

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ULTRA \mathbf{Gauss} : ULTRAFAST GAUSSIAN RECONSTRUCTION OF 3D ULTRASOUND VOLUMES

Anonymous authors

Paper under double-blind review

ABSTRACT

Ultrasound imaging is widely used due to its safety, affordability, and real-time capabilities, but its 2D interpretation is highly operator-dependent, leading to variability and increased cognitive demand. We present **UltraGauss**: an ultrasound-specific Gaussian Splatting framework that serves as an efficient approximation to acoustic image formation. Unlike projection-based splatting, UltraGauss renders by *probe-plane intersection* with in-plane aggregation, aligning with plane-based echo sampling while remaining fast and memory-efficient. A stable parameterisation and compute-aware GPU rasterisation make this method practical at scale. On clinical datasets, UltraGauss delivers state-of-the-art 2D-to-3D reconstructions in minutes on a single GPU (reaching 0.99 SSIM within \sim 20 minutes), and a clinical expert survey rates its reconstructions the most realistic among competing methods. To our knowledge, this is the first Gaussian Splatting approach tailored to ultrasound 2D-to-3D reconstruction. Code will be released upon publication.

1 INTRODUCTION

Ultrasound (US) is a mainstay in medical imaging: it is real-time, low-cost, portable, and non-ionizing. Yet routine use still asks clinicians to infer 3D anatomy from 2D slices: a cognitively demanding step that introduces operator-dependent variability (Benacerraf, 2002; Nelson & Pretorius, 1998), undermining reproducibility and hindering standardized assessment. While volumetric (3D) probes exist, their workflows are largely offline and the hardware is costly, making them uncommon outside well-resourced centres (Merz & Welter, 2005). Enabling 3D reasoning from routine 2D acquisitions offers a software-only path to scale volumetric assessment across sites and resource levels. As 2D is the universal denominator in global healthcare, 2D-to-3D reconstruction will yield standardised volume and surface metrics that are difficult to obtain consistently from single slices.

Recent learning-based approaches span *implicit* NeRF-style representations (Yeung et al., 2024; Wysocki et al., 2023; Eid et al., 2025; Gaits et al., 2024) and *explicit* voxel grids (Solberg et al., 2007). However, implicit fields are computationally heavy; voxel grids incur memory and resolution limits; and many methods adopt light-transport assumptions (ray accumulation with transmittance), which fundamentally mismatch US physics, where waves propagate into tissue and return to the probe (Powles et al., 2018; Aldrich, 2007). Classical Gaussian Splatting (GS) for cameras achieves fast, high-quality rendering in optical settings by projecting 3D Gaussians onto the image plane and blending splats in depth order using alpha compositing (Kerbl et al., 2023). US images, on the other hand, are not perspective renderings: they sample echo intensities within the probe plane (with attenuation), so the camera-style projection and occlusion paradigm does not apply.

We introduce **UltraGauss**: an ultrasound-specific GS reconstruction framework that serves as an efficient approximation to the US image formation model, replacing projection-based rendering with probe-plane intersection rendering, consistent with wave-based acquisition. Instead of marching rays or *projecting* Gaussians into 2D, UltraGauss evaluates anisotropic 3D Gaussians where they *intersect* the probe plane and aggregates intensities in-plane. This removes the need for depth-based occlusion, matches the acquisition geometry of linear *and* curvilinear probes, and enables resolution-free slicing at arbitrary orientations. This approximation preserves plane-based sampling and dominant attenuation behaviour while avoiding expensive wave simulation.

Our technical contributions are:

- **Efficient forward-model approximation:** UltraGauss renders by *probe-plane* intersection rather than camera projection, capturing plane-based sampling and attenuation while side-stepping costly wave simulation.

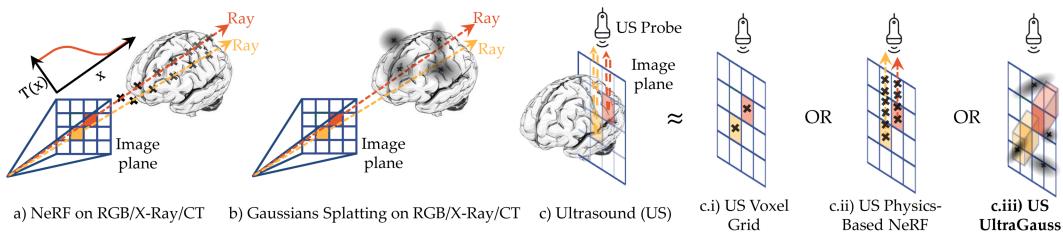


Figure 1: Image formation mechanisms. (a, b) Projection-based (e.g. camera) rendering with depth-ordered compositing and (c) ultrasound. (c.i) Slice sampling from a voxel grid or continuous field. (c.ii) Physics-based ultrasound: path integration with attenuation/backscatter. (c.iii) **UltraGauss**: plane-intersection of 3D Gaussians with in-plane aggregation.

- **Stable parameterisation:** A triangular inverse-covariance (*precision*) factorisation guarantees positive-definiteness and well-conditioned gradients for reconstruction and densification.
- **Compute-aware rasterisation:** *Closed-form χ^2 ellipsoidal bounds* in probe coordinates yield tight per-plane 2D boxes, coupled with a *two-phase, load-balanced CUDA* pipeline that rejects non-intersecting Gaussians and confines work to those boxes.
- **Lightweight attenuation:** A Beer–Lambert term models acoustic shadowing when probe geometry is available, adding realism at negligible cost.
- **Clinical validation:** Experiments on clinical fetal datasets and freehand cine sweeps show high-fidelity 3D reconstructions (strong SSIM) within minutes on a single GPU, with clinician-perceived realism comparable to native 3D scans.

2 RELATED WORK

2D-to-3D reconstruction in medical imaging. Classical pipelines segment 2D slices, register them (e.g. using B-splines), and interpolate a 3D volume (Sarmah et al., 2023). These approaches require substantial manual intervention, struggle with irregular sampling, and scale poorly to high-resolution data. Structure-from-Motion (SfM) methods (e.g. COLMAP (Schönberger & Frahm, 2016)) triangulate depth from photometrically consistent views, but ultrasound violates these assumptions and precise multi-view registration is often impractical.

Learned volumetric representations. In US imaging, three requirements shape learned volumetric methods: coping with irregular plane sampling, avoiding external probe tracking, and adopting a rendering model aligned with plane-based echo acquisition (Fig. 1c) rather than camera projection (Fig. 1a-b). Yeung et al. (2024) (ImplicitVol) reconstructs volumes from routine sweeps without IMU/vision tracking, and Eid et al. (2025) (RapidVol) introduces a tensor-factorised hybrid that accelerates training. These advances improve data efficiency and robustness, yet the underlying rendering remains tied to point/voxel sampling using neural networks.

Ultrasound-specific forward models. Physics-based renderers estimate position-dependent acoustic parameters (e.g. attenuation, scattering) and synthesise images along wave paths (Wysocki et al., 2023; Guo et al., 2024b) (Fig. 1c.ii). These models can be accurate with known probe poses and source geometry, but they necessitate precise acquisition; curvilinear probes (common in obstetrics) further require cone-to-Cartesian transforms, increasing complexity and sensitivity to error.

GS in medical imaging. GS replaces implicit fields with explicit 3D Gaussians and differentiable rasterisation, enabling fast, high-quality rendering. Medical adaptations have largely targeted projection-based modalities: endoscopy (RGB), MRI, and CT (Guo et al., 2024a; Liu et al., 2024; Xie et al., 2024; Bonilla et al., 2024; Peng et al., 2025; Zha et al., 2024; Cai et al., 2025; Nikolakakis et al., 2024). These, therefore, retain camera-style projection and depth compositing (Fig. 1b). Zha et al. (2024) introduce a CUDA-based Gaussian voxel former, but only as a regulariser at small spatial resolutions, and plane sampling from voxels reintroduces interpolation artefacts.

Ultrasound-adapted GS. As far as we know, no ultrasound-adapted GS method exists for 2D-to-3D reconstruction. UltraGauss extends GS to volumetric US by replacing projection-based rendering with probe-plane intersection rendering (Fig. 1c.iii): an efficient approximation to the US image formation model that aligns with plane-based echo sampling and attenuation. The continuous (non-voxelised) Gaussian representation supports arbitrary resolution and orientation slicing without external probe tracking, yielding fast, memory-efficient, and clinically relevant reconstructions.

108 **3 BACKGROUND**
 109

110 For cameras in the visible spectrum, one can render the color c_{RGB} for a given pixel by evaluating a
 111 volumetric model along the corresponding ray, sampling the model’s opacities $\hat{\alpha}(x)$ and colors $\hat{c}(x)$
 112 at m ordered sample points x_j on the ray:

113
$$c_{\text{RGB}} = \sum_{j=1}^m T_j \hat{\alpha}(x_j) \hat{c}(x_j), \quad T_j = \prod_{k=1}^{j-1} (1 - \hat{\alpha}(x_k)), \quad (1)$$

 114
 115

116 where T_j denotes the accumulated transmittance of the material. The main purpose of the accumulated
 117 transmittance is to model occlusions (see Fig. 1a – points at the back are blocked by those in
 118 front, so are given a lower T_j /importance). In the case of GS (Kerbl et al., 2023), the volumetric
 119 model is a weighted sum of n Gaussian functions with means μ_i , covariances Σ_i , colors c_i , and
 120 coefficients (maximum opacities) α_i . In contrast to NeRFs (Mildenhall et al., 2020), the sampled
 121 points x_j can now be reduced to the projections of close Gaussians on to the image plane, which
 122 can be made more efficient. The opacity of the i^{th} Gaussian at a 2D (image-space) point x is then
 123 calculated as

124
$$\hat{\alpha}_i(x) = \alpha_i \exp \left(-\frac{1}{2} \underbrace{(x - \mu_i^{\text{2D}})^T (\Sigma_i^{\text{2D}})^{-1} (x - \mu_i^{\text{2D}})}_{\text{2D squared Mahalanobis distance}} \right), \quad (2)$$

 125
 126

127 with the mean μ_i^{2D} and covariance Σ_i^{2D} projected to the 2D image. To achieve this, the 3D parameters
 128 (μ_i, Σ_i) are translated and rotated to the camera’s reference frame by a view transformation
 129 (extrinsic camera matrix) W , and projected to 2D with the affine approximation of a projective
 130 transformation (using intrinsic camera matrix K) (Zwicker et al., 2001). We can express this formally
 131 with operators to convert to and from homogenous coordinates, $\hbar(u) = [u_1, \dots, u_D, 1]^T$ and
 132 $\hbar^{-1}(u) = [u_1/u_D, \dots, u_{D-1}/u_D]$ respectively:

133
$$\mu_i^{\text{2D}} = \text{proj}(\mu_i) = \hbar^{-1}(K W \hbar(\mu_i)) \quad (3)$$

 134

135
$$\Sigma_i^{\text{2D}} = J_i W \Sigma_i W^T J_i^T, \quad J_i = \frac{\partial \text{proj}(\mu_i)}{\partial \mu_i} \quad (4)$$

 136

137 We can then combine the opacities and colors of all close Gaussians to obtain the rendered opacity
 138 $\hat{\alpha}(x)$ and color $\hat{c}(x)$ for use in Eq. 1:

139
$$\hat{\alpha}(x) = \sum_i^n \hat{\alpha}_i(x), \quad \hat{c}(x) = \frac{1}{\hat{\alpha}(x)} \left(\sum_i^n \hat{\alpha}_i(x) c_i \right) \quad (5)$$

 140
 141

142 In this summary we leave out spherical harmonics, which support directionally-dependent colors
 143 (Ramamoorthi, 2006; Yu et al., 2021).

144 **3.1 OPTIMIZATION**

145 While Eqs. 1-5 can render a Gaussian model, practical use requires optimizations to avoid costly
 146 nested iterations. Namely, Gaussians are tiled and depth-sorted, then distributed across GPU threads
 147 for rasterization onto pixels (Kerbl et al., 2023). Depth sorting and cut-off distances for Gaussians in
 148 the image plane make rendering only approximate (Huang et al., 2024). Heuristics are necessary to
 149 remove/resample Gaussians in overly sparse/dense regions (Rota Bulò et al., 2025; Yu et al., 2023).

150 **4 METHOD**

151 **4.1 IMAGE FORMATION MODEL – ULTRASOUND VS. VISUAL SPECTRUM**

152 While fast splatting is well-developed for RGB cameras, for ultrasound it requires rethinking many
 153 design choices due to differences in image formation. US probes use reflected ultrasound waves
 154 to measure the response of materials at a dense range of depths from the probe (Fig. 1c). Image
 155 cameras, in contrast, measure light that typically reflects off solid objects, and thus are more affected
 156 by occlusions and most often measure surfaces (Fig. 1a-b). Accumulated transmittance T_j in Eq. 1
 157 is therefore unsuitable, as it would treat opaque volumes as occlusions (see $T(x)$ subplot in Fig. 1a).
 158 Instead, the key mechanism for ultrasound is detecting *intersections* with the probe plane (Fig. 1c),
 159 not projections, while accumulated transmittance plays only a secondary role (discussed in Sec. 4.5).

We do this by first opting to include an additional uniform component (background) with color c_{BG} and coefficient α_{BG} , which improves numerical stability by avoiding a division by zero. By reusing the definition of $\hat{c}(x)$ from Eq. 5 to combine the gaussian and background components at a 2D point (pixel) x , our rendering equation simply becomes:

$$\hat{\alpha}(x) = \sum_i^n \hat{\alpha}_i(x) + \alpha_{\text{BG}} \quad (6)$$

$$c_{\text{Ultrasound}}(x) = \hat{c}(x) = \frac{1}{\hat{\alpha}(x)} \left(\sum_i^n \hat{\alpha}_i(x) c_i + \alpha_{\text{BG}} c_{\text{BG}} \right) \quad (7)$$

A large conceptual difference from rendering the visual spectrum is that, instead of projecting the Gaussian parameters to 2D, we need them to *intersect*, or touch, the probe plane in 3D. The opacity $\hat{\alpha}_i(x)$ of the i^{th} Gaussian at a 2D point x (necessary for Eqs. 6-7) is then evaluated as a Mahalanobis distance in 3D space, by *lifting the 2D image point to 3D* in the coordinate-frame of the probe:

$$\hat{\alpha}_i(x) = \alpha_i \exp \left(-\frac{1}{2} \underbrace{\left(x_{|0} - \mu_i^{\text{3D}} \right)^T \left(\Sigma_i^{\text{3D}} \right)^{-1} \left(x_{|0} - \mu_i^{\text{3D}} \right)}_{\text{3D squared Mahalanobis distance}} \right), \quad (8)$$

where $x_{|0} = [x_1, x_2, 0]^T$ and the Gaussian's parameters are moved to the probe's coordinate-frame using its inverse transform matrix W :

$$\mu_i^{\text{3D}} = \hbar^{-1}(W\hbar(\mu_i)), \quad \Sigma_i^{\text{3D}} = W\Sigma_i W^T. \quad (9)$$

Contrast Eq. 2 to Eq. 8: despite the similarities, the former (for RGB images) evaluates a 2D Gaussian after projecting it to image-space, while the later evaluates a 3D Gaussian by doing the reverse operation (for ultrasound images). Using this model, we can now design a fast splatting strategy to avoid the computational expense of simply evaluating Eqs. 6-7 and Eq. 8 (which would result in nested iterations over all pixels and all Gaussians).

4.2 TRIANGULAR COVARIANCE PARAMETERIZATION FOR EFFICIENT INVERSION AND GAUSSIAN SAMPLING

One challenge in optimizing representations with covariance matrices Σ (omitting the subscript i for conciseness) is that they must remain positive-definite (PD, all eigenvalues strictly positive), while gradient-based optimization methods typically only support unconstrained optimization. Kerbl et al. (2023) achieved this by reparameterizing the covariances as a product of a scaling vector s and a quaternion-derived rotation matrix R , as $\Sigma = R \text{diag}(s^2) R^T$. When a Gaussian is projected onto a 2D plane, its covariance becomes 2D, which is easily invertible for use in Eq. 2. However, we found that for our setting the inversion of Σ^{3D} (as opposed to Σ^{2D}) when it is formed by a quaternion (which to represent a 3D rotation must first be normalized), resulted in numerical instabilities. Being a 3×3 matrix, inversion also takes longer. Hence, we propose to learn $(\Sigma^{\text{3D}})^{-1}$ directly. We ensure it is PD by parameterizing it as a product of a matrix M with itself (which is positive-semidefinite), and adding a small multiple of the identity I to ensure positive eigenvalues (with $\beta > 0$):

$$\Sigma'^{-1} = MM^T + \beta I, \quad M = \begin{bmatrix} M_{11} & M_{12} & M_{13} \\ M_{12} & M_{22} & M_{23} \\ M_{13} & M_{23} & M_{33} \end{bmatrix} \quad (10)$$

Note that M itself is symmetric, to achieve the minimal number of degrees of freedom of a 3D covariance (6). No normalization is needed. A remaining challenge is that, in addition to requiring the inverse covariance Σ^{-1} frequently to render pixels (Eq. 8), we also occasionally (e.g. every 100 iterations) must perform heuristic resampling of some Gaussians. This operation requires *inverting* Σ^{-1} *explicitly* to obtain the original covariance Σ , as well as factorizing it to draw a sample from the multivariate gaussian (Sec. 3.1). Both operations can be numerically unstable for ill-conditioned Σ^{-1} . Therefore, we propose instead a more efficient parameterization, as a product of a lower-triangular matrix L :

$$\Sigma^{-1} = LL^T, \quad L = \begin{bmatrix} L_{11}^2 + \beta & 0 & 0 \\ L_{12} & L_{22}^2 + \beta & 0 \\ L_{13} & L_{23} & L_{33}^2 + \beta \end{bmatrix}. \quad (11)$$

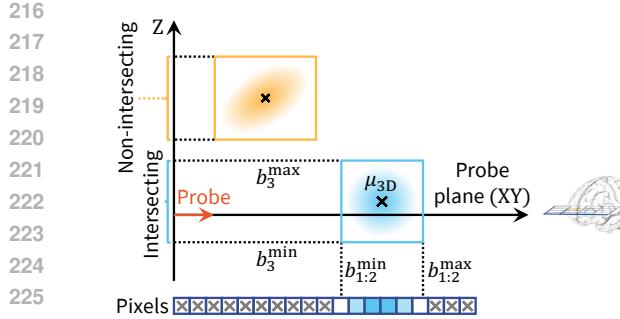


Figure 2: Side view of the Gaussian bounding-box/plane intersection used to derive 2D rasterisation bounds (Sec. 4.3). Boxes that do *not* intersect the probe plane are rejected early (Sec. 4.4); intersecting boxes are rasterized only over their 2D intersection window (Sec. 4.3). Pixels not iterated over are marked with \times . The probe plane (XY) points into the page.

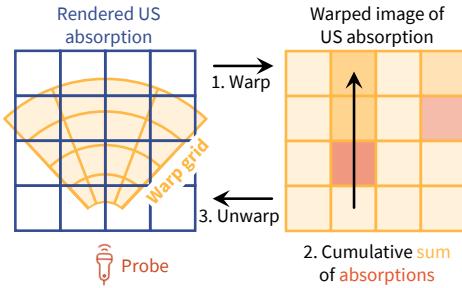


Figure 3: Approximating acoustic shadowing via *warp*–*integrate*–*unwarp*. Left: rendered ultrasound absorption in image space. Right: the same field warped to a beam-aligned grid. (1) Warp to beam space; (2) take a cumulative sum along the beam direction (path integral of absorption); (3) un-warp back to image space to obtain the attenuation map.

Eq. 11 guarantees that Σ^{-1} is PD, since a triangular matrix’s eigenvalues are its diagonal elements, and it can be seen that these are strictly positive (with $\beta > 0$). Moreover, a lower-triangular matrix is extremely efficient to invert via *forward substitution*, which is implemented in most numerical packages. This allows easily computing $\Sigma = (L^{-1})^T L^{-1}$. Finally, another advantage is that this factorization allows sampling from a Gaussian (for the resampling heuristic, Sec. 3.1), by simply projecting a standard normal sample z with the (efficiently) inverted L :

$$y = \mu + L^{-T} z, \quad z \sim \mathcal{N}(0, 1). \quad (12)$$

Eq. 11 supports efficient calculation of Σ^{-1} , and thus opacity $\hat{\alpha}_i(x)$ (Eq. 8), $1.40 \times$ faster than had we used $\Sigma = R \text{diag}(s^2) R^T$. Similarly Σ and the resampling heuristics are computed $1.25 \times$ faster. Yet for most Pixel-Gaussian pairs, opacity is near 0 and so can be ignored. Sec. 4.3 discusses this.

4.3 RASTERIZATION BOUNDARIES

To avoid evaluating Eq. 8 when it is known to yield opacities close to 0, we compute it only inside a bounding box for each Gaussian. The bounding box is defined around the ellipsoid that encompasses $p\%$ of the Gaussian’s probability density. This ellipsoid can be expressed as a function of the 3D squared Mahalanobis distance from Eq. 8, and the chi-squared distribution value for 3 degrees of freedom (omitting the subscript i for clarity):

$$(x_{|0} - \mu^{3D})^T (\Sigma^{3D})^{-1} (x_{|0} - \mu^{3D}) \leq \chi_{3,1-p}^2, \quad (13)$$

which evaluates to $\chi_{3,1-p}^2 = 7.815$ for $p = 95\%$ (note that it is a bound on a squared distance). To calculate this ellipsoid’s bounding box, its bounds are the 3D vectors b^{\min} and b^{\max} :

$$b^{\min/\max} = \mu^{3D} \pm \sqrt{\lambda v}, \quad \lambda = \frac{\chi_{3,1-p}^2}{\det(\Sigma^{3D})}, \quad (14)$$

$$v = \begin{bmatrix} \Sigma_{22}^{3D} \Sigma_{33}^{3D} - (\Sigma_{23}^{3D})^2 \\ \Sigma_{11}^{3D} \Sigma_{33}^{3D} - (\Sigma_{13}^{3D})^2 \\ \Sigma_{11}^{3D} \Sigma_{22}^{3D} - (\Sigma_{12}^{3D})^2 \end{bmatrix},$$

using Σ_{jk}^{3D} to index the $(j, k)^{\text{th}}$ element of a matrix. Eq. 14 can be obtained by writing Eq. 13 as a function of each element of $\delta = x_{|0} - \mu^{3D}$. Fig. 2 illustrates these bounds.

4.4 LOAD BALANCING ACROSS RENDERING THREADS

Note that the bounding box (Eq. 14) can be partitioned into 2 components:

1. The 2D bounding box in the probe plane (first two elements of $b^{\min/\max}$, i.e. $b_{1:2}^{\min/\max}$).

270 2. The 1D segment orthogonal to the probe plane (the third element, b_3^{\min} / \max).

271 This suggests a two-phase process for efficient rendering:

272 1. Reject any Gaussians whose cut-off boundaries do not intersect with the probe plane: $b_3^{\min} > 0$
 273 or $b_3^{\max} < 0$.

274 2. For each accepted Gaussian, only iterate over 2D pixels x inside the bounding box of the plane:
 275 $b_1^{\min} \leq x_1 \leq b_1^{\max}$ and $b_2^{\min} \leq x_2 \leq b_2^{\max}$.

276 This process is shown in Fig. 2, and naturally balances the load across parallel GPU threads. Phase
 277 1 requires iterating through all the Gaussians (but not the pixels), marking them as accepted or re-
 278 jected based on the perpendicular distance, which can be equally partitioned between all the threads.
 279 Standard buffer compaction (Corp., 2020) can then reduce this list of Gaussians to only the accepted
 280 ones. Phase 2 then requires iterating through the compacted list of only accepted Gaussians, and
 281 rasterizing each one only onto the corresponding bounding box of the image buffer, by atomically
 282 adding to pixel accumulators for color $\hat{c}(x)$ and opacity $\hat{\alpha}(x)$ (implementing Eqs. 6-7). The com-
 283 pacted list in Phase 2 is again equally partitioned among the threads, to ensure optimal throughput.

284 4.5 SHADOW MODELLING

285 While other NeRF-based methods model more complex ultrasound physics (Guo et al., 2024b;
 286 Wysocki et al., 2023), we found that a simple and fast approximation can account for the attenu-
 287 ation of ultrasound waves as they traverse from the probe to a point x , producing shadows along
 288 the way. Based on the Beer-Lambert Law, we calculate an intensity reduction factor T for a pixel
 289 located at position x :

$$291 T(x) = \exp^{\int_{s=0}^{s=x} \alpha(s) ds} \approx \exp^{-\sum_{j=0}^{j=N} \hat{\alpha}_j \delta_j} \quad (15) \\ 292$$

293 where $\hat{\alpha}_j$ is the opacity of pixel j as computed in Eq. 6, and δ_j is the distance between successive
 294 $\hat{\alpha}_j$'s, which in our case is simply the pixel spacing ($= 1/\text{ImageHeight}$). This can be efficiently com-
 295 puted for all pixels by using a cumulative sum (cumsum) along the row dimension, after warping
 296 (bilinearly interpolating) the ultrasound cone into a square image (illustrated in Fig. 3). We then
 297 multiply $c_{\text{Ultrasound}}(x)$ from Eq. 7 by $T(x)$ to give the updated image pixel color.

298 4.6 OPTIMIZATION

299 We train UltraGauss end-to-end by backpropagation with Adam, optimising per-Gaussian parame-
 300 ters $\{\mu_i, L_i, c_i, \alpha_i\}$. From our triangular parameterisation, precision and covariance are linked by

$$302 \Lambda_i \equiv \Sigma_i^{-1} = L_i L_i^\top + \varepsilon I \quad (\text{cf. Eq. 11}), \\ 303$$

304 so we optimise L_i (lower triangular) directly and recover $\Sigma_i = \Lambda_i^{-1}$ as needed.

305 **Initialisation.** Unlike camera GS, we do not use COLMAP initialisation (which yields surface-only
 306 points). Instead, we sample μ_i uniformly within the acquisition volume, and initialise $c_i = 0.5$
 307 and $\alpha_i = 0.731$ (constrained to $[0, 1]$ via a sigmoid). For the precision factor, we draw entries
 308 $L_{ij} \sim \mathcal{U}[4, 5]$; under our normalisation this corresponds to small initial marginal variances for Σ_i
 309 (approximately $(1.7\text{--}4.3) \times 10^{-3}$). A small jitter ε ensures $\Lambda_i \succ 0$.

310 **Training details.** Sparsification and densification heuristics follow (Kerbl et al., 2023; Ye
 311 et al., 2024) with small adaptations for plane-intersection rasterisation. We evaluate $N \in$
 312 $\{100k, 200k, 2M\}$ Gaussians; our CUDA kernels sustain $N \approx 2M$ with minute-level reconstructions
 313 (fewer Gaussians trade a little accuracy for speed). Learning rates are 0.05 for all parameters ex-
 314 cept μ_i , which uses an exponentially decaying schedule starting at 1.6×10^{-4} (Yu et al., 2021; Kerbl
 315 et al., 2023). All experiments run on a single NVIDIA RTX-A4000. Code (PyTorch and custom
 316 CUDA kernels for forward process and gradients) will be released upon publication.

317 5 EXPERIMENTS

318 5.1 DATASETS

319 Two clinical datasets were curated to validate UltraGauss in different ultrasound acquisition settings.
 320 **Dataset A** consists of volumetric scans, allowing assessment of reconstruction fidelity across multi-
 321 ple orthogonal views. **Dataset B** comprises freehand US video sequences, evaluating reconstruction
 322 quality in the absence of full 3D coverage, mimicking real-world fetal monitoring scenarios.

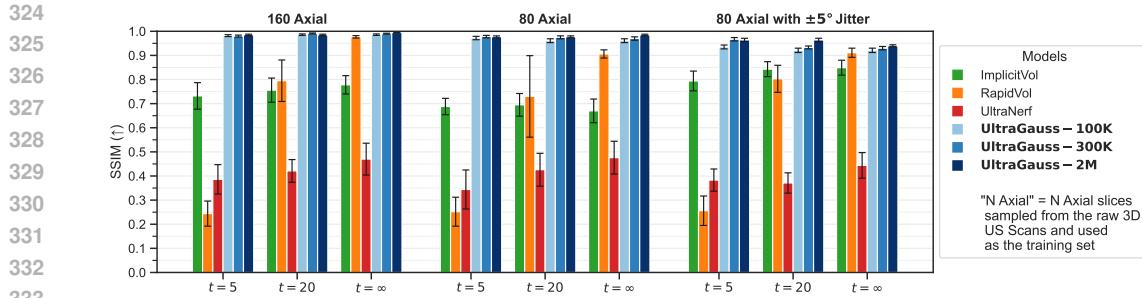


Figure 4: Reconstruction results for 6 different models (ours in **bold**) on 3 different training datasets. Shown at 3 time points over the duration of reconstruction. Higher SSIM is better. Errors bars show ± 1 standard deviation amongst fetuses.

Dataset A – 3D Ultrasound Volumes: This dataset includes twelve 3D fetal brain ultrasound volumes ($160 \times 160 \times 160$ voxels, $0.6 \times 0.6 \times 0.6$ mm 3 resolution), obtained from the INTERGROWTH-21st study (Papageorgiou et al., 2018). Acquisitions were performed between 14 and 26 gestational weeks: spanning a critical period of brain maturation (Namburete et al., 2015; 2023), and the standard time for fetal anomaly screening (Salomon et al., 2022). The scans were collected using a Philips HD9 curvilinear probe (2.5 MHz wave frequency) by multiple sonographers, which introduces variability in probe positioning and image appearance.

Dataset B – 2D Freehand Video Sequences: Three freehand 2D ultrasound videos of fetal brain acquisitions were collected at 19 and 20 weeks’ gestational age at Leiden University Medical Center using a GE Voluson E10 ultrasound scanner. Each video consists of ~ 100 frames, with each frame cropped and resized to 160×160 pixels, and resampled to a resolution of 0.6×0.6 mm 2 . Fig. D1 (Appendix) illustrates how UltraGauss uses Datasets A and B.

5.2 EVALUATION OF RECONSTRUCTION QUALITY AND SPEED

To assess maximum achievable reconstruction quality, we provide UltraGauss with *oracle coverage*: 160 evenly spaced axial slices (sampled from a 3D Scan in Dataset A). This exposes the full $160 \times 160 \times 160$ volume. We evaluate reconstructions by rendering 160 slices in each of the axial, coronal, and sagittal cross-sections, comparing against the native slices using SSIM (\uparrow) (Wang et al., 2004), PSNR (\uparrow) and LPIPS (\downarrow) (Zhang et al., 2018). As reconstruction time t is a key consideration for clinical adoption, we report metrics at $t = 5$ and $t = 20$ minutes, and at convergence ($t = \infty$).

Limited and perturbed inputs. We repeat the evaluation with only 50% coverage (80 evenly spaced axial slices), and add small, random rotations about the x and y axes ($\theta \sim U(-5^\circ, +5^\circ)$) to mimic the sonographer’s hand motion and fetal movement encountered in practice.

Baselines and protocol. We benchmark against three SOTA *ultrasound* reconstruction models: UltraNerf (Wysocki et al., 2023) (physics-informed ray tracing); RapidVol (Eid et al., 2025) (hybrid implicit-explicit); and ImplicitVol (Yeung et al., 2024) (fully implicit, NeRF-like). To satisfy UltraNerf’s parallel-ray assumption, curvilinear inputs are converted from polar to Cartesian.

5.3 CLINICIANS’ SURVEY OF RECONSTRUCTION QUALITY

We assess perceived clinical realism and fidelity of UltraGauss’ reconstructions in a two-part reader study with expert sonographers. All survey material was de-identified.

Study A: Pairwise realism preference. Participants compared natively acquired 3D ultrasound images with reconstructions from RapidVol, ImplicitVol, and UltraGauss variants. We evaluated four routinely assessed fetal brain planes at the 20-week anomaly scan: mid-axial, transthalamic, transventricular, and mid-coronal (Salomon et al., 2022). Assessments spanned five models (ImplicitVol, RapidVol, UltraGauss-100K, UltraGauss-300K, UltraGauss-2M), four fetal scans, and three training budgets $t \in \{5, 20, \infty\}$ minutes. Each participant viewed 10 randomized image pairs (order of questions and left/right placement shuffled) and selected the most realistic scan, or indicated “no preference” if indistinguishable. Respondents reported speciality training, years in practice, and confidence in fetal brain assessment.

Study B: Temporal fidelity (“Turing”) test. To examine convergence over time, a second survey presented nine UltraGauss reconstructions of a mid-coronal plane generated at $t = \{0.5, 1, 2, 3, 4, 5, 10, 15, 20\}$ minutes, alongside two identical ground-truth images. Inputs com-

378 prised 80 axial slices, covering only 50% of the 3D volume and orthogonal to the target plane.
 379 Participants labelled images as either a real ultrasound scan or an “AI reconstruction”. We randomly
 380 included duplicated ground-truth images as a control, to quantify variability in expert judgements.
 381

382 5.4 ULTRAGAUSS IN AN END-TO-END CLINICAL PIPELINE

383 We evaluate UltraGauss in a practical point-of-care workflow: a clinician acquires a freehand video
 384 (cinesweep) with a standard sensorless 2D probe, and then reconstructs a full 3D volume for retro-
 385 spective multiplanar review. This enables imputation of planes that were missed during scanning.

386 **Pipeline.** Each video frame first receives a 6-DoF pose via an ultrasound pose-estimation model
 387 (e.g., Ramesh et al. (2024); Di Vece et al. (2024); Yeung et al. (2022)). Frames and predicted poses
 388 are then fed into a suitable 3D reconstruction model (i.e., UltraGauss), or to a baseline (RapidVol).

389 **Evaluation protocol.** Ground-truth 3D for these fetal cinesweeps is unavailable, so we adopt a
 390 frame hold-out cross-validation scheme: for each video, we randomly partition the video frames into
 391 80% training and 20% testing. After reconstruction, we render the held-out slices at their predicted
 392 poses and compare the syntheses against the corresponding acquired frames. The residual captures
 393 aggregate error from both pose estimation and reconstruction inaccuracies. We run this end-to-end
 394 pipeline on all three fetal cinesweeps in Dataset B.

395 5.5 SHADOW MODELLING

396 We compare UltraGauss with this feature enabled against UltraNerf (Wysocki et al., 2023), using
 397 their publicly available Synthetic Liver linear-probe dataset.

399 6 RESULTS

401 6.1 QUALITY AND SPEED EVALUATION

403 **Summary.** We evaluate on three test sets of increasing difficulty (Fig. 4, more in Fig. B1). Across
 404 all datasets and training budgets $t \in \{5, 20, \infty\}$, UltraGauss outperforms RapidVol and ImplicitVol.

405 **Accuracy vs. time.** For near-real-time use ($t=5$ min), UltraGauss exceeds the best baseline by
 406 ≥ 0.20 SSIM. At convergence, UltraGauss-2M attains a mean SSIM of **0.995**. Variability across
 407 gestational ages is markedly lower: at least $10\times$ smaller variance than RapidVol and ImplicitVol.

408 **Capacity-time trade-offs.** Model capacity interacts with the time budget: (i) at $t=5$ min, ~ 100 K
 409 Gaussians perform best; (ii) for long runs ($\mathcal{O}(\text{hours})$), 2M Gaussians yield the highest final accuracy;
 410 (iii) ~ 300 K Gaussians offer a strong balance across budgets.

411 6.2 CLINICIANS’ SURVEY

413 **Participants.** We invited 12 expert sonographers specializing in fetal, pediatric, and general ultra-
 414 sound imaging, from hospitals in four countries (UK, Ghana, Denmark, and the Netherlands). Ten
 415 experts responded (consultant fetal surgeons and senior sonographers), with a mean of 18 years’
 416 experience (range: 7–30 years).

417 **Comparison of reconstruction methods:** Under matched short training budgets ($t \leq 20$ minutes),
 418 all clinicians preferred UltraGauss over RapidVol and ImplicitVol. At convergence ($t=\infty$), no clin-
 419 ician rated UltraGauss worse than the alternatives (Appendix C.2). Within UltraGauss, preference
 420 increased with capacity ($2M > 300K > 100K$), indicating that larger Gaussian sets yield more real-
 421 istic reconstructions.

422 **Temporal realism.** In the progressive (“Turing”) test, where the rendered novel-view was orthogo-
 423 nal to the input data, 70% of clinicians selected UltraGauss as more realistic than the ground-truth
 424 after just 4 minutes of training, rising to 80% after 15 minutes (Fig. 5; examples in Fig. B2). As a
 425 control for judgement variability, two identical ground-truth images were included; 40% of partic-
 426 ipants labelled one of these as AI, underscoring the inherent variability in expert assessments.

427 6.3 END-TO-END PIPELINE RESULTS

428 We evaluate end-to-end by rendering held-out frames at their *predicted* poses and comparing against
 429 the acquired images (Tab. 1; examples in Fig. 6; temporal curves and full visualizations in Ap-
 430 pendix D.2). Across videos and training budgets ($t \in \{5, 20, \infty\}$), UltraGauss achieves the highest
 431 SSIM/PSNR and lowest LPIPS, with especially large gains at short budgets which is relevant for
 time-critical scans and brief consultations. Averaged over all held-out frames, UltraGauss reaches

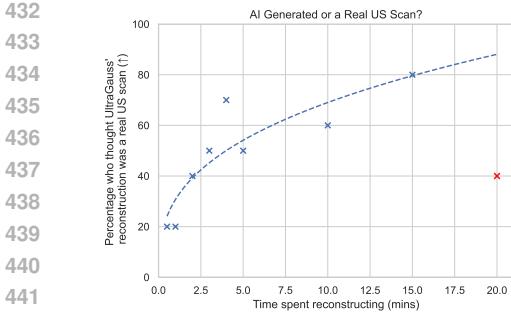


Figure 5: Clinicians’ survey results, asking whether an image is real or “AI Generated” (i.e. UltraGauss-generated).

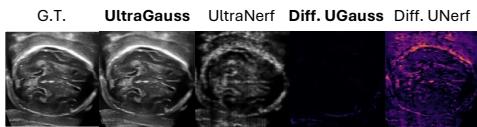


Figure 6: UltraGauss and UltraNerf being used to predict one of the withheld US video frames at its estimated pose. UltraGauss is nearly identical to the Ground Truth (G.T.) Image, as can be seen by the near-black absolute difference image (“Diff. UGauss”).

SSIM 0.91. Note that these scores reflect the *combined* error from both pose estimation and reconstruction.

6.4 SHADOW MODELLING

Enabling shadow modelling on UltraGauss improves SSIM by 0.005, PSNR by 0.2 dB and LPIPS by 0.004, with practically no time penalty. Whilst the metric gains are minimal, it can be seen in Fig. 7 that it leads to straighter, and more realistic acoustic shadows. More are shown in Fig. A1. Despite these improvements, we still lack some of the speckle arising from multiple-scattering which UltraNerf is able to capture, thanks to their inclusion of scattering density and intensity parameters. However, when the ground truth poses are not known due to acquisition from sensorless 2D probes (Sec. 5.4), or if the ultrasound rays are not perfectly straight (e.g. due to using a curvilinear probe and then transforming the image from polar to Cartesian space, as in Fig. 3), then UltraGauss can yield more robust results, since it is less dependent on a calibrated physics model. This is evident by the results in Secs. 6.1 to 6.3. UltraGauss is also much quicker – $6.94 \times$ faster than UltraNerf.

7 CONCLUSION

We introduced **UltraGauss**, an ultrasound-specific Gaussian Splatting framework that serves as an *efficient approximation* to acoustic image formation via *probe-plane intersection* rendering. We derived closed-form χ^2 bounds for plane-intersection rasterisation and a two-phase, load-balanced CUDA pipeline, and proposed a triangular inverse-covariance (precision) parameterisation that stabilises optimisation at scale. On clinical datasets, UltraGauss delivers minute-level 2D \rightarrow 3D reconstructions with high fidelity, and expert sonographers consistently prefer its realism to prior methods. In an end-to-end cinesweep workflow, it enables retrospective multiplanar review from routine 2D acquisitions without additional hardware. By aligning the rendering model with ultrasound physics while retaining the efficiency of Gaussian splatting, UltraGauss provides a practical route to standardised volumetry and more accessible 3D ultrasound, particularly valuable across diverse clinical settings. Future work includes joint pose–reconstruction optimization, richer acoustic effects beyond Beer–Lambert attenuation, and broader evaluation across anatomies and scanners.

	Model	Video 1	Video 2	Video 3	Avg.	Std.
Test SSIM \uparrow	ImplicitVol	0.674	0.797	0.772	0.747	0.065
	RapidVol	0.745	0.799	0.760	0.768	0.028
	UltraNerf	0.446	0.626	0.521	0.531	0.091
	UltraGauss	0.928	0.905	0.910	0.914	0.012
Train SSIM \uparrow	ImplicitVol	0.914	0.899	0.893	0.902	0.010
	RapidVol	0.871	0.884	0.880	0.878	0.007
	UltraNerf	0.489	0.591	0.598	0.559	0.061
	UltraGauss	0.959	0.957	0.939	0.952	0.011

Table 1: Quantitative results of reconstruction performance on cinesweep videos. The SSIM scores shown are the average across all the frames held-out for testing/used in training at $t = \infty$. Best scores are highlighted in **bold**. \uparrow indicates higher is better.

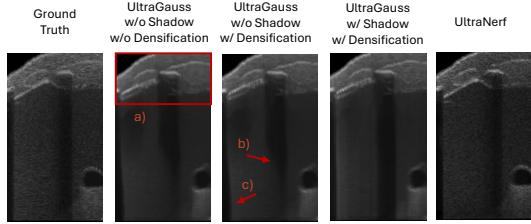


Figure 7: Test-set images rendered with UltraGauss (toggling Gaussian densification or shadow modelling), and UltraNerf. Without densification, speckle is lacking (a). Without shadow modelling, acoustic shadows can abruptly stop (b), or bulge out and become curved instead of straight (c).

486 REFERENCES
487

488 John E. Aldrich. Basic physics of ultrasound imaging:. *Critical Care Medicine*, 35(Suppl):S131–
489 S137, May 2007. ISSN 0090-3493. doi: 10.1097/01.CCM.0000260624.99430.22. URL <http://journals.lww.com/00003246-200705001-00003>.

490 Beryl R. Benacerraf. Three-dimensional fetal sonography. *Journal of Ultrasound in Medicine*,
491 21(10):1063–1067, 2002. doi: <https://doi.org/10.7863/jum.2002.21.10.1063>. URL <https://onlinelibrary.wiley.com/doi/abs/10.7863/jum.2002.21.10.1063>.

492 Sierra Bonilla, Shuai Zhang, Dimitrios Psychogios, Danail Stoyanov, Francisco Vasconcelos,
493 and Sophia Bano. Gaussian Pancakes: Geometrically-Regularized 3D Gaussian Splatting for
494 Realistic Endoscopic Reconstruction. In Marius George Linguraru, Qi Dou, Aasa Feragen,
495 Stamatia Giannarou, Ben Glocker, Karim Lekadir, and Julia A. Schnabel (eds.), *Medical Im-*
496 *age Computing and Computer Assisted Intervention – MICCAI 2024*, volume 15006, pp. 274–
497 283. Springer Nature Switzerland, Cham, 2024. ISBN 978-3-031-72088-8 978-3-031-72089-5.
498 doi: 10.1007/978-3-031-72089-5_26. URL https://link.springer.com/10.1007/978-3-031-72089-5_26. Series Title: Lecture Notes in Computer Science.

499 Yuanhao Cai, Yixun Liang, Jiahao Wang, Angtian Wang, Yulun Zhang, Xiaokang Yang, Zongwei
500 Zhou, and Alan Yuille. Radiative Gaussian Splatting for Efficient X-Ray Novel View Synthesis. In
501 Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gürkaynak Varol (eds.),
502 *Computer Vision – ECCV 2024*, volume 15059, pp. 283–299. Springer Nature Switzerland, Cham,
503 2025. ISBN 978-3-031-73231-7 978-3-031-73232-4. doi: 10.1007/978-3-031-73232-4_16. URL
504 https://link.springer.com/10.1007/978-3-031-73232-4_16. Series Title:
505 Lecture Notes in Computer Science.

506 Nvidia Corp. Thrust library. <https://github.com/NVIDIA/thrust?tab=readme-ov-file>, 2020.

507 Chiara Di Vece, Maela Le Lous, Brian Dromey, Francisco Vasconcelos, Anna L. David, Donald
508 Peebles, and Danail Stoyanov. Ultrasound Plane Pose Regression: Assessing Generalized
509 Pose Coordinates in the Fetal Brain. *IEEE Transactions on Medical Robotics and Bionics*, 6(1):
510 41–52, February 2024. ISSN 2576-3202. doi: 10.1109/TMRB.2023.3328638. URL <https://ieeexplore.ieee.org/document/10302329/authors#authors>. Conference
511 Name: IEEE Transactions on Medical Robotics and Bionics.

512 Mark C. Eid, Pak-Hei Yeung, Madeleine K. Wyburd, João F. Henriques, and Ana I.L. Namburete.
513 Rapidvol: Rapid Reconstruction of 3D Ultrasound Volumes from Sensorless 2D Scans. In 2025
514 *IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5, April 2025. doi: 10.
515 1109/ISBI60581.2025.10980994. URL <https://ieeexplore.ieee.org/abstract/document/10980994>. ISSN: 1945-8452.

516 François Gaits, Nicolas Mellado, and Adrian Basarab. Ultrasound Volume Reconstruction from
517 2D Freehand Acquisitions Using Neural Implicit Representations. In 2024 *IEEE International*
518 *Symposium on Biomedical Imaging (ISBI)*, pp. 1–5, May 2024. doi: 10.1109/ISBI56570.2024.
519 10635582. URL <https://ieeexplore.ieee.org/document/10635582>. ISSN:
520 1945-8452.

521 Jiaxin Guo, Jiangliu Wang, Di Kang, Wenzhen Dong, Wenting Wang, and Yun-hui Liu. Free-
522 SurGS: SfM-Free 3D Gaussian Splatting for Surgical Scene Reconstruction. In Marius George
523 Linguraru, Qi Dou, Aasa Feragen, Stamatia Giannarou, Ben Glocker, Karim Lekadir, and Ju-
524 lia A. Schnabel (eds.), *Medical Image Computing and Computer Assisted Intervention – MIC-*
525 *CAI 2024*, volume 15007, pp. 350–360. Springer Nature Switzerland, Cham, 2024a. ISBN
526 978-3-031-72103-8 978-3-031-72104-5. doi: 10.1007/978-3-031-72104-5_34. URL https://link.springer.com/10.1007/978-3-031-72104-5_34. Series Title: Lecture
527 Notes in Computer Science.

528 Ziwen Guo, Zi Fang, and Zhuang Fu. UIRe-NeRF: 3D Ultrasound Imaging through Neural
529 Rendering with Ultrasound Reflection Direction Parameterization, September 2024b. URL
530 <http://arxiv.org/abs/2408.00860>. arXiv:2408.00860 [cs].

540 Linde S. Hesse, Moska Aliasi, Felipe Moser, Monique C. Haak, Weidi Xie, Mark Jenkinson, and
 541 Ana I. L. Namburete. Subcortical segmentation of the fetal brain in 3D ultrasound using deep
 542 learning. *NeuroImage*, 254:119117, July 2022. ISSN 1053-8119. doi: 10.1016/j.neuroimage.
 543 2022.119117. URL <https://www.sciencedirect.com/science/article/pii/S1053811922002452>.

545 Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting
 546 for geometrically accurate radiance fields. In *SIGGRAPH 2024 Conference Papers*. Association
 547 for Computing Machinery, 2024. doi: 10.1145/3641519.3657428.

548 Raphael Hunger. Floating Point Operations in Matrix-Vector Calculus. Technical report, Technische
 549 Universitat Munchen (TUM), 2007.

551 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3D Gaussian
 552 Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4):1–14,
 553 August 2023. ISSN 0730-0301, 1557-7368. doi: 10.1145/3592433. URL <https://dl.acm.org/doi/10.1145/3592433>.

555 Qi Li, Shaheer U. Saeed, Yuliang Huang, Mingyuan Luo, Zhongnuo Yan, Jiongquan Chen, Xin
 556 Yang, Dong Ni, Nektarios Winter, Phuc Nguyen, Lucas Steinberger, Caelan Haney, Yuan Zhao,
 557 Mingjie Jiang, Bowen Ren, SiYeoul Lee, Seonho Kim, MinKyung Seo, MinWoo Kim, Yimeng
 558 Dou, Zhiwei Zhang, Yin Li, Tomy Varghese, Dean C. Barratt, Matthew J. Clarkson, Tom Ver-
 559 cauteren, and Yipeng Hu. TUS-REC2024: A Challenge to Reconstruct 3D Freehand Ultrasound
 560 Without External Tracker, November 2025. URL <http://arxiv.org/abs/2506.21765>.
 561 arXiv:2506.21765 [eess].

562 Hengyu Liu, Yifan Liu, Chenxin Li, Wuyang Li, and Yixuan Yuan. LGS: A Light-Weight 4D
 563 Gaussian Splatting for Efficient Surgical Scene Reconstruction. In Marius George Linguraru,
 564 Qi Dou, Aasa Feragen, Stamatia Giannarou, Ben Glocker, Karim Lekadir, and Julia A. Schnabel
 565 (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume
 566 15003, pp. 660–670. Springer Nature Switzerland, Cham, 2024. ISBN 978-3-031-72383-4 978-3-
 567 031-72384-1. doi: 10.1007/978-3-031-72384-1_62. URL https://link.springer.com/10.1007/978-3-031-72384-1_62. Series Title: Lecture Notes in Computer Science.

569 E. Merz and C. Welter. 2D and 3D Ultrasound in the evaluation of normal and abnormal
 570 fetal anatomy in the second and third trimesters in a level III center. *Ultraschall in Der
 571 Medizin (Stuttgart, Germany: 1980)*, 26(1):9–16, February 2005. ISSN 0172-4614. doi:
 572 10.1055/s-2004-813947.

573 Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and
 574 Ren Ng. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis, August 2020.
 575 URL <http://arxiv.org/abs/2003.08934>. arXiv:2003.08934 [cs].

577 Ana I. L. Namburete, Bartłomiej W. Papież, Michelle Fernandes, Madeleine K. Wyburd, Linde S.
 578 Hesse, Felipe A. Moser, Leila Cheikh Ismail, Robert B. Gunier, Waney Squier, Eric O. Ohuma,
 579 Maria Carvalho, Yasmin Jaffer, Michael Gravett, Qingqing Wu, Ann Lambert, Adele Win-
 580 sey, María C. Restrepo-Méndez, Enrico Bertino, Manorama Purwar, Fernando C. Barros, Alan
 581 Stein, J. Alison Noble, Zoltán Molnár, Mark Jenkinson, Zulfiqar A. Bhutta, Aris T. Papa-
 582 georgiou, José Villar, and Stephen H. Kennedy. Normative spatiotemporal fetal brain matu-
 583 ration with satisfactory development at 2 years. *Nature*, pp. 1–9, October 2023. ISSN 1476-
 584 4687. doi: 10.1038/s41586-023-06630-3. URL <https://www.nature.com/articles/s41586-023-06630-3>. Publisher: Nature Publishing Group.

586 Ana I.L. Namburete, Richard V. Stebbing, Bryn Kemp, Mohammad Yaqub, Aris T. Papageorgiou,
 587 and J. Alison Noble. Learning-based prediction of gestational age from ultrasound images
 588 of the fetal brain. *Medical Image Analysis*, 21(1):72–86, 2015. ISSN 1361-8415. doi:
 589 <https://doi.org/10.1016/j.media.2014.12.006>. URL <https://www.sciencedirect.com/science/article/pii/S136184151400190X>.

591 Thomas R. Nelson and Dolores H. Pretorius. Three-dimensional ultrasound imaging. *Ultra-
 592 sound in Medicine and Biology*, 24(9):1243–1270, Dec 1998. ISSN 0301-5629. doi: 10.
 593 1016/S0301-5629(98)00043-X. URL [https://doi.org/10.1016/S0301-5629\(98\)00043-X](https://doi.org/10.1016/S0301-5629(98)00043-X).

594 Emmanouil Nikolakakis, Utkarsh Gupta, Jonathan Vengosh, Justin Bui, and Razvan Marinescu.
 595 GaSpCT: Gaussian Splatting for Novel CT Projection View Synthesis, April 2024. URL <http://arxiv.org/abs/2404.03126>. arXiv:2404.03126 [eess].
 596

597 Aris T. Papageorgiou, Stephen H. Kennedy, Laurent J. Salomon, Douglas G. Altman, Eric O.
 598 Ohuma, William Stones, Michael G. Gravett, Fernando C. Barros, Cesar Victora, Manorama
 599 Purwar, Yasmin Jaffer, Julia A. Noble, Enrico Bertino, Ruyan Pang, Leila Cheikh Ismail, Ann
 600 Lambert, Zulfiqar A. Bhutta, and José Villar. The INTERGROWTH-21st fetal growth standards:
 601 toward the global integration of pregnancy and pediatric care. *American Journal of Obstetrics
 602 & Gynecology*, 218(2):S630–S640, February 2018. ISSN 0002-9378, 1097-6868. doi: 10.
 603 1016/j.ajog.2018.01.011. URL [https://www.ajog.org/article/S0002-9378\(18\)30012-7/fulltext](https://www.ajog.org/article/S0002-9378(18)30012-7/fulltext). Publisher: Elsevier.
 604

605 Tengya Peng, Ruyi Zha, Zhen Li, Xiaofeng Liu, and Qing Zou. Three-Dimensional MRI Recon-
 606 struction with Gaussian Representations: Tackling the Undersampling Problem, February 2025.
 607 URL <http://arxiv.org/abs/2502.06510>. arXiv:2502.06510 [eess].
 608

609 Alexander Ej. Powles, David J. Martin, Irving Tp. Wells, and Christopher R. Goodwin. Physics
 610 of ultrasound. *Anaesthesia & Intensive Care Medicine*, 19(4):202–205, April 2018. ISSN
 611 14720299. doi: 10.1016/j.mpaic.2018.01.005. URL <https://linkinghub.elsevier.com/retrieve/pii/S1472029918300171>.
 612

613 Ravi Ramamoorthi. MODELING ILLUMINATION VARIATION WITH SPHERICAL HAR-
 614 MONICS. In *Face Processing*, pp. 385–424. Elsevier, 2006. ISBN 978-0-12-088452-
 615 0. doi: 10.1016/B978-012088452-0/50013-3. URL <https://linkinghub.elsevier.com/retrieve/pii/B9780120884520500133>.
 616

617 Jayroop Ramesh, Nicola Dinsdale, Pak-Hei Yeung, and Ana I. L. Namburete. Geometric Transfor-
 618 mation Uncertainty for Improving 3D Fetal Brain Pose Prediction from Freehand 2D Ultrasound
 619 Videos. In Marius George Linguraru, Qi Dou, Aasa Feragen, Stamatia Giannarou, Ben Glocker,
 620 Karim Lekadir, and Julia A. Schnabel (eds.), *Medical Image Computing and Computer Assisted
 621 Intervention – MICCAI 2024*, volume 15001, pp. 419–429. Springer Nature Switzerland, Cham,
 622 2024. ISBN 978-3-031-72377-3 978-3-031-72378-0. doi: 10.1007/978-3-031-72378-0_39. URL
 623 https://link.springer.com/10.1007/978-3-031-72378-0_39. Series Title:
 624 Lecture Notes in Computer Science.
 625

626 Samuel Rota Bulò, Lorenzo Porzi, and Peter Kotschieder. Revising Densification in Gaussian
 627 Splatting. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and
 628 Gü̈l Varol (eds.), *Computer Vision – ECCV 2024*, pp. 347–362, Cham, 2025. Springer Nature
 629 Switzerland. ISBN 978-3-031-73036-8. doi: 10.1007/978-3-031-73036-8_20.

630 L. J. Salomon, Z. Alfirevic, V. Berghella, C. M. Bilardo, G. E. Chalouhi, F. Da Silva Costa,
 631 E. Hernandez-Andrade, G. Malinger, H. Munoz, D. Paladini, F. Prefumo, A. Sotiriadis, A. Toi,
 632 and W. Lee. ISUOG Practice Guidelines (updated): performance of the routine mid-trimester
 633 fetal ultrasound scan. *Ultrasound in Obstetrics & Gynecology*, 59(6):840–856, June 2022. ISSN
 634 0960-7692, 1469-0705. doi: 10.1002/uog.24888. URL <https://obgyn.onlinelibrary.wiley.com/doi/10.1002/uog.24888>.
 635

636 Mriganka Sarmah, Arambam Neelima, and Heisnam Rohen Singh. Survey of methods and
 637 principles in three-dimensional reconstruction from two-dimensional medical images. *Visual
 638 Computing for Industry, Biomedicine, and Art*, 6(1):15, July 2023. ISSN 2524-4442. doi:
 639 10.1186/s42492-023-00142-7. URL <https://vciba.springeropen.com/articles/10.1186/s42492-023-00142-7>.
 640

641 Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Confer-
 642 ence on computer vision and pattern recognition (CVPR)*, 2016.
 643

644 Ole Vegard Solberg, Frank Lindseth, Hans Torp, Richard E. Blake, and Toril A. Nagelhus Hernes.
 645 Freehand 3D Ultrasound Reconstruction Algorithms—A Review. *Ultrasound in Medicine
 646 & Biology*, 33(7):991–1009, July 2007. ISSN 0301-5629. doi: 10.1016/j.ultrasmedbio.
 647 2007.02.015. URL <https://www.sciencedirect.com/science/article/pii/S0301562907001081>.

648 Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image Quality Assessment: From Error
 649 Visibility to Structural Similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, April
 650 2004. ISSN 1057-7149. doi: 10.1109/TIP.2003.819861. URL <http://ieeexplore.ieee.org/document/1284395/>.

652 653 Magdalena Wysocki, Mohammad Farid Azampour, Christine Eilers, Benjamin Busam, Mehrdad
 654 Salehi, and Nassir Navab. Ultra-NeRF: Neural radiance fields for ultrasound imaging. In *MDL*,
 655 2023.

656 657 Weixing Xie, Junfeng Yao, Xianpeng Cao, Qiqin Lin, Zerui Tang, Xiao Dong, and Xiaohu Guo.
 658 SurgicalGaussian: Deformable 3D Gaussians for High-Fidelity Surgical Scene Reconstruction.
 659 In Marius George Linguraru, Qi Dou, Aasa Feragen, Stamatia Giannarou, Ben Glocker, Karim
 660 Lekadir, and Julia A. Schnabel (eds.), *Medical Image Computing and Computer Assisted In-
 661 tervention – MICCAI 2024*, volume 15006, pp. 617–627. Springer Nature Switzerland, Cham,
 662 2024. ISBN 978-3-031-72088-8 978-3-031-72089-5. doi: 10.1007/978-3-031-72089-5_58. URL
 663 https://link.springer.com/10.1007/978-3-031-72089-5_58. Series Title:
 664 Lecture Notes in Computer Science.

665 Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari,
 666 Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An Open-Source Library
 667 for Gaussian Splatting, September 2024. URL <http://arxiv.org/abs/2409.06765>.
 arXiv:2409.06765 [cs].

668 Pak-Hei Yeung, Moska Aliasi, Monique Haak, The INTERGROWTH-21st Consortium, Weidi Xie,
 669 and Ana I. L. Namburete. Adaptive 3D Localization of 2D Freehand Ultrasound Brain Images,
 670 September 2022. URL <http://arxiv.org/abs/2209.05477>. arXiv:2209.05477 [cs,
 671 eess].

672 Pak-Hei Yeung, Linde Hesse, Moska Aliasi, Monique Haak, Weidi Xie, Ana IL Namburete, and
 673 others. ImplicitVol: Sensorless volumetric reconstruction of fetal brain freehand ultrasound scans
 674 with deep implicit representation. *Medical Image Analysis*, 94:103147, 2024. Publisher: Elsevier.

675 Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo
 676 Kanazawa. Plenoxels: Radiance Fields without Neural Networks, December 2021. URL
 677 <http://arxiv.org/abs/2112.05131>. arXiv:2112.05131 [cs].

678 Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-Splatting:
 679 Alias-free 3D Gaussian Splatting, November 2023. URL <http://arxiv.org/abs/2311.16493>. arXiv:2311.16493 [cs].

680 Ruyi Zha, Tao Jun Lin, Yuanhao Cai, Jiwen Cao, Yanhao Zhang, and Hongdong Li. R²-
 681 gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In A. Globerson,
 682 L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in
 683 neural information processing systems*, volume 37, pp. 44907–44934. Curran Associates, Inc.,
 684 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/4f550cb7b30b59553e50cd08a9dbf068-Paper-Conference.pdf.

685 Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unrea-
 686 sonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Confer-
 687 ence on Computer Vision and Pattern Recognition*, pp. 586–595, Salt Lake City, UT, June
 688 2018. IEEE. ISBN 978-1-5386-6420-9. doi: 10.1109/CVPR.2018.00068. URL <https://ieeexplore.ieee.org/document/8578166/>.

689 M. Zwicker, H. Pfister, J. Van Baar, and M. Gross. EWA volume splatting. In *Proceedings Visu-
 690 alization, 2001. VIS '01.*, pp. 29–538, San Diego, CA, USA, 2001. IEEE. ISBN 978-0-7803-
 691 7200-9. doi: 10.1109/VISUAL.2001.964490. URL <http://ieeexplore.ieee.org/document/964490/>.

692 693 694 695 696 697 698 699 700 701

702

A ABLATIONS

703

A.1 c AND α OR ONE COMBINED PARAMETER?

704 Since ultrasound is a grayscale modality, requiring only one color channel, it is possible to combine
 705 each Gaussian's single-channel color c_i (which ranges from 0 to 1) and opacity α_i (also ranges
 706 from 0 to 1) into a single learnable parameter, say ρ_i , as for example is done by Zha et al. (2024).
 707 We could then multiply this by Gaussian i 's probability at point \mathbf{x} . This would then be repeated
 708 for all the other Gaussians and they can then either sum be summed together (but stopping if the
 709 sum reaches 1), or have the average taken. Since we do not splat Gaussians onto an image plane
 710 and sort them by depth from front to back, it is not obvious how to decide which 3D Gaussians to
 711 exclude from the sum if it reaches 1, so we would have to take the average. Expressed formally
 712 this potential approach is:

$$714 \hat{\rho}_i(x) = \rho_i \exp \left(-\frac{1}{2} \underbrace{(x_{|0} - \mu_i^{3D})^T (\Sigma_i^{3D})^{-1} (x_{|0} - \mu_i^{3D})}_{\text{3D squared Mahalanobis distance}} \right) \\ 715 \\ 716 \hat{c}(x) = \frac{\sum_i^n \hat{\rho}_i(x)}{n} \quad (A1)$$

720 Instead, we propose a formulation as in Eq. 7. Ignoring α_{BG} and c_{BG} for now, we are essentially
 721 proposing:

$$722 \hat{c}(x) = \frac{\sum_i^n \hat{\alpha}_i(x) c_i}{\sum_i^n \hat{\alpha}_i(x)} \quad (A2)$$

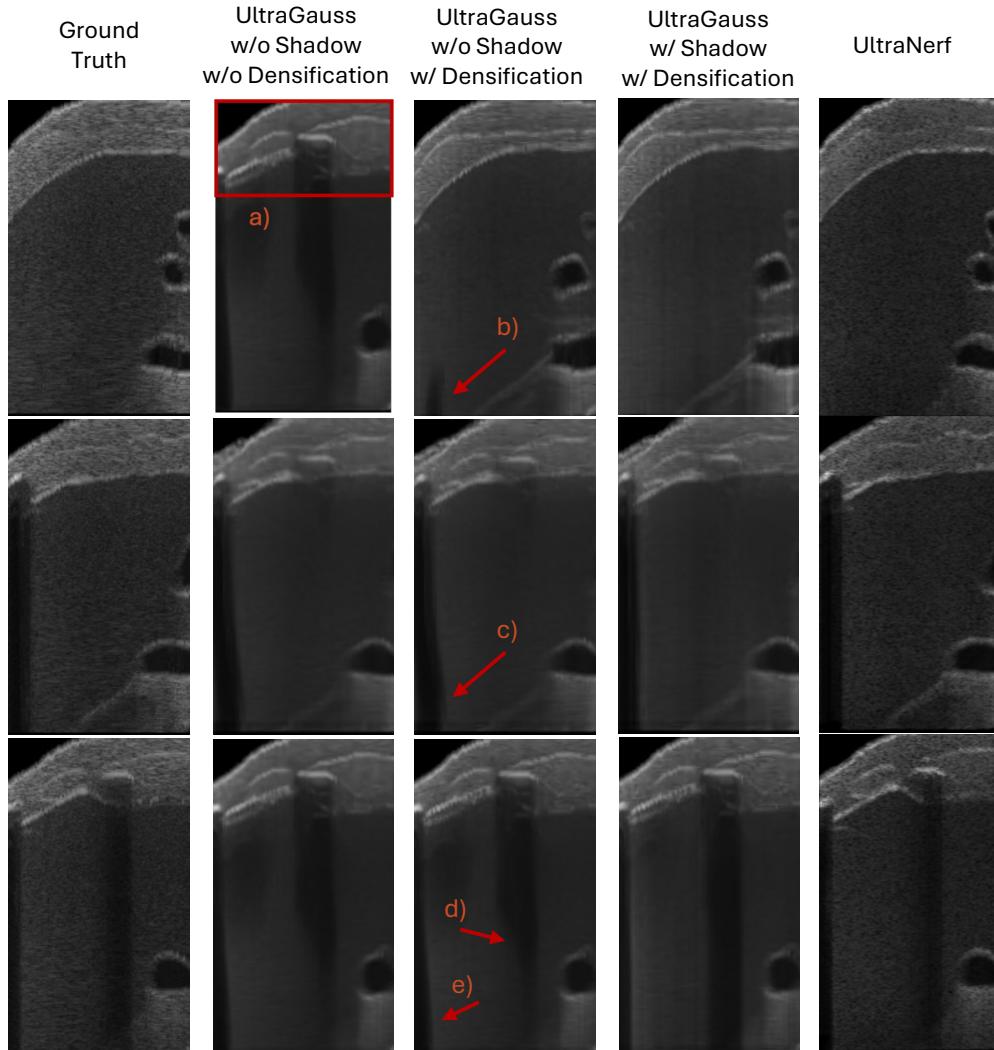
724 which mathematically is a *weighted* average. Compare this with the *arithmetic* average when one
 725 parameter is used. The former is naturally more powerful, and Table A1 confirms this. In our
 726 application, intuitively, α_i dictates the importance of each Gaussian, and how strongly its color
 727 should feature in the final pixel color \hat{c} .

728 Table A1: Using c_i and α_i vs. only one parameter. Scores are the avg. across all three standard test
 729 views, for four 20 week fetuses.

	Input Set = 80 Axial Slices			Input Set = 160 Axial Slices		
	SSIM (↑)	PSNR (↑)	LPIPS (↓)	SSIM (↑)	PSNR (↑)	LPIPS (↓)
One Param.	0.931 ± 0.012	3.42 ± 0.100	0.097 ± 0.026	0.947 ± 0.011	3.66 ± 0.120	0.085 ± 0.026
c_i and α_i	0.967 ± 0.012	3.70 ± 0.193	0.032 ± 0.015	0.990 ± 0.009	4.42 ± 0.382	0.011 ± 0.009

756 A.2 THE IMPACT OF SHADOW MODELLING AND DENSIFICATION
757

758 Fig. A1 visibly shows the beneficial impact adaptive Gaussian densification and sparsification, once
759 modified and implemented appropriately, can have. Greater sharpness and speckle can be seen in
760 the top quarter of the images (a) when Densification is turned on (c.f. 2nd column with 3rd column).
761 When Shadow Modelling is turned on, view-dependent acoustic shadows become more realistic.
762 Comparing columns 3 and 4, one can see that shadows or dark patches no longer appear where they
763 shouldn't (b), shadows no longer bulge out and curve, but rather become straight (c & e), and they
764 also no longer abruptly stop (d).



798 Figure A1: Test-set images rendered using UltraGauss-300K, with (w/) and without (w/o) shadow
799 modelling or adaptive Gaussian densification & sparsification. A fully ultrasound physics based
800 model, UltraNerf (Wysocki et al., 2023) serves as a baseline.

801
802
803
804
805
806
807
808
809

B EXPERIMENT 5.2 (3D SCANS) - ADDITIONAL RESULTS

Input Set	Time t	Model	Training SSIM (↑)	Test Slices, 160 Linearly Spaced:								Average Across all 3 views			
				Axial		Coronal		Sagittal		LPIPS (↑)		SSIM (↑)	PSNR (↑)	SSIM (↑)	PSNR (↑)
5 mins	ImplicitVol	ImplicitVol	0.850 ± 0.043	0.203 ± 0.056	26.44 ± 4.07	0.698 ± 0.052	23.07 ± 2.04	0.866 ± 0.056	0.707 ± 0.056	25.78 ± 2.99	0.288 ± 0.066	0.732 ± 0.052	16.03 ± 1.03	0.262 ± 0.060	
		RapidVol	0.240 ± 0.047	0.240 ± 0.051	15.96 ± 1.39	0.648 ± 0.059	0.235 ± 0.052	15.83 ± 1.29	0.616 ± 0.036	0.258 ± 0.052	16.30 ± 1.74	0.647 ± 0.051	0.244 ± 0.052	16.03 ± 1.47	0.632 ± 0.042
		UltraNerf	0.531 ± 0.038	0.494 ± 0.037	19.39 ± 2.35	0.318 ± 0.027	0.412 ± 0.049	16.71 ± 1.46	0.412 ± 0.028	0.252 ± 0.056	14.47 ± 0.54	0.585 ± 0.015	0.386 ± 0.061	16.85 ± 1.45	0.388 ± 0.024
	UltraGauss - 100K	UltraGauss - 100K	0.988 ± 0.003	0.982 ± 0.003	43.07 ± 2.92	0.913 ± 0.002	0.982 ± 0.004	40.28 ± 2.05	0.979 ± 0.004	44.25 ± 2.46	0.938 ± 0.004	42.72 ± 2.46	0.922 ± 0.004	42.72 ± 2.46	0.926 ± 0.005
		ImplicitVol	0.988 ± 0.002	0.987 ± 0.002	41.17 ± 1.85	0.983 ± 0.003	0.984 ± 0.003	39.31 ± 1.29	0.942 ± 0.005	0.983 ± 0.003	42.11 ± 2.08	0.949 ± 0.006	0.985 ± 0.003	40.86 ± 1.74	0.938 ± 0.005
		RapidVol	0.855 ± 0.051	0.855 ± 0.051	29.57 ± 4.19	0.140 ± 0.046	0.704 ± 0.047	23.47 ± 1.55	0.265 ± 0.052	0.708 ± 0.051	26.22 ± 2.90	0.273 ± 0.064	0.756 ± 0.050	26.42 ± 2.88	0.226 ± 0.054
160 Linearly spaced Axial Slices	ImplicitVol	ImplicitVol	0.590 ± 0.060	0.561 ± 0.047	24.67 ± 4.88	0.225 ± 0.106	0.788 ± 0.087	23.96 ± 3.64	0.261 ± 0.096	0.789 ± 0.086	26.61 ± 4.13	0.222 ± 0.077	0.756 ± 0.086	26.78 ± 4.28	0.238 ± 0.093
		UltraNerf	0.590 ± 0.002	0.991 ± 0.002	45.37 ± 3.22	0.005 ± 0.001	0.987 ± 0.003	43.44 ± 1.79	0.413 ± 0.056	0.260 ± 0.055	14.47 ± 0.76	0.588 ± 0.013	0.421 ± 0.047	17.44 ± 1.32	0.432 ± 0.041
		UltraGauss - 100K	0.998 ± 0.001	0.997 ± 0.001	46.61 ± 2.88	0.993 ± 0.001	0.994 ± 0.002	43.44 ± 1.79	0.912 ± 0.002	0.986 ± 0.003	45.91 ± 3.88	0.928 ± 0.005	0.986 ± 0.003	45.51 ± 2.90	0.016 ± 0.003
	UltraGauss - 300K	ImplicitVol	0.986 ± 0.002	0.987 ± 0.002	41.17 ± 1.75	0.031 ± 0.003	0.978 ± 0.004	39.36 ± 1.20	0.050 ± 0.005	0.979 ± 0.004	42.10 ± 1.89	0.055 ± 0.006	0.980 ± 0.004	40.60 ± 1.62	0.045 ± 0.005
		RapidVol	0.882 ± 0.004	0.882 ± 0.004	42.01 ± 3.72	0.103 ± 0.003	0.975 ± 0.005	40.63 ± 2.41	0.022 ± 0.003	0.974 ± 0.006	45.82 ± 3.45	0.027 ± 0.006	0.977 ± 0.005	42.82 ± 3.19	0.021 ± 0.004
		UltraNerf	0.680 ± 0.080	0.680 ± 0.080	27.78 ± 3.78	0.004 ± 0.001	0.988 ± 0.003	42.48 ± 1.95	0.015 ± 0.003	0.977 ± 0.004	44.70 ± 0.99	0.030 ± 0.006	0.986 ± 0.003	44.64 ± 2.98	0.016 ± 0.003
20 mins	Convergence	UltraGauss - 100K	0.991 ± 0.002	0.982 ± 0.002	45.76 ± 1.91	0.004 ± 0.002	0.993 ± 0.002	44.50 ± 1.79	0.009 ± 0.001	0.979 ± 0.004	46.69 ± 4.41	0.023 ± 0.004	0.950 ± 0.002	47.19 ± 2.26	0.011 ± 0.002
		UltraGauss - 300K	0.998 ± 0.000	0.998 ± 0.000	50.38 ± 3.59	0.001 ± 0.000	0.993 ± 0.000	48.26 ± 2.86	0.000 ± 0.000	0.995 ± 0.001	46.29 ± 1.68	0.007 ± 0.001	0.989 ± 0.002	49.63 ± 3.52	0.011 ± 0.002
		UltraGauss - 2M	0.999 ± 0.000	1.000 ± 0.000	54.26 ± 2.86	0.000 ± 0.000	0.995 ± 0.001	45.29 ± 1.68	0.007 ± 0.001	0.990 ± 0.001	46.29 ± 1.68	0.007 ± 0.001	0.995 ± 0.001	50.06 ± 2.89	0.006 ± 0.001
	ImplicitVol	ImplicitVol	0.560 ± 0.015	0.960 ± 0.016	38.17 ± 4.38	0.041 ± 0.014	0.687 ± 0.042	36.69 ± 1.42	0.243 ± 0.041	0.685 ± 0.052	36.79 ± 2.39	0.246 ± 0.053	0.688 ± 0.053	38.95 ± 1.03	0.176 ± 0.036
		RapidVol	0.982 ± 0.004	0.982 ± 0.004	42.01 ± 3.72	0.103 ± 0.003	0.975 ± 0.005	40.63 ± 2.41	0.022 ± 0.003	0.974 ± 0.006	45.82 ± 3.45	0.027 ± 0.006	0.977 ± 0.005	42.82 ± 3.19	0.021 ± 0.004
		UltraNerf	0.680 ± 0.080	0.680 ± 0.080	27.78 ± 3.78	0.004 ± 0.001	0.988 ± 0.003	42.48 ± 1.95	0.015 ± 0.003	0.977 ± 0.004	44.70 ± 4.41	0.028 ± 0.006	0.986 ± 0.003	44.64 ± 2.98	0.016 ± 0.003
80 Linearly sliced Axial Slices	Convergence	UltraGauss - 100K	0.991 ± 0.003	0.982 ± 0.003	40.99 ± 2.34	0.014 ± 0.003	0.977 ± 0.004	38.03 ± 2.20	0.025 ± 0.004	0.975 ± 0.005	40.70 ± 2.85	0.033 ± 0.005	0.989 ± 0.007	38.09 ± 2.05	0.046 ± 0.005
		UltraGauss - 300K	0.998 ± 0.000	0.991 ± 0.003	43.99 ± 2.34	0.014 ± 0.002	0.977 ± 0.002	37.36 ± 1.82	0.050 ± 0.004	0.974 ± 0.006	40.28 ± 2.41	0.056 ± 0.006	0.977 ± 0.007	38.65 ± 2.74	0.027 ± 0.005
		UltraGauss - 2M	0.986 ± 0.002	0.981 ± 0.003	39.63 ± 1.91	0.032 ± 0.004	0.974 ± 0.005	36.30 ± 0.006	0.036 ± 0.008	0.965 ± 0.008	37.60 ± 2.09	0.030 ± 0.006	0.972 ± 0.007	38.65 ± 2.47	0.024 ± 0.005
	ImplicitVol	ImplicitVol	0.929 ± 0.013	0.789 ± 0.039	26.09 ± 2.63	0.172 ± 0.020	0.659 ± 0.052	23.34 ± 1.75	0.341 ± 0.031	0.659 ± 0.050	26.48 ± 2.64	0.331 ± 0.037	0.695 ± 0.047	25.30 ± 3.34	0.281 ± 0.029
		RapidVol	0.622 ± 0.153	0.622 ± 0.168	23.88 ± 1.53	0.025 ± 0.007	0.723 ± 0.169	23.06 ± 1.90	0.306 ± 0.011	0.724 ± 0.156	24.71 ± 1.56	0.321 ± 0.017	0.730 ± 0.169	24.71 ± 1.58	0.285 ± 0.119
		UltraNerf	0.611 ± 0.091	0.578 ± 0.097	21.70 ± 2.35	0.307 ± 0.048	0.445 ± 0.059	18.46 ± 1.12	0.068 ± 0.046	0.456 ± 0.049	14.23 ± 0.65	0.584 ± 0.028	0.426 ± 0.068	17.63 ± 1.37	0.432 ± 0.041
80 Linearly sliced Axial Slices with ±5° Jitter	Convergence	UltraGauss - 100K	0.994 ± 0.001	0.974 ± 0.005	38.59 ± 3.16	0.019 ± 0.003	0.961 ± 0.009	36.22 ± 1.81	0.051 ± 0.010	0.949 ± 0.010	40.99 ± 4.00	0.061 ± 0.010	0.961 ± 0.008	38.60 ± 2.99	0.043 ± 0.007
		UltraGauss - 300K	0.998 ± 0.000	0.981 ± 0.003	41.45 ± 3.10	0.009 ± 0.002	0.975 ± 0.006	38.62 ± 2.10	0.023 ± 0.005	0.965 ± 0.008	42.12 ± 3.31	0.032 ± 0.006	0.977 ± 0.006	40.73 ± 2.83	0.021 ± 0.004
		UltraGauss - 2M	2.000 ± 0.000	0.991 ± 0.002	44.40 ± 3.36	0.005 ± 0.001	0.984 ± 0.004	40.24 ± 2.17	0.034 ± 0.003	0.980 ± 0.004	45.88 ± 3.91	0.018 ± 0.003	0.985 ± 0.003	45.51 ± 3.15	0.012 ± 0.002
	ImplicitVol	ImplicitVol	0.568 ± 0.008	0.745 ± 0.042	26.13 ± 2.60	0.171 ± 0.020	0.632 ± 0.054	22.97 ± 1.77	0.533 ± 0.030	0.632 ± 0.052	26.36 ± 2.77	0.535 ± 0.039	0.670 ± 0.050	26.36 ± 2.77	0.292 ± 0.029
		RapidVol	0.694 ± 0.107	0.688 ± 0.102	26.13 ± 2.83	0.070 ± 0.008	0.744 ± 0.055	22.97 ± 1.79	0.537 ± 0.031	0.639 ± 0.055	26.36 ± 2.77	0.535 ± 0.039	0.670 ± 0.050	26.36 ± 2.77	0.292 ± 0.029
		UltraNerf	0.594 ± 0.001	0.594 ± 0.001	38.59 ± 3.16	0.019 ± 0.003	0.961 ± 0.009	36.22 ± 1.81	0.051 ± 0.010	0.949 ± 0.010	40.99 ± 4.00	0.061 ± 0.010	0.961 ± 0.008	38.60 ± 2.99	0.043 ± 0.007
80 Linearly sliced Axial Slices with ±5° Jitter	Convergence	UltraGauss - 100K	0.994 ± 0.000	0.982 ± 0.004	38.59 ± 3.16	0.012 ± 0.002	0.971 ± 0.007	37.71 ± 2.16	0.032 ± 0.006	0.957 ± 0.009	42.23 ± 3.74	0.042 ± 0.007	0.970 ± 0.007	40.27 ± 3.18	0.029 ± 0.005
		UltraGauss - 300K	0.998 ± 0.000	0.986 ± 0.006	43.99 ± 2.75	0.007 ± 0.001	0.986 ± 0.004	39.30 ± 2.52	0.067 ± 0.008	0.987 ± 0.006	44.70 ± 3.06	0.032 ± 0.006	0.988 ± 0.005	45.51 ± 3.15	0.012 ± 0.002
		UltraGauss - 2M	0.990 ± 0.001	0.985 ± 0.008	38.39 ± 2.52	0.067 ± 0.007	0.961 ± 0.008	36.30 ± 2.06	0.076 ± 0.008	0.974 ± 0.006	37.32 ± 2.80	0.067 ± 0.006	0.984 ± 0.007	37.00 ± 2.12	0.069 ± 0.006
	ImplicitVol	ImplicitVol	0.577 ± 0.007	0.804 ± 0.028	30.42 ± 2.99	0.102 ± 0.019	0.842 ± 0.033	30.07 ± 1.67	0.143 ± 0.025	0.841 ± 0.033	33.72 ± 3.19	0.161 ± 0.032	0.849 ± 0.032	33.41 ± 2.68	0.135 ± 0.025
		RapidVol	0.973 ± 0.006	0.925 ± 0.016	32.28 ± 3.70	0.074 ± 0.014	0.905 ± 0.021	31.32 ± 2.53	0.127 ± 0.019	0.903 ± 0.022	36.09 ± 3.49	0.119 ± 0.022	0.911 ± 0.019	34.43 ± 2.24	0.107 ± 0.018
		UltraNerf	0.625 ± 0.051	0.477 ± 0.040	30.91 ± 2.17	0.265 ± 0.036	0.745 ± 0.038	17.90 ± 1.04	0.373 ± 0.021	0.747 ± 0.038	14.63 ± 0.85	0.584 ± 0.043	0.444 ± 0.053	17.61 ± 1.35	0.407 ± 0.022
80 Linearly sliced Axial Slices with ±5° Jitter	Convergence	UltraGauss - 100K	0.997 ± 0.001	0.929 ± 0.007	31.82 ± 0.94	0.078 ± 0.006	0.917 ± 0.010	29.55 ± 0.46	0.088 ± 0.008	0.918 ± 0.010	30.98 ± 0.54	0.092 ± 0.009	0.921 ± 0.009	30.45 ± 0.65	0.086 ± 0.008
		UltraGauss - 300K	0.999 ± 0.000	0.934 ± 0.006	32.34 ± 1.26	0.064 ± 0.008	0.927 ± 0.008	29.13 ± 0.95	0.059 ± 0.007	0.929 ± 0.008	29.55 ± 0.93	0.065 ± 0.006	0.931 ± 0.009	30.34 ± 1.05	0.063 ± 0.007
		UltraGauss - 2M	2.000 ± 0.000	0.942 ± 0.004	33.21 ± 1.22	0.064 ± 0.007	0.937 ± 0.004	29.12 ± 0.78	0.054 ± 0.006	0.941 ± 0.004	29.79 ± 0.91	0.056 ± 0.005	0.940 ± 0.004	30.71 ± 0.97	0.058 ± 0.006

</

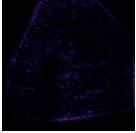
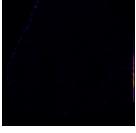
	Reconstruction Time (mins)	Rendered Image	Absolute Difference Image
864			
865			
866			
867			
868			
869			
870	Ground Truth		N/A
871			
872			
873			
874			
875	$t = 0.5$		
876			
877			
878			
879			
880	$t = 1$		
881			
882			
883			
884			
885	$t = 2$		
886			
887			
888			
889	$t = 3$		
890			
891			
892			
893			
894	$t = 4$		
895			
896			
897			
898	$t = 5$		
899			
900			
901			
902			
903	$t = 10$		
904			
905			
906			
907			
908	$t = 15$		
909			
910			
911			
912	$t = 20$		
913			
914			
915			

Figure B2: Progression of UltraGauss as a 3D reconstruction is formed from 80 Axial Slices (sampled from a 20 GW fetal scan). The image requested to be rendered image is an orthogonal, mid-coronal cross-section. Clinicians were then asked to chose whether each one of these was a real or "AI Generated" scan. The results to this are in Fig. 5

C EXPERIMENT 5.3 (CLINICIANS' SURVEY)

C.1 THE SURVEY

Below is the questionnaire which Clinicians answered:

PRELIMINARY

1. What is your current position?
2. How many years of medical experience do you have?
3. On a scale of 1 to 5, how would you rate your confidence in assessing **Fetal Brains in Ultrasound**? (5 is very confident, 1 is not at all)

PART I

Participants answered a random selection of 10 pair-wise comparisons. An example of one is shown below:

Please choose the most realistic ultrasound image out of the two.



- Option 1



- Option 2



o No Preference

PART II

Participants then answered 10 of the following questions. The order of these 10 questions and the order of the two options within each question was randomly shuffled for each participant.

The ultrasound scans below show mid-sagittal views of the same fetal brain. Some are real scans acquired using standard ultrasound equipment, while others are AI-generated reconstructions. For each pair of images, please **identify whether the scan is real or AI-generated**.

Real ultrasound scan, or AI reconstruction?



- Real
- AI-generated

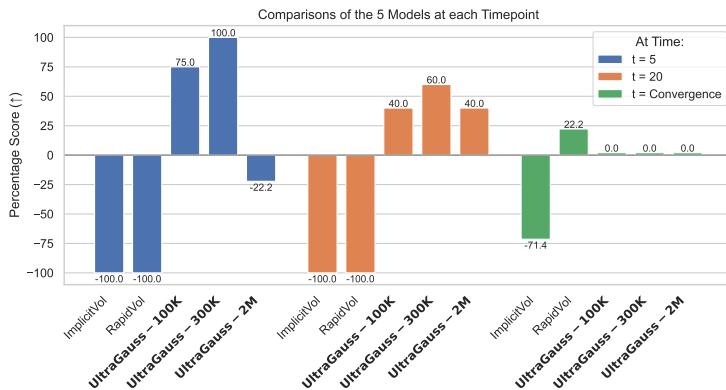
972 C.2 CLINICIANS' SURVEY ADDITIONAL RESULTS
973

Figure C1: Survey results: Comparison of all 5 models at each time point. Actively choosing a model gave it 1 point, not choosing it -1, and choosing “no preference” gave it 0 points. Each model’s score was then divided by the total number of points it could have gained, to give a percentage. Thus +100% means that all participants actively selected and preferred that model. -100% means that all participants actively disliked this model (and selected the other one). 0% means that the model is neither liked or disliked, as the participant chose the “no preference” option.

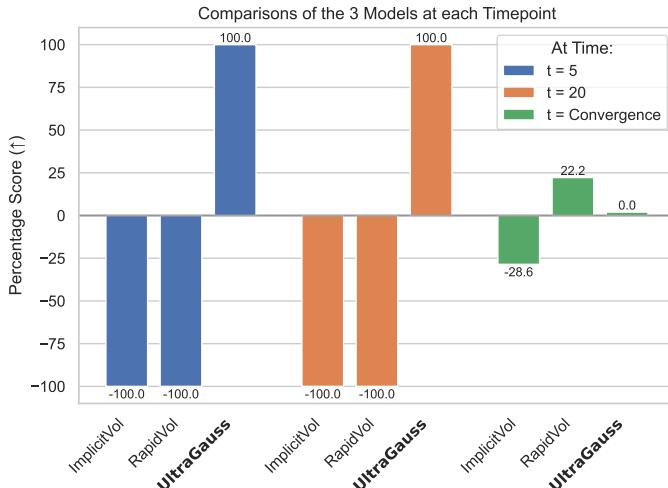
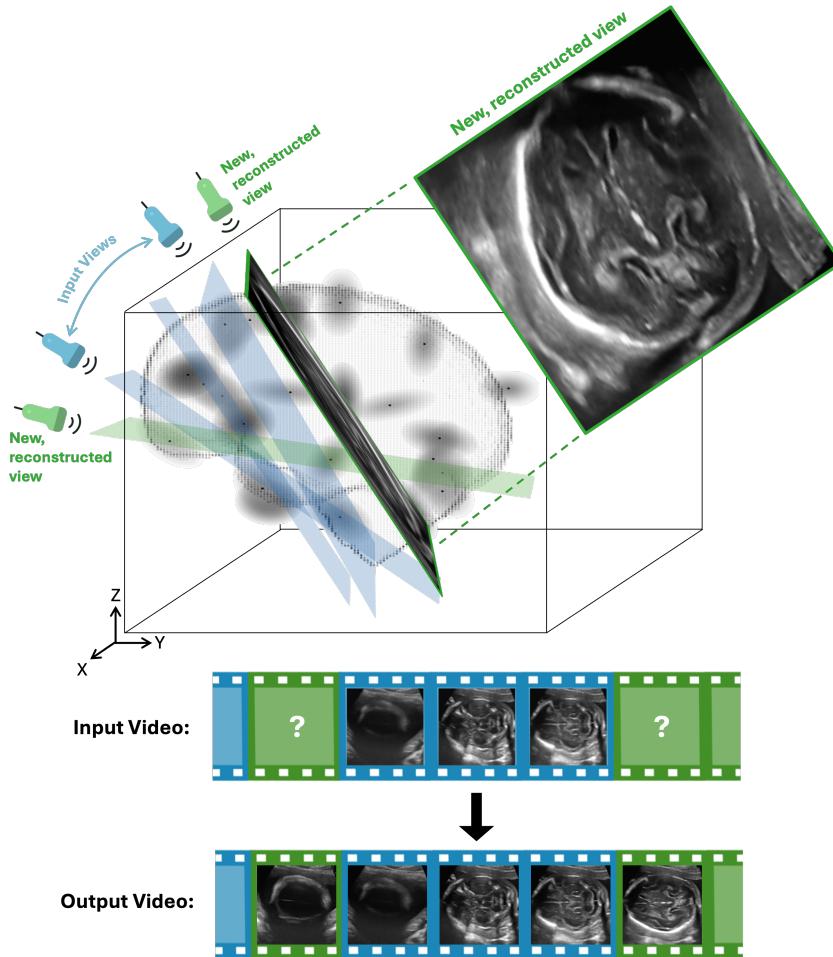


Figure C2: Survey results: Comparison of ImplicitVol, RapidVol and the better of the three Ultra-Gauss Models, at each time point. The same scoring system applies as Fig. C1

1026 **D EXPERIMENT 5.4 (US VIDEOS)**
10271028 **D.1 FURTHER EXPERIMENTAL SETUP DETAILS**

1029 In this experiment we use ultrasound cinesweep videos acquired using standard sensorless probes. As
1030 such, the poses of their frames have to first be estimated using a suitable pose estimation that works
1031 with sensorless probes (e.g. (Ramesh et al., 2024; Di Vece et al., 2024; Yeung et al., 2022)). 20% of
1032 the frames are then randomly removed and withheld for testing. The remaining frames (blue frames
1033 in Input Video on Fig. D1) and their corresponding poses (blue probes/planes) are then inputted
1034 into UltraGauss to form a 3D reconstruction of the individual fetal brain within 5 minutes. Images
1035 are then sampled from the reconstructed volume at the poses of the withheld video frames (green
1036 probes/planes), and are shown in green in the Output Video tape.

1037 A similar method also applies to forming 3D reconstructions from Dataset A. Rather than the blue
1038 Input Views being video frames and their predicted poses, they are instead slices sampled from the
1039 volumetric scans of Dataset A and their corresponding ground truth poses.



1071 Figure D1: Input images of the fetal brain are acquired (blue Input Video frames) and their poses
1072 estimated (blue probes/planes). UltraGauss then uses these to form a total 3D reconstruction within 5
1073 minutes. A few of the 3D Gaussians are shown as grey ellipses. Cross-sectional views at previously
1074 unseen poses can then be sampled from the reconstructed 3D volume to form a complete cinesweep
1075 or be viewed individually. These can be seen in the green Output Video cinefilm frames.
1076
1077
1078
1079

1080

1081 D.2 ADDITIONAL RESULTS FROM VIDEOS
1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

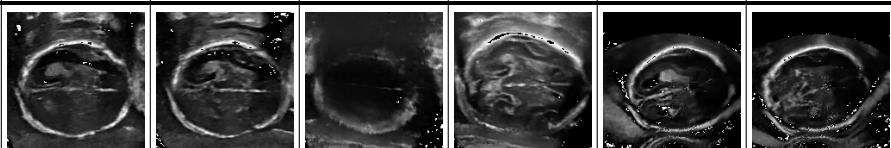
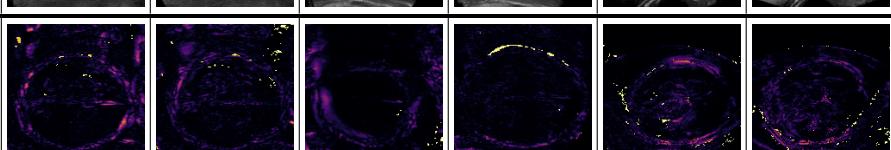
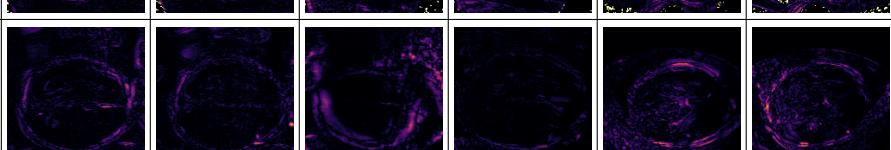
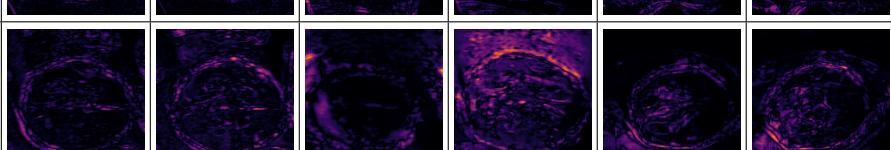
1105

1106

1107

1108

1109

Ground Truth**ImplicitVol****RapidVol****UltraNerf****UltraGauss****ImplicitVol Diff.****RapidVol Diff.****UltraNerf Diff.****UltraGauss Diff.**

Continued below...

1129

1130

1131

1132

1133

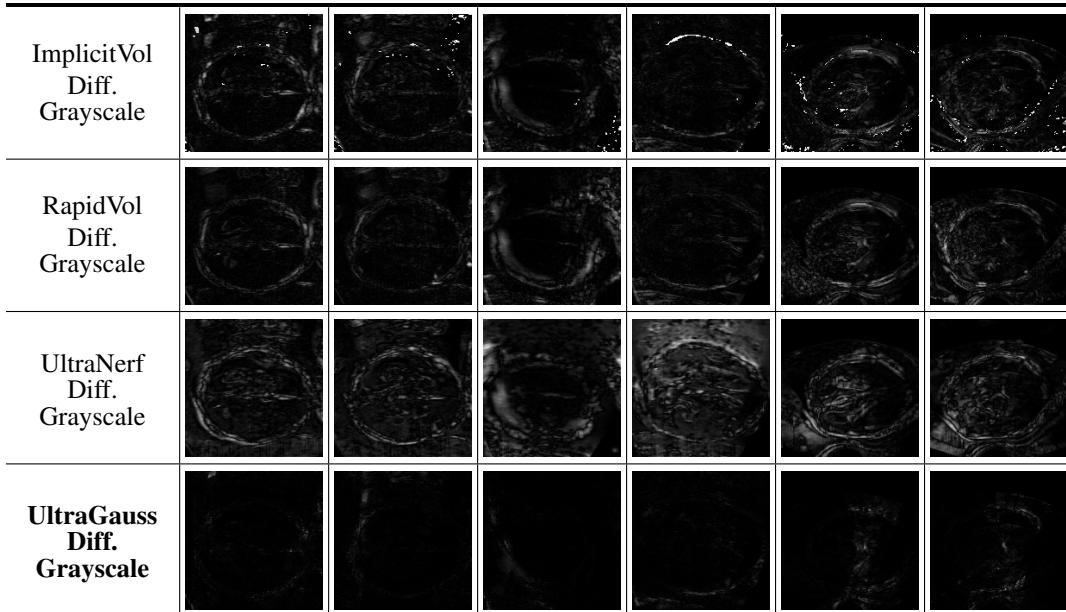
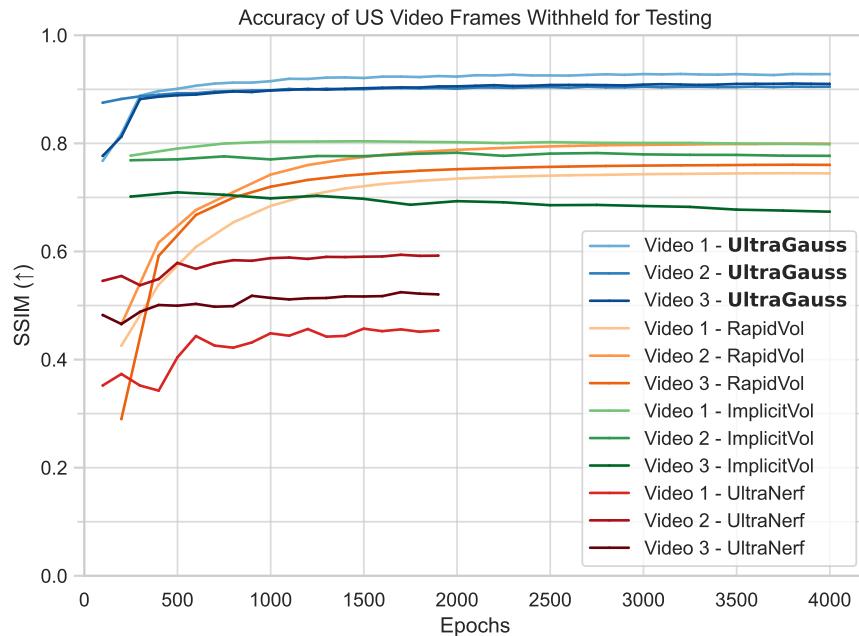


Figure D2: For each video (V_n), two ground truth images (I_{1j}) from the test set are shown, alongside the corresponding image generated by ImplicitVol, RapidVol, UltraNerf and UltraGauss - 300K (ours). The absolute difference (diff.) between the predicted and ground truth image is also shown, in both the `inferno` colored heatmap and grayscale. A pure black diff. image means perfect similarity, and so the darker the diff. image the better.



1188 **E PHYSICAL MEANING OF L_{ij} TERMS IN EQ. 11**
 1189

1190 In order to gain an intuitive understanding of the physical meaning of the components of L in Eq. 11,
 1191 the 3×3 covariance matrix Σ needs to be algebraically calculated, as only then can the standard
 1192 deviations and correlation coefficients in the x, y, z dimensions be extracted and provide us with
 1193 geometric insight. Ignoring the small value β from Eq. 11 for simplicity:

$$1194 \quad L = \begin{bmatrix} L_{11}^2 & 0 & 0 \\ 1195 \quad L_{12} & L_{22}^2 & 0 \\ 1196 \quad L_{13} & L_{23} & L_{33}^2 \end{bmatrix} \quad (E1)$$

$$1197 \quad \Sigma^{-1} = LL^T = \begin{bmatrix} L_{11}^4 & L_{11}^2 L_{12} & L_{11}^2 L_{13} \\ 1198 \quad L_{11}^2 L_{12} & L_{12}^2 + L_{22}^4 & L_{23} L_{22}^2 + L_{12} L_{13} \\ 1199 \quad L_{11}^2 L_{13} & L_{23} L_{22}^2 + L_{12} L_{13} & L_{13}^2 + L_{23}^2 + L_{33}^4 \end{bmatrix} \quad (E2)$$

$$1200 \quad \Sigma = [\Sigma^{-1}]^{-1} = \begin{bmatrix} \sigma_{xx}^2 & \rho_{xy} \sigma_{xx} \sigma_{yy} & \rho_{xz} \sigma_{xx} \sigma_{zz} \\ 1201 \quad \rho_{xy} \sigma_{xx} \sigma_{yy} & \sigma_{yy}^2 & \rho_{yz} \sigma_{yy} \sigma_{zz} \\ 1202 \quad \rho_{xz} \sigma_{xx} \sigma_{zz} & \rho_{yz} \sigma_{yy} \sigma_{zz} & \sigma_{zz}^2 \end{bmatrix} \quad (E3)$$

1203 By inverting Σ^{-1} of Eq. E2 and comparing to the standard format of Σ as in Eq. E3, we are then able
 1204 to relate how a gaussian's size in the x, y, z dimensions (governed by the three standard deviations),
 1205 and shape/rotation (governed by the three correlation coefficients) are affected by the six L_{ij} terms
 1206 of L .

1207 Unfortunately, due to the nature of our efficient decomposition, and which is optimised to primarily
 1208 learn Σ^{-1} rather than Σ , as can be seen below the six L_{ij} terms are all heavily coupled and it is hard
 1209 to get a physical intuition of what each term is responsible for. Nevertheless, the sensitivity graphs
 1210 centered around 4 (the initial values of L_{ij}) do provide some insight (see Fig. E1).

1211 It is also evident from the equations below that all six of a gaussian's defining parameters are
 1212 solely a function of L , and so L can be thought of as the 6 Degree of Freedom matrix directly
 1213 responsible for defining a 3D Gaussian's 6 DoF shape and size. Whilst the *individual* terms of L
 1214 are not as interpretable as the original gaussian splatting's parametrization, which used parameters
 1215 $q_r, q_i, q_j, q_k, s_x, s_y, s_z$ to rotate and scale the gaussian, L as a whole is still responsible for the shape
 1216 and size of a gaussian, and nothing else. We believe this slight loss in interpretability is overwhelm-
 1217 ingly outweighed by the efficiency gains demonstrated (see Appendix F for more details).

1218 For choosing initialisation values, it is also helpful to note that generally, increasing all six L_{ij} values
 1219 causes L and subsequently Σ^{-1} to be larger. The covariance Σ , equal to the inverse of Σ^{-1} , will
 1220 thus be smaller, and so the 3D gaussians will start off smaller.

$$1221 \quad \sigma_{xx} = \frac{\sqrt{L_{12}^2 L_{23}^2 + L_{12}^2 L_{33}^4 - 2L_{12} L_{22}^2 L_{13} L_{23} + L_{22}^4 L_{13}^2 + L_{22}^4 L_{33}^4}}{L_{11}^2 L_{22}^2 L_{33}^2}$$

$$1222 \quad \sigma_{yy} = \frac{\sqrt{L_{23}^2 + L_{33}^4}}{L_{22}^2 L_{33}^2}$$

$$1223 \quad \sigma_{zz} = \frac{1}{L_{33}^2}$$

$$1224 \quad \rho_{xy} = \rho_{yx} = -\frac{L_{12}(L_{23}^2 + L_{33}^4) - L_{22}^2 L_{13} L_{23}}{\sqrt{L_{23}^2 + L_{33}^4} \sqrt{L_{12}^2 L_{23}^2 + L_{12}^2 L_{33}^4 - 2L_{12} L_{22}^2 L_{13} L_{23} + L_{22}^4 L_{13}^2 + L_{22}^4 L_{33}^4}}$$

$$1225 \quad \rho_{xz} = \rho_{zx} = -\frac{-L_{13} L_{22}^2 + L_{12} L_{23}}{\sqrt{L_{12}^2 L_{23}^2 + L_{12}^2 L_{33}^4 - 2L_{12} L_{22}^2 L_{13} L_{23} + L_{22}^4 L_{13}^2 + L_{22}^4 L_{33}^4}}$$

$$1226 \quad \rho_{yz} = \rho_{zy} = \frac{L_{23}}{\sqrt{L_{23}^2 + L_{33}^4}}$$

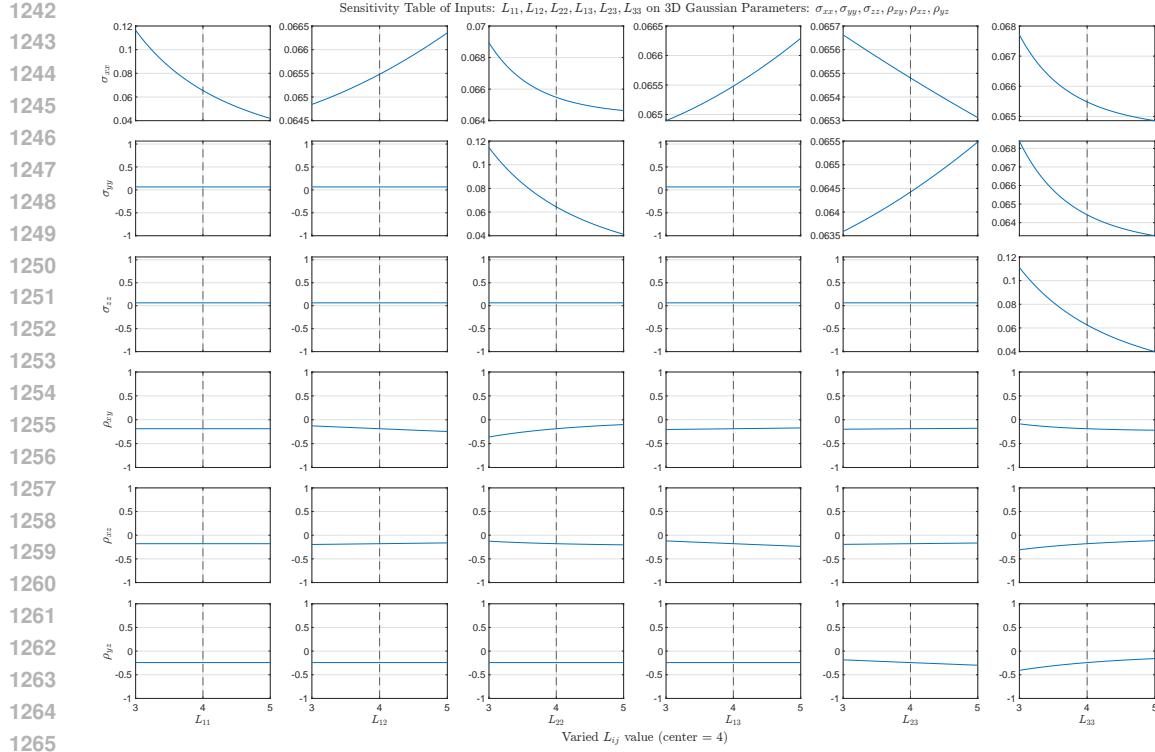


Figure E1: Sensitivity table showing the effect each L_{ij} term in Eq. 11 has on a Gaussian's 6 defining parameters as it is perturbed and all other L_{ij} values are held constant. Perturbations are shown here centered about 4 as that is the initialisation value of all L_{ij} 's, however during the optimization process each L_{ij} term is free to take any real value.

1296 **F DETAILED ANALYSIS OF OUR COVARIANCE PARAMETRIZATION**
 1297

1298 **F.1 SPEED IMPROVEMENT**

1299 The number of FLOPs (Floating Point Operations) to compute $(\Sigma^{3D})^{-1}$, which is required **every** it-
 1300 eration, and Σ^{3D} , which is only required every $N_{\text{densification}}$ iterations, can be calculated theoretically
 1301 (Hunger, 2007) as follows:

1303 **EXISTING PARAMETRIZATIONS APPLIED TO 3D**

1304 1. To form the Rotation matrix R from the 4 quaternions $\{q_r, q_i, q_j, q_k\} \in \mathbb{R}^N$ (where N is the
 1305 Number of Gaussians), we apply the standard quaternion to rotation matrix conversion and see
 1306 that it requires $30N$ FLOPs. Forming the diagonal Scaling matrix S from the 3 scaling values
 1307 $\{s_x, s_y, s_z\} \in \mathbb{R}^N$ requires 0 FLOPs.

1308 2. Next, to form Σ^{3D} from $\Sigma^{3D} = R^T S^T S R$, we are first required to form $A = S R$. Since S is
 1309 diagonal, this only requires $3 \times 3 = 9$ FLOPs. We now have $\Sigma^{3D} = A^T A$ and can make use of the
 1310 symmetric nature of the product's result, so only have to compute the leading diagonal terms and
 1311 one of the off-diagonal triangles. This means for a 3×3 matrix, 30 FLOPs are required to compute
 1312 this. Thus in total $9 + 30 = 39N$ FLOPs for this step.

1313 3. Finally, to compute $(\Sigma^{3D})^{-1}$, we can exploit the properties of the Covariance matrix, namely
 1314 that it is symmetric and positive definite, and use Cholesky decomposition to invert it in 39 FLOPs.
 1315 This means that overall, Σ^{3D} requires $30N + 0N + 9N + 30N = 69N$ FLOPs, and to form
 1316 $(\Sigma^{3D})^{-1}$, $69N + 39N = 108N$ FLOPs.

1317 **OUR PARAMETRIZATION**

1318 1. To form

$$L = \begin{bmatrix} L_{11}^2 + \beta & 0 & 0 \\ L_{12} & L_{22}^2 + \beta & 0 \\ L_{13} & L_{23} & L_{33}^2 + \beta \end{bmatrix} \quad (11)$$

1319 from our 6 covariance parameters $\{L_{11}, L_{12}, L_{13}, L_{22}, L_{23}, L_{33}\} \in \mathbb{R}^N$ (where N is the Number
 1320 of Gaussians), it requires $3N$ multiplications and $3N$ additions, so $6N$ FLOPs in total.

1321 2. Next, to form $(\Sigma^{3D})^{-1}$ from $(\Sigma^{3D})^{-1} = LL^T$, as L is a $M \times M$ Lower Triangular Matrix with
 1322 $M = 3$, we only require $\frac{1}{3}M^3 + \frac{1}{2}M^2 + \frac{1}{6}M = \frac{1}{3}3^3 + \frac{1}{2}3^2 + \frac{1}{6}3 = 14$ FLOPs. So $14N$ FLOPs in
 1323 total. Therefore every iteration, we require $6N + 14N = 20N$ FLOPs to compute $(\Sigma^{3D})^{-1}$ from
 1324 our 6 parameters and use it in Eq. 8 .

1325 3. Finally, to form Σ^{3D} from $\Sigma^{3D} = (L^{-1})^T L^{-1}$, we first need to invert L . However since
 1326 L is Lower Diagonal, this can be done very efficiently using forward substitution and only takes
 1327 $\frac{1}{3}M^3 + \frac{2}{3}M = \frac{1}{3}3^3 + \frac{2}{3}3 = 11$ FLOPs. The product of two lower triangular matrices (this time L^{-1}
 1328 rather than L) is, as before, 14 FLOPs. As such, we need $11N + 14N = 25N$ additional FLOPs to
 1329 form Σ^{3D} .

1330 Importantly, however, is that during densification we only need σ_{xx} , σ_{yy} & σ_{zz} , and not the full
 1331 covariance matrix. This means we can simply take the squared L2-norm of $(L^{-1})^T$ along the
 1332 column dimension, only requiring $9N$ FLOPs. When we then need to sample from the gaussian
 1333 distributions, we can also make use of Eq. 12, which again does not require Σ^{3D} but $(L^{-1})^T$ -
 1334 and that has already been calculated. Therefore, every $N_{\text{densification}}$ iterations we only require $20N$
 1335 additional FLOPs.

1336 The above can be nicely summarised in Tab.E1

1337 As can be seen, our formulation is much more efficient than had we used the 3D version of 3DGS'
 1338 formulation. Not only do we require $5.4 \times$ less FLOPs to compute Σ^{-1} each iteration, even if we
 1339 did need the full covariance matrix Σ , our formulation still computes this using $1.73 \times$ less FLOPs.
 1340 In practice, we measured speed-ups of $1.4 \times$ on the forward pass and $1.2 \times$ on the backward pass
 1341 when calculating Σ^{-1} . As Σ^{-1} is required every iteration, this is a meaningful speed improvement.
 1342 When Σ is required for densification every $N_{\text{Densification}}$ iterations, our formulation computes this
 1343 $1.25 \times$ faster in the forward pass and $6.65 \times$ faster in the backward pass than had we used the 3DGS

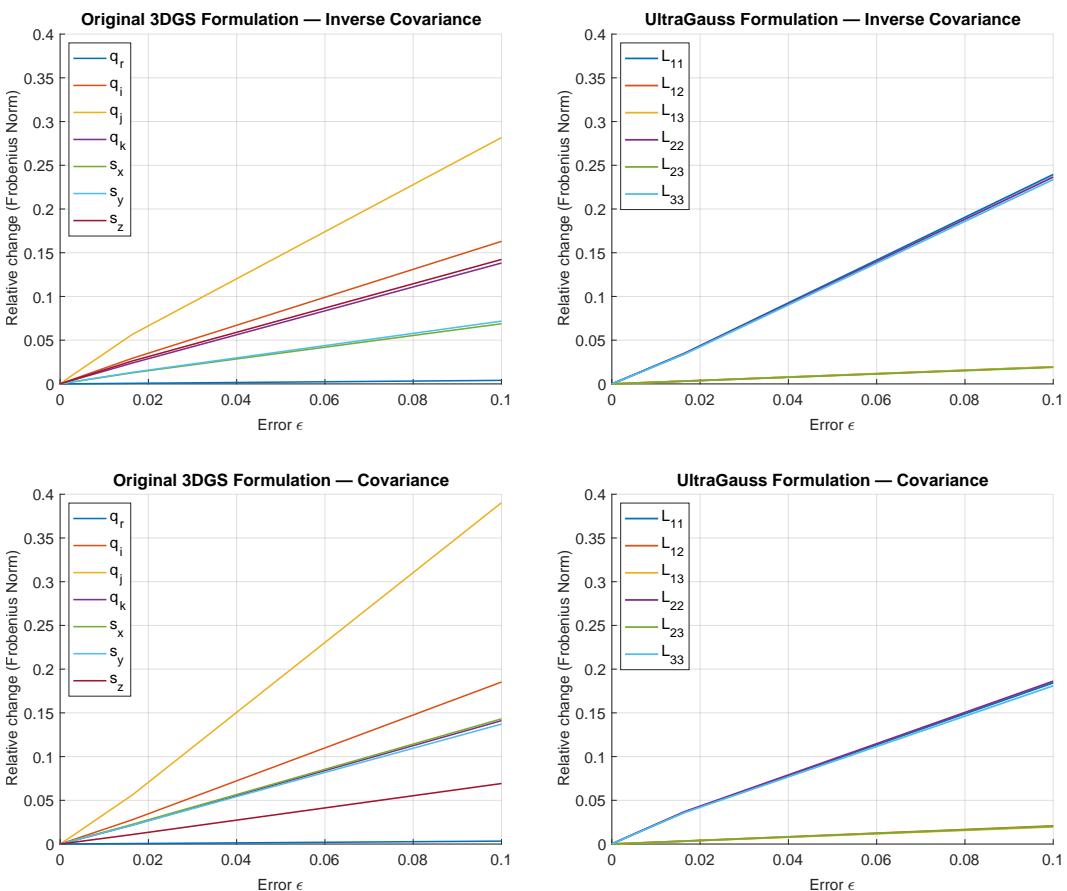
1350 formulation (although in practice we only need L^{-T} , not the full covariance matrix, so speed-up is
 1351 even quicker).
 1352

Method	FLOPs for Σ^{-1} (computed every iter)	FLOPs for Σ (computed every $N_{\text{densification}} \text{ iters}$)
Original 3DGS Formulation (e.g., Kerbl et al.; Zha et al.)	$108N$	$69N$
UltraGauss (ours)	$20N$	$+20N$ in practice* ($+25N$ if Σ is actually needed*)

1353
 1354 Table E1: Comparison of theoretical FLOPs required to compute Σ and Σ^{-1} for different parameterizations.
 1355 *For densification, only the variances are needed - not the full covariance matrix - so we
 1356 can take the squared L2 norm of L^{-T} , saving 5 FLOPs.
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372

F.2 NUMERICAL STABILITY

1373 In terms of numerical stability, we also performed a numerical sensitivity analysis on the two
 1374 formulations. Each input parameter was subjected to a small perturbation, to simulate numerical
 1375 errors/rounding, and the relative Frobenius Norm difference between the error-containing
 1376 inverse/covariance matrix and the error-free matrix is shown.
 1377
 1378



1401 Figure E1: Numerical analysis of the effect a small perturbation or error in each parameter has on the
 1402 final Covariance or Inverse Covariance Matrix, for the original parameterization and our proposed
 1403 one.
 1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413

1404 From Fig. E1 it is evident that our proposed formulation is more robust to numerical error, when
1405 computing both the inverse covariance matrix and covariance matrix from the input parameters.

1406 We also found empirically that sometimes, when trying to invert the 3×3 covariance matrix formed
1407 by the four quaternions and three scaling parameters (which is never done in classical GS, as only
1408 the *projected* 2×2 covariance matrix is inverted), if any of the eigenvalues became close to zero we
1409 would get numerical instabilities. Our formulation (by adding a small β), ensures all eigenvalues
1410 are always $\geq \beta$ so we never run into this issue.

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

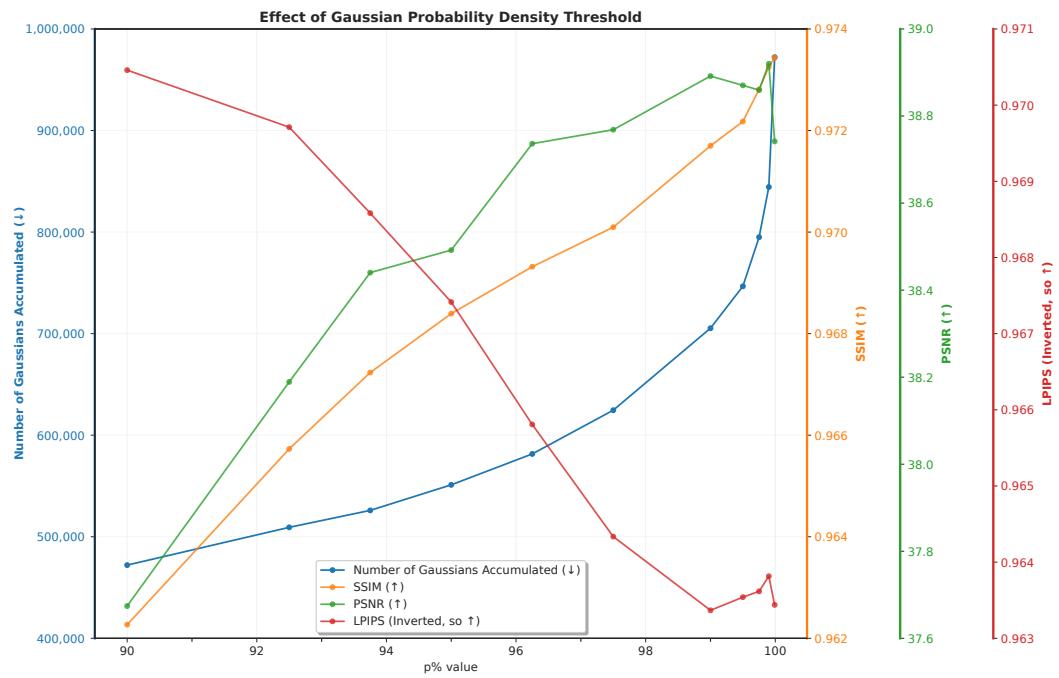
1455

1456

1457

1458 G EFFECT OF A GAUSSIAN'S PROBABILITY DENSITY p VALUE
1459

1460 We expected that as the gaussian probability density threshold value was increased, the size of a
1461 gaussian's ellipsoid of influence would increase, and so more pixels would consider it to be "close"
1462 and accumulate it. Applied to all N gaussians, this would mean more atomic additions and so
1463 increase the time taken, but lead to potentially higher accuracy as each pixel's final grayscale value
1464 is now controlled by more gaussians. To confirm this we conducted an ablation on UltraGauss-300K
1465 applied to one of the 3D scans from Dataset A, and varied the $p\%$ value (see Eq. 13). Metrics are
1466 averaged across all 3 test views. Our ablation confirmed our expectations and as such we chose
1467 $p = 95\%$ as a good compromise between the number of gaussians accumulated each iteration and
1468 reconstruction accuracy (considering all three metrics).



1491 Figure G1: Ablation study on UltraGauss in which the p value of the gaussians' bounding probability
1492 density ellipsoids are increased. Metrics are averaged across all 3 test views.

1512 **H THE EFFECT OF POSE ESTIMATION ERROR**

1513

To attempt to disentangle the error as a consequence of having inaccurate poses from the error due
1514 to the 2D-to-3D reconstruction process, we conducted the following ablation. For 80 linearly spaced
1515 axial slices sampled from the 3D volumes of Dataset A, rather than providing these images along
1516 with their ground truth poses as we did in Sec. 5.2, we now intentionally add some uniform random
1517 noise to the poses, to simulate UltraGauss being given inaccurate poses as would be the case when
1518 using the cinesweeps in our end-to-end pipeline (Sec. 5.4).

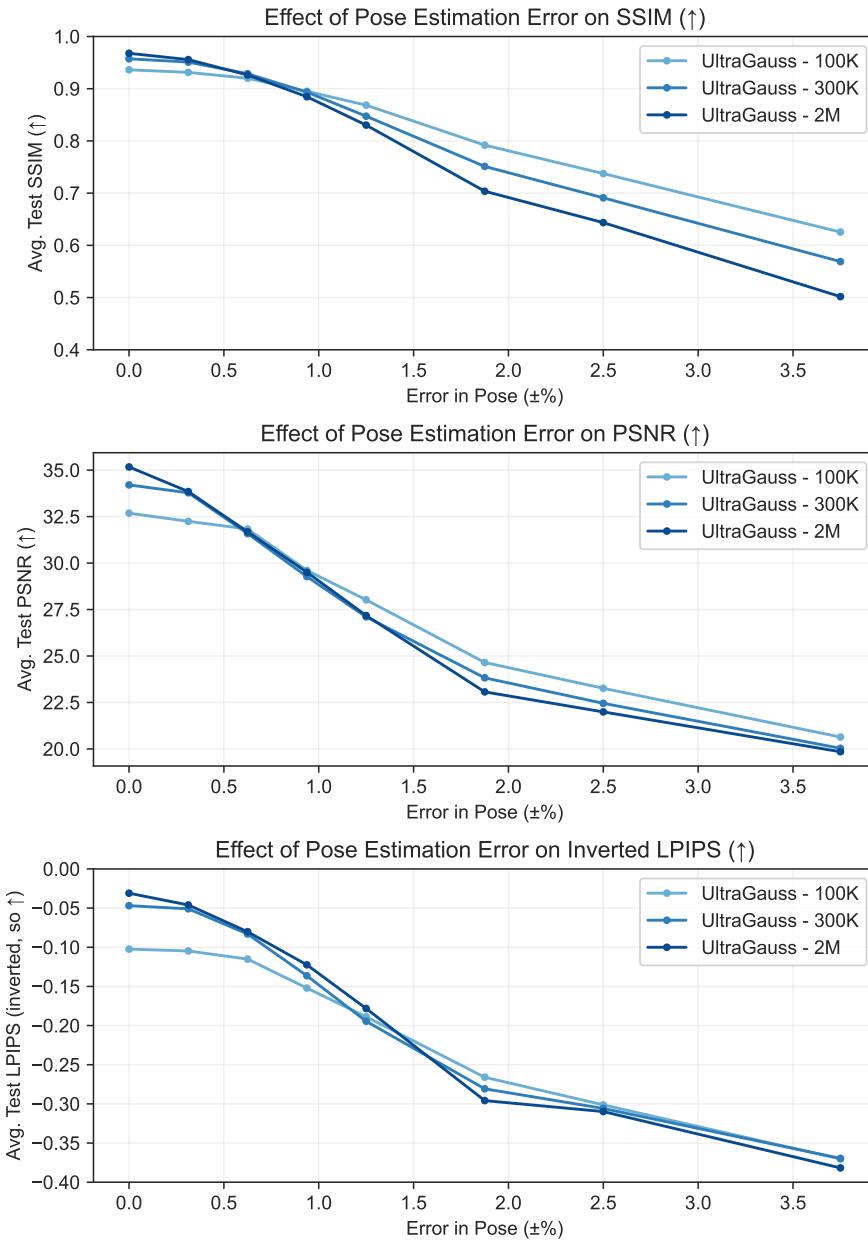

1559

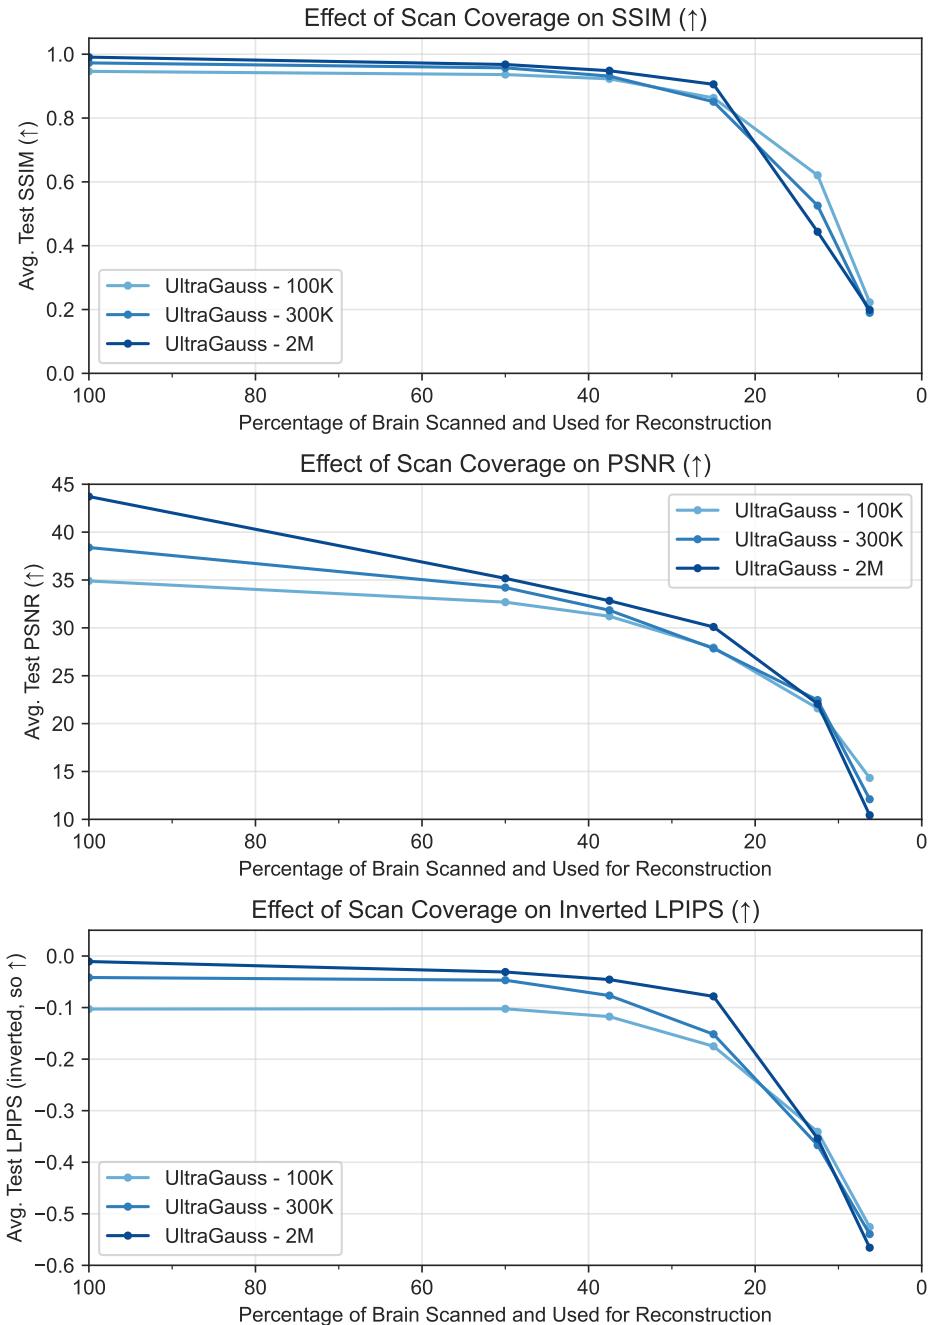
Figure H1: Ablation study on how the different variants of UltraGauss cope as error in the pose
1560 of the reconstruction input images increases. LPIPS is shown inverted here, so that like SSIM and
1561 PSNR, a higher value signifies better accuracy.

1562

From Fig. H1 it can be seen as is typical with reconstruction methods which are overfitted to a
1563 particular scene and are not pre-trained, they do require fairly accurate registration of images to
1564 a fixed co-ordinate system. For UltraGauss the error limit is about 1.5% before reconstruction
1565 accuracy degrades.

1566 I THE EFFECT OF SCAN COVERAGE
1567

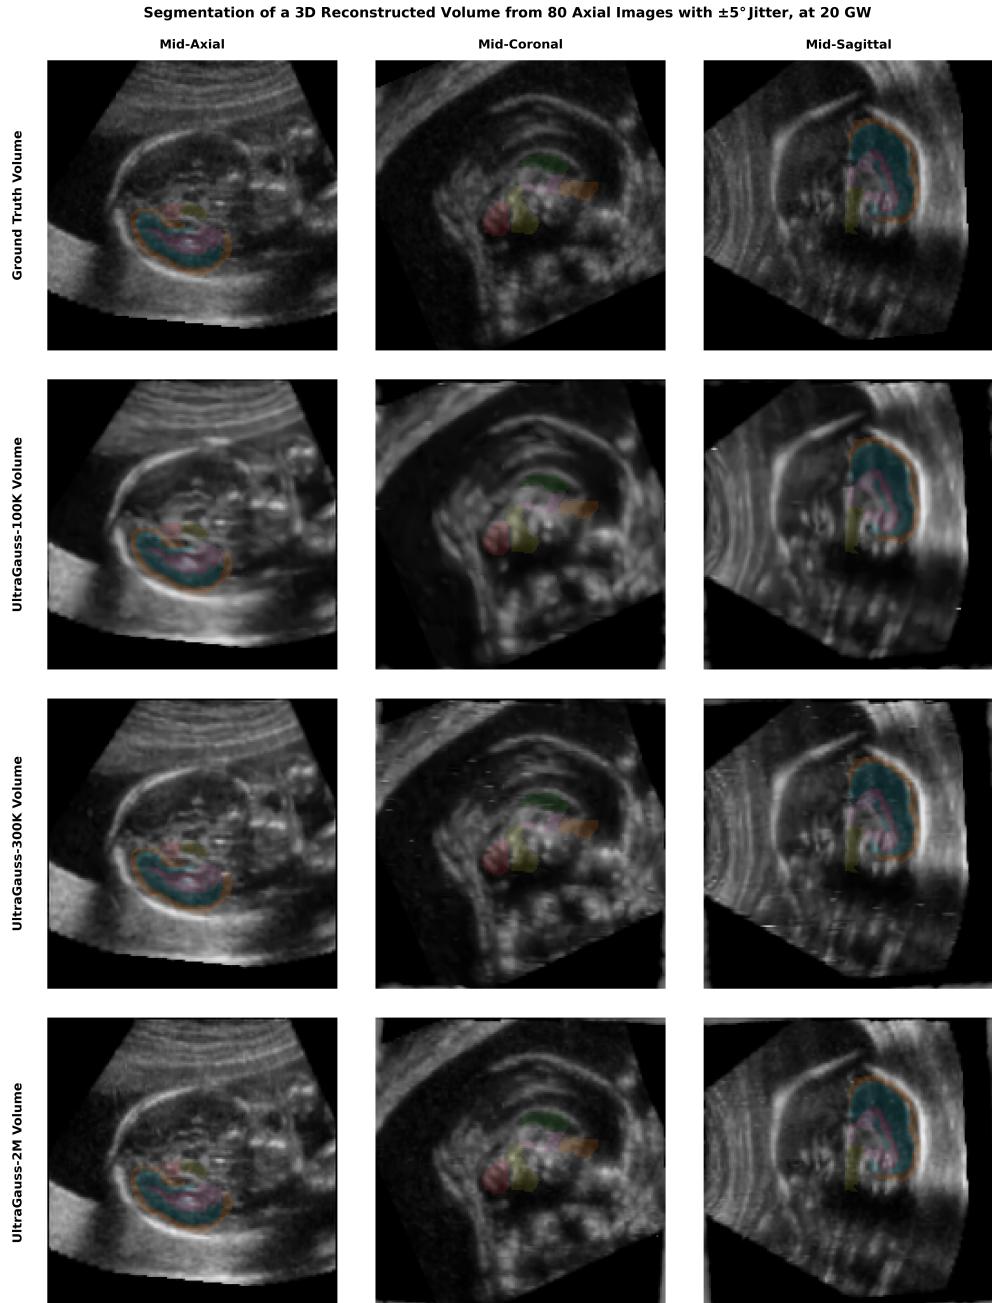
1568 Here we investigate how well UltraGauss copes with incomplete scan coverage. We train on less
1569 and less axial slices, and report the average SSIM, PSNR and Inverted LPIPS averaged across the
1570 novel axial, sagittal and coronal views generated. Generally, quality is good until around 25%
1571 coverage of the subject, beyond which it drops significantly. Since UltraGauss is not pre-trained,
1572 and like most 3D reconstruction methods over-optimises to the specific scene, this behaviour is
1573 somewhat expected. Crucially, it also ensures that UltraGauss does not hallucinate or create non-
1574 existent anatomy, something that is crucial for use in medical settings and we actively considered
1575 when tuning our heuristics.



1618 Figure I1: Ablation study on the quality of UltraGauss reconstructions as scan coverage decreases.
1619 LPIPS is shown inverted here, so that like SSIM and PSNR, a higher value signifies better accuracy.

1620 **J ANATOMICAL ACCURACY OF ULTRAGAUSS' 3D RECONSTRUCTIONS**
1621

1622 In Figs. J1 and J2 we overlay segmentation masks (Namburete et al., 2023; Hesse et al., 2022) over
 1623 the UltraGauss reconstructions in the hardest of our 3 scenarios - using 80 Axial slices with $\pm 5^\circ$
 1624 Jitter (to mimic the sonographer's natural hand motion). We also segment the original 3D volume as
 1625 a benchmark. We show mid-axial, mid-coronal and mid-sagittal views, at both 20 and 26 Gestational
 1626 Weeks (no segmentation masks were available for 14 GW). All the UltraGauss reconstructions show
 1627 good anatomical accuracy.



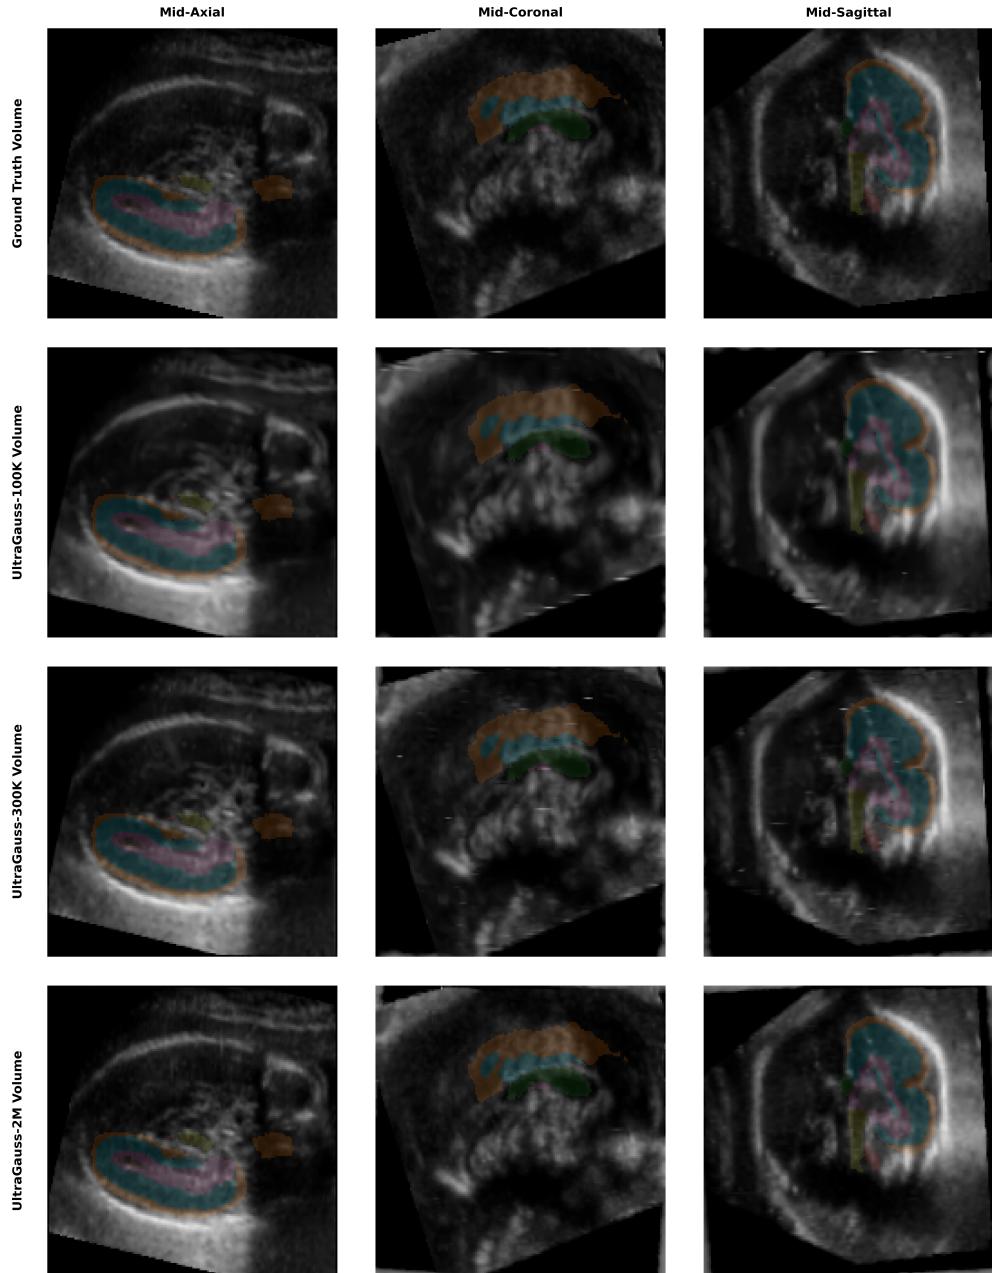
1669 **Figure J1:** Segmentation mask applied to the Original and UltraGauss-reconstructed volumes of a
 1670 20 GW fetus. The input to the reconstructions was 80 Axial slices with $\pm 5^\circ$ of rotational jitter in
 1671 the plane.
 1672

1673

1674
1675
1676
1677
1678

1679 **Segmentation of a 3D Reconstructed Volume from 80 Axial Images with $\pm 5^\circ$ Jitter, at 26 GW**

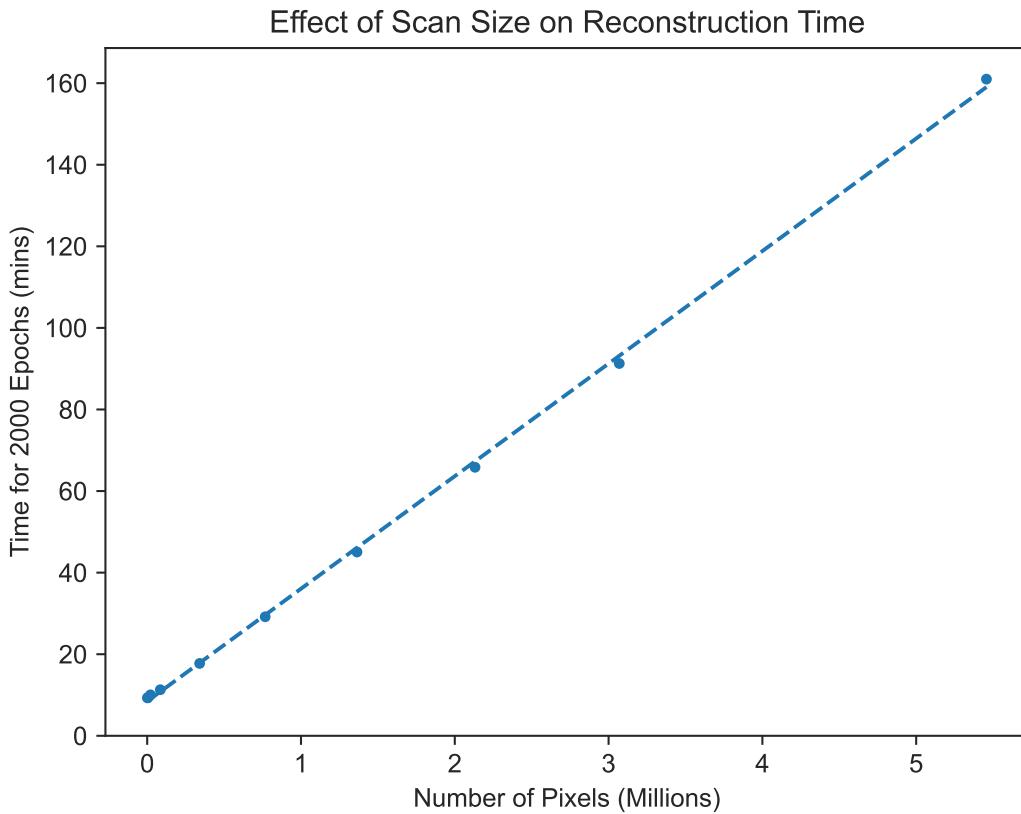
1680



1721 Figure J2: Segmentation mask applied to the Original and UltraGauss-reconstructed volumes of a
1722 26 GW fetus. The input to the reconstructions was 80 Axial slices with $\pm 5^\circ$ of rotational jitter in
1723 the plane.

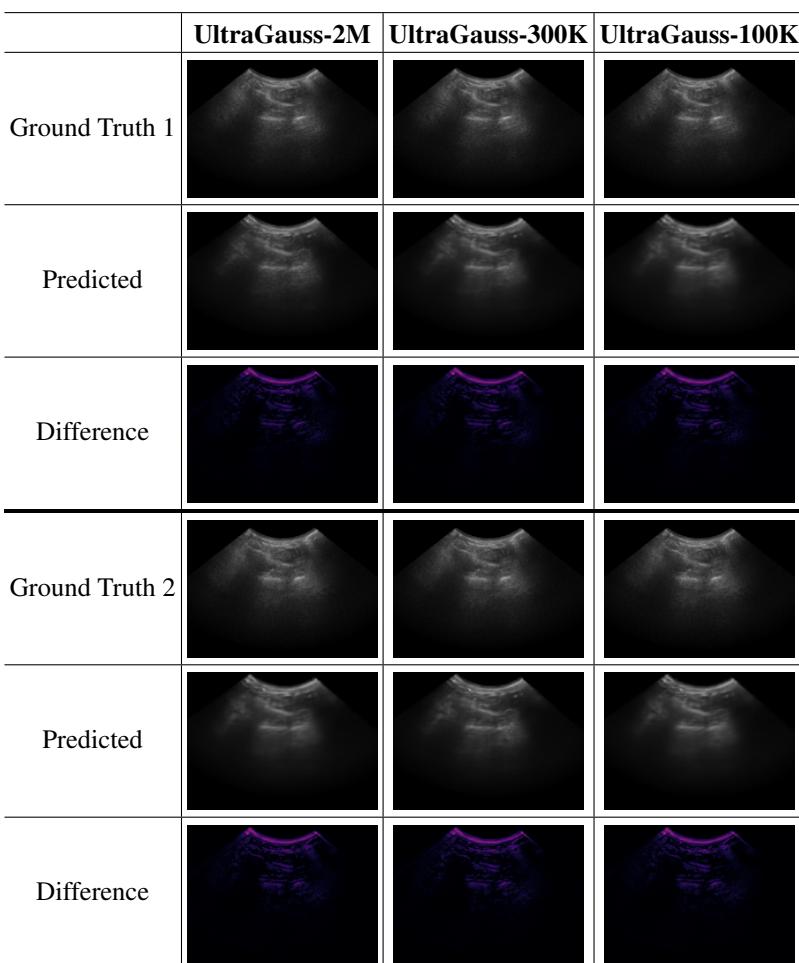
1724
1725
1726
1727

1728 **K THE EFFECT OF SCAN SIZE ON RECONSTRUCTION TIME**
1729

1730 To confirm our expectation that reconstruction time scales linearly with the number of pixels (i.e. the
1731 area of the 2D scans), we conducted an ablation using UltraGauss-300K. Fig. K1 evidently confirms
1732 this.

1759 Figure K1: Ablation study on how scans size (and thus the number of pixels), affects reconstruction
1760 time of UltraGauss (shown here for UltraGauss with 300,000 initial gaussians).

1782 **L PRELIMINARY RESULTS ON ULTRASOUND CINESWEEPS OF FOREARMS**
1783

1784 To further demonstrate that UltraGauss is not organ specific and can be applied to general 2D
 1785 ultrasound cinesweeps, we also validate on the TUS-REC Dataset (Li et al., 2025). This dataset of
 1786 human forearm ultrasound cinesweeps was acquired using an Ultrasonix machine equipped with a
 1787 curvilinear probe (4DC7-3/40). A NDI Polaris Vicra optical tracking system was used to provide
 1788 ground truth poses, and standard ultrasound post-processing speckle reduction was **not** applied. S
 1789 shape trajectory scans and poses were used as input into UltraGauss, and then novel views were
 1790 sampled from the 3D reconstruction at the poses of the straight line trajectory scans. Preliminary
 1791 results can be found in Fig. L1. It is evident that UltraGauss still forms good 3D reconstructions,
 1792 and that the main structural content and anatomy is modelled accurately. Most of the pixel-wise
 1793 difference seen in Fig. L1 is in areas with a lot of fine speckle and which would require many,
 1794 extremely small gaussians. In practice there would not be this much speckle in the ultrasound
 1795 cinesweeps, as post-processing speckle reduction would be applied, and so the reconstructions
 1796 would be even better. It is only in this dataset the authors intentionally chose not to do this, in order
 1797 to help the participants in their pose-prediction challenge.



1828 Figure L1: UltraGauss - {2M, 300K, 100K} was applied to form a 3D reconstruction from an input
 1829 cinesweep of a human participant’s forearm. Two novel views were then sampled and are shown
 1830 here alongside their ground truth. The pixel-wise absolute difference between the predicted and
 1831 ground truth images are shown in the perceptually uniform `inferno` colored heatmap. A pure
 1832 black difference image means perfect similarity, and so the darker the difference image the better.
 1833
 1834
 1835