each subject contains three orthogonal stacks ($0.8 \times 0.8 \times 2.2\text{mm}^3$). **2**) Clinical Data: Real-world performance was assessed on 90 fetal subjects (GA: 28-37w) acquired on a 3T Philips Ingenia scanner (T2-weighted SSh-TSE, $0.71 \times 0.71 \times 3.0\text{mm}^3$, TR/TE=15000/183ms). We compared G^{3}R against three representative methods: SVoRT [26], NeSVoR [25], and GSVR [3]. All experiments were conducted on a single NVIDIA RTX 4090 GPU (24GB) to ensure a fair comparison.

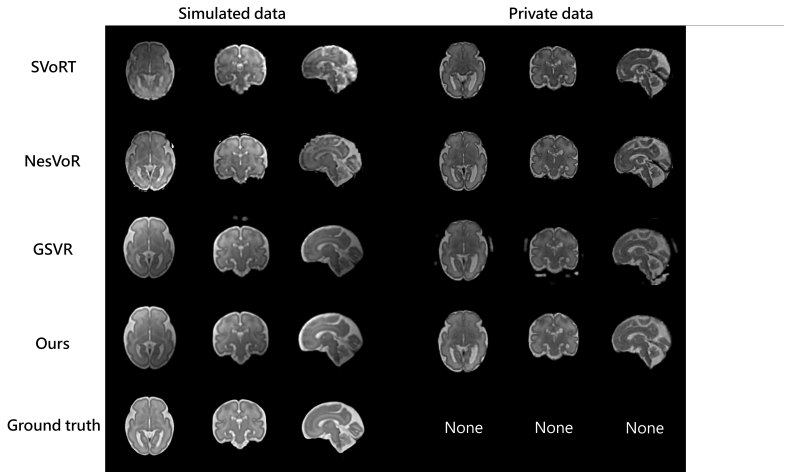

Fig. 3. The reconstruction results of two typical cases in the simulated data (left) and private clinical data (right).

Table 1. Quantitative results on the simulated dataset (mean±std).(**bold**:best)

Method	Time (s)↓	PSNR(dB)↑	SSIM↑	NCC↑	RMSE↓	LPIPS↓
SVoRT	∼2	22.47 ± 2.86	0.814 ± 0.069	0.944 ± 0.011	0.079 ± 0.023	0.112 ± 0.034
NeSVoR	∼175	24.30 ± 2.83	0.902 ± 0.037	**0.960** ± 0.007	0.064 ± 0.019	0.099 ± 0.017
GSVR	∼85	24.40 ± 2.93	0.898 ± 0.036	0.959 ± 0.008	0.063 ± 0.019	**0.085** ± 0.013
G^{3}R	∼90	**24.56** ± 2.87	**0.905** ± 0.033	**0.960** ± 0.008	**0.062** ± 0.018	**0.085** ± 0.012

Comparison on simulated dataset. Table 1 and Fig. 3 highlight the superior structural fidelity of our framework. While SVoRT tends to over-smooth, implicit (NeSVoR) and explicit (GSVR) baselines suffer from high-frequency artifacts and noise. In contrast, G^{3}R ensures coherent anatomy, achieving the best quantitative balance (PSNR **24.56 dB**, LPIPS **0.085**). Crucially, by distilling the robustness prior, G^{3}R effectively regularizes the optimization, eliminating the geometric instability inherent in unconstrained Gaussians. Efficiency-wise, G^{3}R retains the rapid inference of 3DGS (∼90s), significantly outpacing NeSVoR.

Robustness in Extreme Clinical Scenarios. To validate the effectiveness of our geometric guidance, we focus on performance in challenging clinical scenarios. Specifically, we utilize single-stack reconstruction as an extreme stress test to evaluate robustness. While specialized single-stack methods such as FC-SVR [27] rely on a pre-trained interpolation network, our framework ensures structural reliability through the observation-guided prior. As shown in Fig.4, while unconstrained baselines suffer severe geometric collapse in single-stack scenarios, G^{3}R maintains high stability, outperforming GSVR by a substantial margin of

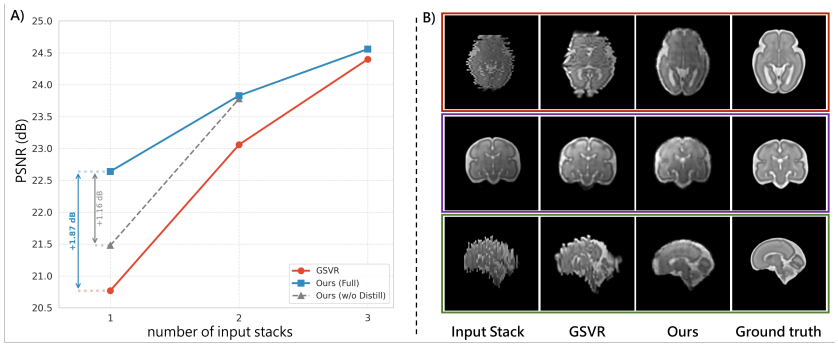

Fig. 4. Robustness analysis extreme clinical scenarios. **(A)** PSNR drop-off as input stacks decrease from 3 to 1. **(B)** Visual comparison in the extreme single-stack stress test. Unconstrained GSVR suffers from severe structural distortion due to the lack of orthogonal views, whereas G^{3}R preserves coherent 3D anatomy.

Table 2. Ablation on the DGPN prior. DGPN w/o M_ρ: DGR without density gating. Metrics are evaluated on the prior volume ($\hat{V}_{prior}$) before 3DGS optimization.

Setting	Configuration	PSNR(dB)↑	SSIM↑	LPIPS↓
Single-Stack	DGPN w/o M_ρ	22.09 ± 1.11	0.242 ± 0.068	0.088 ± 0.019
	Full DGPN ($\hat{V}_{prior}$)	**22.11** ± 2.19	**0.781** ± 0.053	**0.080** ± 0.019
Triple-Stack	DGPN w/o M_ρ	**23.39** ± 2.02	0.367 ± 0.082	0.079 ± 0.024
	Full DGPN ($\hat{V}_{prior}$)	22.75 ± 2.28	**0.826** ± 0.062	**0.070** ± 0.018

1.87 dB (22.64 vs. 20.77 dB). This confirms that our method prevents structural distortion, ensuring reliability even in the most challenging scenarios (Fig.4B).

Ablation Studies. Table 2 demonstrates the critical role of the density map (M_ρ). Removing it causes the DGPN to default to over-smoothed "mean regression," minimizing MSE but catastrophically reducing SSIM (0.367 vs. 0.826). By explicitly distinguishing valid anatomy from voids, our DGR preserves high-frequency cortical details. Moreover, PSNR increases from 22.75 dB to 24.56 dB after 3DGS optimization, confirming its contribution beyond the DGPN prior. Regarding the "Guidance" mechanism (dashed line in Fig. 1), removing the distillation constraint ($\mathcal{L}_{distill}$) leads to a 1.16 dB drop in single-stack scenarios (Fig. 4A), as the Gaussians degenerate into overfitting slice thickness without structural anchoring. Notably, as the number of input stacks increases, this performance gap narrows, indicating that the reconstruction naturally shifts from relying on the prior to extracting fine details directly from the abundant slice observations.

4 Conclusion

In this work, we introduced G^{3}R to enhance the fidelity and structural consistency of 3DGS-based fetal brain MRI reconstruction. By synergizing a geometry-guided prior with explicit Gaussian representation, our approach transforms the optimization into a structurally regularized process. This yields superior reconstruction quality with preserved high-frequency details, demonstrating remarkable robustness and efficiency. G^{3}R highlights the potential of integrating physically grounded priors into fast reconstruction frameworks to achieve reliable clinical imaging, paving the way for future uncertainty-aware safety mechanisms.

Acknowledgments. This study was funded by the Key-Area Research and Development Program of Guangdong Province (2023B0303040001), the General Program of the National Natural Science Foundation of China (NSFC-62471185), the Guangdong Basic and Applied Basic Research Foundation (2024A1515010180 and 2025A1515012836), and the Guangdong Provincial Key Laboratory of Human Digital Twin (2022B1212010004).

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

References

1. Bao, Y., Ding, T., Huo, J., Liu, Y., Li, Y., Li, W., Gao, Y., Luo, J.: 3d gaussian splatting: Survey, technologies, challenges, and opportunities. IEEE Transactions on Circuits and Systems for Video Technology **35**(7), 6832–6852 (2025)
2. Bethlehem, R.A., Seidlitz, J., White, S.R., Vogel, J.W., Anderson, K.M., Adamson, C., Adler, S., Alexopoulos, G.S., Anagnostou, E., Areces-Gonzalez, A., et al.: Brain charts for the human lifespan. Nature **604**(7906), 525–533 (2022)
3. Dannecker, M., Jia, S., Stolt-Ansó, N., Girard, N., Auzias, G., Rousseau, F., Rueckert, D.: Fast and explicit: Slice-to-volume reconstruction via 3d gaussian primitives with analytic point spread function modeling. arXiv preprint arXiv:2512.11624 (2025)
4. Ebner, M., Wang, G., Li, W., Aertsen, M., Patel, P.A., Aughwane, R., Melbourne, A., Doel, T., Dymarkowski, S., De Coppi, P., et al.: An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain mri. NeuroImage **206**, 116324 (2020)
5. Gholipour, A., Estroff, J.A., Sahin, M., Prabhu, S.P., Warfield, S.K.: Maximum a posteriori estimation of isotropic high-resolution volumetric mri from orthogonal thick-slice scans. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 109–116. Springer (2010)
6. Gholipour, A., Estroff, J.A., Warfield, S.K.: Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain mri. IEEE transactions on medical imaging **29**(10), 1739–1758 (2010)
7. Gholipour, A., Rollins, C.K., Velasco-Annis, C., Ouaalam, A., Akhondi-Asl, A., Afacan, O., Ortinau, C.M., Clancy, S., Limperopoulos, C., Yang, E., et al.: A normative spatiotemporal mri atlas of the fetal brain for automatic segmentation and analysis of early brain growth. Scientific reports **7**(1), 476 (2017)

8. Hou, B., Alansary, A., McDonagh, S., Davidson, A., Rutherford, M., Hajnal, J.V., Rueckert, D., Glocker, B., Kainz, B.: Predicting slice-to-volume transformation in presence of arbitrary subject motion. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 296–304. Springer (2017)
9. Huang, S., Chen, G., Sun, K., Cui, Z., Zhang, X., Xue, P., Zhang, X., Zhang, H., Shen, D.: Super-resolution reconstruction of fetal brain mri with prior anatomical knowledge. In: International Conference on Information Processing in Medical Imaging. pp. 428–441. Springer (2023)
10. Jiang, S., Xue, H., Glover, A., Rutherford, M., Rueckert, D., Hajnal, J.V.: Mri of moving subjects using multislice snapshot images with volume reconstruction (svr): application to fetal, neonatal, and adult brain studies. IEEE transactions on medical imaging **26**(7), 967–980 (2007)
11. Kainz, B., Steinberger, M., Wein, W., Kuklisova-Murgasova, M., Malamateniou, C., Keraudren, K., Torsney-Weir, T., Rutherford, M., Aljabar, P., Hajnal, J.V., et al.: Fast volume reconstruction from motion corrupted stacks of 2d slices. IEEE transactions on medical imaging **34**(9), 1901–1913 (2015)
12. 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)
13. Khader, F., Müller-Franzes, G., Tayebi Arasteh, S., Han, T., Haarburger, C., Schulze-Hagen, M., Schad, P., Engelhardt, S., Baeßler, B., Foersch, S., et al.: Denoising diffusion probabilistic models for 3d medical image generation. Scientific reports **13**(1), 7303 (2023)
14. Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M.A., Hajnal, J.V., Schnabel, J.A.: Reconstruction of fetal brain mri with intensity matching and complete outlier removal. Medical image analysis **16**(8), 1550–1564 (2012)
15. Lv, Y., Cai, Z., Chen, S., Yang, C., Zhang, X.: D-svg: Diffusion based slice-to-volume generation with implicit neural representation for fetal brain mri. In: International Workshop on Preterm, Perinatal and Paediatric Image Analysis. pp. 15–25. Springer (2025)
16. Marzol, K., Kolton, I., Smolak-DyĹL'ewska, W., Kaleta, J., Mazur, M., Spurek, P., et al.: Medgs: Gaussian splatting for multi-modal 3d medical imaging. arXiv preprint arXiv:2509.16806 (2025)
17. 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)
18. Payette, K., de Dumast, P., Kebiri, H., Ezhov, I., Paetzold, J.C., Shit, S., Iqbal, A., Khan, R., Kottke, R., Grehten, P., et al.: An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset. Scientific data **8**(1), 167 (2021)
19. Rousseau, F., Glenn, O., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D., Barkovich, J., Studholme, C.: A novel approach to high resolution fetal brain mr imaging. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 548–555. Springer (2005)
20. Rousseau, F., Glenn, O.A., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D.B., Barkovich, J.A., Studholme, C.: Registration-based approach for reconstruction of high-resolution in utero fetal mr brain images. Academic radiology **13**(9), 1072–1081 (2006)
21. Uus, A.U., Egloff Collado, A., Roberts, T.A., Hajnal, J.V., Rutherford, M.A., Deprez, M.: Retrospective motion correction in foetal mri for clinical applications: existing methods, applications and integration into clinical practice. The British journal of radiology **96**(1147), 20220071 (2023)

22. Wang, H., Zhang, X., Li, T., Wan, Y., Chen, T., Sun, J.: Dmplug: A plug-in method for solving inverse problems with diffusion models. Advances in Neural Information Processing Systems **37**, 117881–117916 (2024)
23. 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)
24. Wu, J., Wei, H., Zhang, Y.: Svrmamba: Slice-to-volume reconstruction from multiple mri stacks with slice sequence guided mamba. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 8404–8412 (2025)
25. Xu, J., Moyer, D., Gagoski, B., Iglesias, J.E., Grant, P.E., Golland, P., Adalsteinsson, E.: Nesvor: implicit neural representation for slice-to-volume reconstruction in mri. IEEE transactions on medical imaging **42**(6), 1707–1719 (2023)
26. Xu, J., Moyer, D., Grant, P.E., Golland, P., Iglesias, J.E., Adalsteinsson, E.: Svort: Iterative transformer for slice-to-volume registration in fetal brain mri. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 3–13. Springer (2022)
27. Young, S.I., Balbastre, Y., Fischl, B., Golland, P., Iglesias, J.E.: Fully convolutional slice-to-volume reconstruction for single-stack mri. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11535–11545 (2024)
28. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)