3D Cerebrovascular Shape Completion from Biplane Angiography and CTA Prior

Janik Jehkul[1,2], Sarah Frisken[1], Vivek Gopalakrishnan[3], Daniel Rückert[2,3], and Nazim Haouchine[1]

[1] Harvard Medical School, Brigham and Women's Hospital, Boston, MA, USA
[2] Chair for AI in Healthcare and Medicine, Technical University of Munich and TUM University Hospital, Munich, Germany
[3] Department of Computing, Imperial College London, UK
[4] Massachusetts Institute of Technology, Cambridge, MA, USA
Corresponding author: janik.jehkul@tum.de

Abstract. We propose a novel approach to bridge the resolution and dimensionality gap between CTA, which provides 3D vascular geometry but often misses small vessels, and biplanar 2D DSA, which offers higher spatial resolution but lacks 3D structural information, by formulating the problem as a 3D shape completion task. Our method represents vasculature using a set of 3D Gaussians initialized from CTA-derived vessel geometry and augmented with additional spatial and opacity primitives seeded from two DSA projections. These Gaussians jointly encode geometry and attenuation and are optimized to fit the observed DSA images while remaining consistent with the original CTA anatomy. Experiments on synthetic and clinical cerebrovascular data demonstrate improved 3D reconstruction of small vessel branches at submillimetric resolution, validating the use of DSA to complement missing CTA vessel anatomy. Code and data are available here: https://github.com/janik-j/cta-dsa-fusion

Keywords: Angiography · 3D Reconstruction · Gaussian splatting.

1 Introduction

Computed tomography angiography (CTA) plays an important role in the planning and treatment of neurovascular diseases by providing rich 3D information about brain vascular structure; however, its resolution, typically between 0.5 and 1 mm limits visualization of small distal vessels, obscuring thereby subtle pathologies. In contrast, digital subtraction angiography (DSA), frequently acquired as biplanar 2D projections, captures submillimetric vascular detail, with pixel sizes on the order of 0.1 to 0.2 mm [5, 2]. This higher resolution enables improved depiction of small-caliber vessels and vascular lesions, and yields superior diagnostic performance compared to CTA, particularly for small lesions (such as aneurysms below 3 mm [12]). This limitation has motivated the development of multiple reconstruction approaches aimed at building a high-resolution 3D vascular network from DSA images.

When rotational 3D DSA is available, the dominant method for reconstructing three-dimensional vasculature relies on Feldkamp–Davis–Kress (FDK) filtered backprojection [4]. This method was originally developed for dense acquisitions, typically requiring more than a hundred projections, but recent attempts have been made to adapt it to sparse-view settings reducing the number of acquisitions [19, 20, 11, 10]. The authors in [19] reported 3D-DSA reconstruction from ultra-sparse 2D projections with substantial dose reduction relative to clinical protocols [19]. TiAVox [20] extended this direction to 4D DSA using time-aware attenuation voxels. VPAL [11] introduced a vessel-probability field to decompose static background and dynamic contrast flow, improving sparse-view dynamic DSA reconstruction quality. 4DRGS [10] replaced attenuation voxel fields with Gaussian kernels following the Gaussian splatting representation [9], reporting comparable quality with substantially faster optimization. In particular, Gaussian representations have previously been shown to be well suited for X-ray rasterization in the synthesis of DSA images, as demonstrated by R^2-Gaussian [18]. In this formulation, each Gaussian jointly encodes spatial support, via its mean and covariance defining local geometry, and radiographic contribution, via an associated density term modeling X-ray attenuation. This explicit coupling of geometry and opacity aligns naturally with X-ray image formation, where measured intensities arise from the integration of attenuation along projection rays.

The clinical availability of rotational 3D DSA providing angular diversity across dozens of viewpoints is not always guaranteed. In routine biplane DSA, the system typically acquires only frontal (AP) and lateral (LAT) projections, with an additional oblique view obtained selectively, motivating a separate line of 3D reconstruction methods that operate solely on biplane images. In [6], the authors introduced a constrained two-view DSA reconstruction that back-projects distance fields of 2D centerlines and enforces agreement of DSA-derived vessel radii, bolus-arrival timing, and projected velocity, reporting strong ghosting suppression in digital phantoms with evaluation on clinical data [5]. In [3], DSA reconstruction is formulated as a variational optimization problem that relies on a pre-existing 3D map as a prior for regularization and constraint. Other approaches have explored learned priors through conditional GANs [21], denoising autoencoders on back-projected volumes [16], or combining 3D nnUNet segmentation with MAP optimization and a connectivity prior [2]. These methods improve two-view reconstruction through explicit constraints or learned priors, but do not incorporate patient-specific preoperative CTA geometry as a direct reconstruction prior or assume that DSA-visible vessel anatomy is fully represented in a 3D CTA prior, an assumption that is rarely satisfied in practice.

Contribution. We propose a novel approach to address the inherent CTA-DSA resolution–dimensionality trade-off by formulating the problem as a 3D *shape completion* task, which aims at restoring anatomy that is absent, or not visible, from a partial anatomy. In our setting, missing vessels are reconstructed to complete a CTA-based vascular structure from as few as two biplane DSA projections. Our approach parameterizes the vascular field as a set of 3D Gaussians drawn from two complementary sources that are jointly optimized, *partial*

CTA vessel mesh surface and *complete* bi-plane DSA projection images. Unlike existing methods, our approach requires no rotational acquisition or temporal data and leverages existing patient-specific 3D geometry without requiring large training datasets. We evaluate on both synthetic and clinical data and show that our method outperforms the evaluated baselines, with the ability to grow missing 3D vessel branches from CTA using DSA at submillimetric resolution. Code and data will be made available.

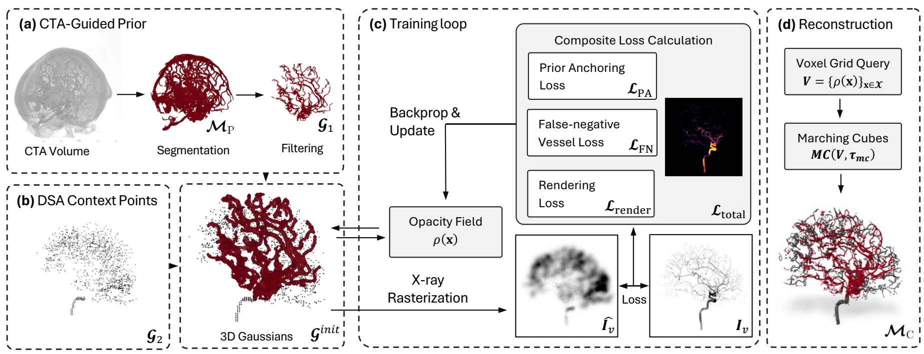

Fig. 1. Method overview: CTA priors obtained by segmentation in **(a)** and DSA back-projections context points in **(b)** are used to initialize Gaussians that are optimized under a composite loss in **(c)**, leading to the final vessel reconstruction in **(d)**.

2 Methods

2.1 Problem Formulation

We consider two DSA maximum-intensity projections I_v, with $v \in \{\mathrm{AP}, \mathrm{LAT}\}$, together with an incomplete vessel surface mesh $\mathcal{M}_{\mathrm{P}}$ extracted from a preoperative CTA volume, which captures major vessels but misses finer branches visible only in DSA. For each view v, the corresponding C-arm projection geometry P_v is obtained by co-registering $\mathcal{M}_{\mathrm{P}}$ to I_v using [7]. We seek a completed vessel mesh $\mathcal{M}_{\mathrm{C}}$ that recovers the missing shape so that:

$$\mathcal{M}_{\mathrm{C}}^{*} = \arg\min_{\mathcal{M}_{\mathrm{C}}} \mathcal{L}_{\mathrm{total}}(\mathcal{M}_{\mathrm{P}},\ P_v,\ I_v)\,, \tag{1}$$

where $\mathcal{L}_{\mathrm{total}}$ combines a rendering fidelity term, a prior anchoring regularizer, and an explicit vessel-completion loss (Sec. 2.3). To solve Eq. 1, we chose in our method, illustrated in Fig. 1, to parameterize the vessel field as a set of 3D Gaussians, an explicit representation that can be initialized directly from the CTA-derived vessel mesh and rendered via differentiable rasterization [9], and which, once optimized, can be converted into a complete vessel mesh $\mathcal{M}_{\mathrm{C}}$.

2.2 Representation

Gaussian parameterization of 3D vessel mesh. To parameterize the vessel mesh $\mathcal{M}_{\mathrm{P}}$ for differentiable optimization, we represent each mesh vertex as the center of a 3D Gaussian kernel. Additional Gaussians are introduced to capture shape absent from the prior (Sec. 2.3). Together, these form a set of N Gaussians $\mathcal{G} = \{g_i\}_{i=1}^{N}$. Following [9], each Gaussian g_i carries a center $\boldsymbol{\mu}_i \in \mathbb{R}^3$, a rotation quaternion $\mathbf{q}_i \in \mathbb{R}^4$, and an anisotropic scale vector $\mathbf{s}_i \in \mathbb{R}^3_+$, which together define a covariance matrix $\boldsymbol{\Sigma}_i \in \mathbb{R}^{3\times 3}$. Scales are activated from raw parameters $\tilde{\mathbf{s}}_i \in \mathbb{R}^3$ via a bounded sigmoid [10, 17], $\mathbf{s}_i = (s_{\max} - s_{\min})\,\sigma(\tilde{\mathbf{s}}_i) + s_{\min}$, which confines all scales to the interval $[s_{\min}, s_{\max}]$ and prevents degenerate kernels that would collapse or explode. Each Gaussian additionally carries an opacity $\alpha_i = \phi\big(h(\boldsymbol{\mu}_i)\big) \in \mathbb{R}_{\geq 0}$, predicted from its center by a multi-resolution hash encoding h [13] followed by an MLP ϕ (2 hidden layers, 64 neurons each). Given any 3D coordinate $\mathbf{x} \in \mathbb{R}^3$, the opacity response due to Gaussian g_i is defined as [10, 18]:

$$O(\mathbf{x};\, g_i) = \alpha_i\, \exp\!\Big(-\tfrac{1}{2}(\mathbf{x}-\boldsymbol{\mu}_i)^\top \boldsymbol{\Sigma}_i^{-1}(\mathbf{x}-\boldsymbol{\mu}_i)\Big). \tag{2}$$

Summing over all Gaussians yields a continuous volumetric opacity field of space: $\rho(\mathbf{x}) = \sum_{g_i \in \mathcal{G}} O(\mathbf{x};\, g_i)$, from which the completed vessel mesh $\mathcal{M}_{\mathrm{C}}$ is extracted.

Generating DSA images from Gaussians. Because Gaussian Splatting projects 3D kernels directly onto the image plane, it enables efficient, fully differentiable rendering without discretizing a volume during training. Given a cone-beam projection geometry P_v with point source $\mathbf{o}_v$, each pixel defines a ray $\mathbf{r}(a) = \mathbf{o}_v + a\,\mathbf{d}$, where $\mathbf{d}$ is the unit-length direction toward that pixel and $a \in \mathbb{R}_{\geq 0}$ parametrizes distance along the ray. Since DSA removes background anatomy via mask subtraction, the rendered pixel value reduces to a line integral of the opacity field ρ under the Beer–Lambert law [8]: $\hat{I}(\mathbf{r}) = \int_{a_n}^{a_f} \rho(\mathbf{r}(a))\, \mathrm{d}a$, where a_n and a_f denote the near and far clipping distances. This integral is computed efficiently via X-ray rasterization [18], which splats each Gaussian $g_i \in \mathcal{G}$ onto the detector plane and accumulates contributions in a single forward pass. The complete rendered image for view v is denoted $\hat{I}_v = \mathcal{R}(\mathcal{G}, P_v)$.

2.3 Model Optimization and Output

Dual-source initialization. Solving Eq. 1 to minimize $\mathcal{L}_{\text{total}}$ requires a careful initialization of the Gaussian set $\mathcal{G}$. Therefore, we construct two complementary subsets, $\mathcal{G}_1$ and $\mathcal{G}_2$, such that $\mathcal{G}^{\text{init}} = \mathcal{G}_1 \cup \mathcal{G}_2$ with $N = N_1 + N_2$. The subset $\mathcal{G}_1$ is formed by sampling N_1 center points from the surface of the prior mesh $\mathcal{M}_{\mathrm{P}}$. The subset $\mathcal{G}_2$ is formed by sampling N_2 center points from a residual point cloud obtained by intersecting the two DSA back-projections using FDK [4] and retaining only points outside of $\mathcal{M}_{\mathrm{P}}$. This strategy seeds Gaussians in regions supported by the prior as well as in DSA-visible vessels not covered by $\mathcal{M}_{\mathrm{P}}$, providing meaningful gradients during early optimization rather than relying solely on densification.

Rendering loss. For each view v, we measure the discrepancy between the rendered image $\hat{I}_v$ and the target MIP I_v following [9]:

$$\mathcal{L}_{\text{render}} = \sum_v \Big[(1-\lambda_{\text{ssim}})\,\|\hat{I}_v - I_v\|_1 + \lambda_{\text{ssim}}\,\big(1-\text{SSIM}(\hat{I}_v,\, I_v)\big)\Big]. \tag{3}$$

where λ_{ssim} is set to 0.07 . The L_1 term drives global intensity agreement while structural similarity [15] preserves local vessel contrast.

False-negative vessel loss. With only two views, thin branches occupy few pixels and receive weak gradients under $\mathcal{L}_{\text{render}}$ alone. We add a false-negative loss over vessel regions, averaging pixel p by vessel area rather than image area:

$$\mathcal{L}_{\text{FN}} = \sum_v \frac{1}{|\Omega_v|} \sum_{p \in \Omega_v} \big[I_v(p) - \hat{I}_v(p)\big]_+, \tag{4}$$

where $\Omega_v = \{p : I_v(p) \geq \tau_{\text{v}}\}$ is the set of vessel pixels obtained by thresholding the MIP at pre-defined threshold τ_{v}. The operator $[\cdot]_+$ clamps negative values to zero so that only under-predicted pixels incur a penalty, hence the name false-negative. Over-predictions are penalized by $\mathcal{L}_{\text{render}}$. Dividing by $|\Omega_v|$ rather than the total pixel count ensures branches receive supervision proportional to their anatomical importance.

Prior-anchoring regularizer. We assume that Gaussians that lie close to the CTA vessel surface encode anatomy that is correct; without constraint, gradient pressure from the two DSA views can displace these centers from their reliable prior positions. For each center $\boldsymbol{\mu}_i$, we find its nearest point $\text{nn}(\boldsymbol{\mu}_i) \in \mathcal{M}_{\text{P}}$ and collect all centers within distance r into the set $\mathcal{A} = \{\boldsymbol{\mu}_i : \|\boldsymbol{\mu}_i - \text{nn}(\boldsymbol{\mu}_i)\| \leq r\}$. Distant Gaussians remain free to discover new vessels. Each center in $\mathcal{A}$ is tethered to the local mean $\bar{\mathbf{p}}_i$ of its k nearest points in $\mathcal{M}_{\text{P}}$, which provides a smoother target than a single neighbor, leading to the prior-anchoring loss:

$$\mathcal{L}_{\text{PA}} = \frac{1}{|\mathcal{A}|} \sum_{\boldsymbol{\mu}_i \in \mathcal{A}} \|\boldsymbol{\mu}_i - \bar{\mathbf{p}}_i\|^2. \tag{5}$$

The squared norm permits small adjustments for registration imprecision while penalizing large displacements disproportionately.

Training objective. The full objective combines the two data-fidelity losses with the prior-anchoring regularizer:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{render}} + \lambda_{\text{FN}}\,\mathcal{L}_{\text{FN}} + \lambda_{\text{PA}}\,\mathcal{L}_{\text{PA}}. \tag{6}$$

where λ_{FN} and λ_{PA} are fixed hyperparameters weights that we set to 1.5 and 0.0025 respectively. Gradients of $\mathcal{L}_{\text{total}}$ update the per-Gaussian parameters ($\boldsymbol{\mu}_i$, $\mathbf{q}_i$, $\tilde{\mathbf{s}}_i$) and the opacity network (h, ϕ) jointly. During training we follow the densification strategy of [9, 18]: Gaussians with large position gradients are cloned (if small) or split (if large). Following [10], we prune low-opacity kernels to remove Gaussians that do not contribute to the reconstructed volume.

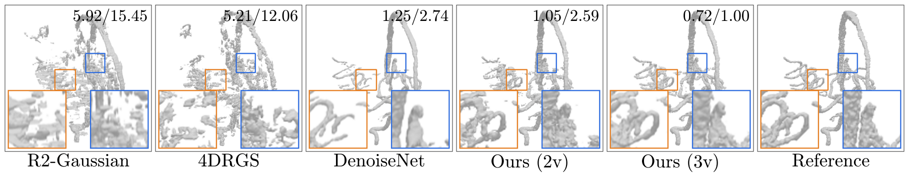

Fig. 2. Qualitative 3D reconstruction comparison (TopBrain case 001). Each panel shows the predicted vessel mesh with CD/HD (mm) annotated; corner insets highlight fine vessel detail. All methods show 2-view input; rightmost shows 3-view.

Output 3D Anatomy Completion. Finally, once the Gaussians are optimized, we obtain the complete vessels mesh $\mathcal{M}_\mathrm{C}$ by evaluating the opacity field $\rho(\mathbf{x}) = \sum_{g_i \in \mathcal{G}} O(\mathbf{x}; g_i)$ on a regular voxel grid $\mathcal{X}$ at target resolution to obtain the volume $V = \{\rho(\mathbf{x})\}_{\mathbf{x}\in\mathcal{X}}$ [18], and extract the completed vessel surface via marching cubes $\mathrm{MC}(\cdot)$ at iso-threshold τ_mc set to 0.008, yielding $\mathcal{M}_\mathrm{C} = \mathrm{MC}(V, \tau_\mathrm{mc})$.

3 Experiments

Dataset and Metrics. To evaluate our method, we use a subset of 10 CTA scans from the TopBrain dataset [14]. Each CTA volume is resampled to 256^3 and forward-projected with TIGRE [1] onto a 512×512 detector. DSA images are obtained by subtracting the mask from the fill projection. We first evaluate using AP and LAT views and then with an additional oblique view. We simulate incomplete priors $\mathcal{M}_\mathrm{P}$, by filtering the ground truth CTA complete mesh $\mathcal{M}_\mathrm{C}$ by vessel diameter at thresholds of 1.0, 1.5, and 2.0 mm. Reconstruction accuracy (in mm) is measured using Chamfer Distance (CD) and Hausdorff Distance (HD) between the reconstructed vessel mesh and the ground-truth complete mesh $\mathcal{M}_\mathrm{C}$. We also report CD_ROI and HD_ROI, by filtering vessels from the ground-truth surface at the same diameter applied to construct $\mathcal{M}_\mathrm{P}$. Vessels removed by filtering define the missing region, while vessels retained define the prior region.

Table 1. Comparison with state-of-the-art methods. Best in **bold**.

	AP, LAT		AP, LAT, Oblique	
	CD_ROI (mm)	HD_ROI (mm)	CD_ROI (mm)	HD_ROI (mm)
DenoiseNet	2.30 ± 0.48	6.30 ± 1.82	–	–
R^2-Gaussian	8.78 ± 1.97	24.70 ± 10.75	4.87 ± 1.93	18.71 ± 10.23
4DRGS	8.35 ± 2.57	23.40 ± 12.69	2.56 ± 1.38	8.50 ± 6.85
Ours	$\mathbf{2.10 \pm 0.35}$	$\mathbf{4.53 \pm 0.50}$	$\mathbf{0.86 \pm 0.11}$	$\mathbf{1.68 \pm 0.63}$

Table 2. Loss component ablation. Best in **bold**.

Views | Configuration		CD↓	HD↓	CD_{ROI}↓	HD_{ROI}↓	CD_{prior}↓	HD_{prior}↓
2	$\mathcal{L}_{render}$	1.17 ± 0.14	2.95 ± 0.66	2.02 ± 0.35	4.44 ± 0.59	0.76 ± 0.16	1.28 ± 0.91
	$\mathcal{L}_{render} + \mathcal{L}_{PA}$	1.13 ± 0.15	3.09 ± 0.67	2.14 ± 0.40	4.61 ± 0.61	0.67 ± 0.16	1.32 ± 1.14
	$\mathcal{L}_{render} + \mathcal{L}_{FN}$	1.16 ± 0.14	**2.89 ± 0.68**	**1.99 ± 0.34**	**4.34 ± 0.58**	0.76 ± 0.15	**1.27 ± 0.87**
	$\mathcal{L}_{render} + \mathcal{L}_{PA} + \mathcal{L}_{FN}$	**1.11 ± 0.14**	3.03 ± 0.64	2.10 ± 0.35	4.53 ± 0.50	**0.66 ± 0.15**	1.30 ± 1.08
3	$\mathcal{L}_{render}$	0.71 ± 0.09	1.09 ± 0.50	0.92 ± 0.13	1.91 ± 0.80	0.61 ± 0.07	0.69 ± 0.11
	$\mathcal{L}_{render} + \mathcal{L}_{PA}$	0.69 ± 0.09	1.15 ± 0.57	0.95 ± 0.13	2.06 ± 0.81	0.56 ± 0.06	0.61 ± 0.08
	$\mathcal{L}_{render} + \mathcal{L}_{FN}$	0.67 ± 0.08	**0.93 ± 0.32**	**0.85 ± 0.11**	**1.56 ± 0.60**	0.60 ± 0.07	0.69 ± 0.09
	$\mathcal{L}_{render} + \mathcal{L}_{PA} + \mathcal{L}_{FN}$	**0.64 ± 0.08**	0.94 ± 0.38	0.86 ± 0.11	1.68 ± 0.63	**0.54 ± 0.06**	**0.60 ± 0.07**

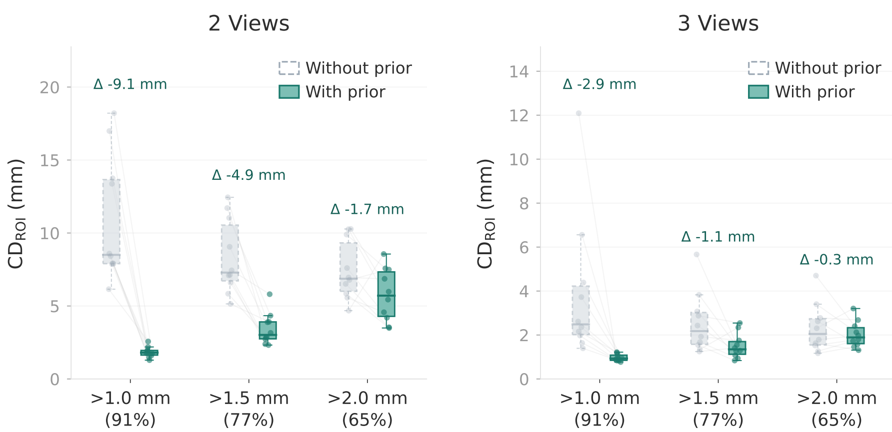

Fig. 3. Reconstruction accuracy with varying prior vessel thickness thresholds.

Comparison. We compare against three methods: R^2-Gaussian [18], a Gaussian-splatting method for tomographic reconstruction adapted to our synthetic DSA setting, 4DRGS [10], a multi-view, dynamic DSA reconstruction based on radiative Gaussian splatting, and a 3D nnU-Net denoiser (DenoiserNet) similar to the ones used in [2, 16] (bi-plane only). Our method is optimized in ≈4 min on an A100 GPU using Adam over 30,000 iterations with 60,000 Gaussians.

As shown in Table 1, our method achieves the lowest CD_{ROI} and HD_{ROI} across both settings. Reducing CD_{ROI} from 8.35 to 2.10 mm and HD_{ROI} from 23.40 to 4.53 mm for 2 views and from 2.56 to 0.86 mm and from 8.50 to 1.68 mm for 3 views, compared to the best prior baseline.

Sensitivity and ablation study. We perform an ablation study to demonstrate that each auxiliary loss targets a different error mode (see Table 2). Adding $\mathcal{L}_{FN}$ improves the missing region, lowering HD_{ROI} from 1.91 to 1.56 mm at 3 views, while $\mathcal{L}_{PA}$ tightens the prior region, reducing CD_{prior} from 0.61 to 0.56 mm. Combining both yields the best overall CD at both view counts.

We furthermore show in Table 3 that residual intersected FDK initialization outperforms all alternatives, reducing CD_{ROI} from 3.62 to 2.10 mm at 2 views and from 1.26 to 0.86 mm at 3 views compared to using no context points.

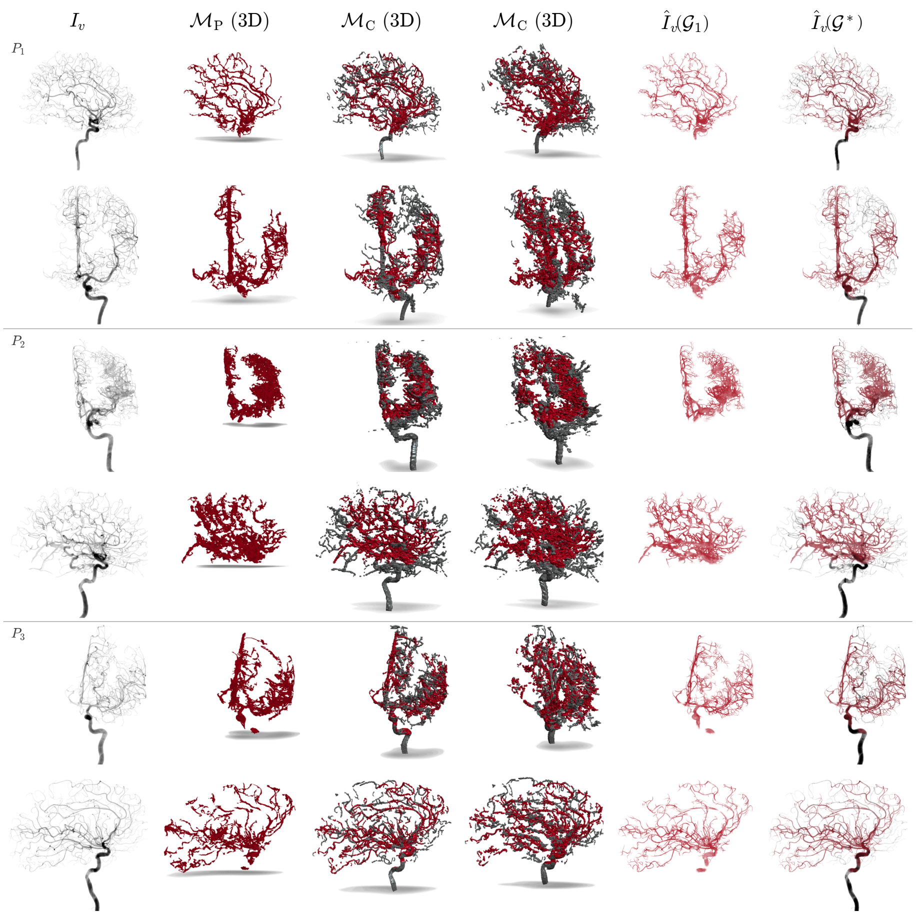

Fig. 4. Qualitative evaluation on clinical data. Each patient (P_1, P_2, P_3) is shown in AP and LAT views across two rows. From left to right: input DSA projections I_v, prior CTA mesh $\mathcal{M}_{\mathrm{P}}$, completed mesh $\mathcal{M}_{\mathrm{C}}$ from two viewpoints (prior vessels in red, newly recovered vessels in gray) and Gaussian renderings of $\mathcal{M}_{\mathrm{P}}$ and $\mathcal{M}_{\mathrm{C}}$.

Finally, Fig. 3 shows the effect of CTA prior completeness. With 2 views, adding the prior reduces $\mathrm{CD}_{\mathrm{ROI}}$ by 9.1 mm at the ≥ 1.0 mm threshold, by 4.9 mm at ≥ 1.5 mm, and by 1.7 mm at the coarsest ≥ 2.0 mm level. With 3 views the gains are smaller but consistent, with reductions of 2.9, 1.1, and 0.3 mm respectively. Adding a third view improves metrics across all prior levels.

Results on Clinical Data. To validate clinical applicability, we apply our method to clinical biplane DSA acquisitions and CTA from three patients. CTA-derived prior are obtained via manual segmentation and CTA-to-DSA registra-

Table 3. Context-point initialization methods ablation. Best in **bold**.

	AP, LAT		AP, LAT, Oblique	
	CD_{ROI} (mm)	HD_{ROI} (mm)	CD_{ROI} (mm)	HD_{ROI} (mm)
No M2 (M1 only)	3.62 ± 1.30	9.49 ± 3.50	1.26 ± 0.50	2.69 ± 1.69
Uniform	3.97 ± 1.11	10.31 ± 2.62	1.37 ± 0.42	3.07 ± 2.03
Residual FDK	2.99 ± 0.66	8.46 ± 2.34	0.95 ± 0.17	1.96 ± 1.21
Residual Intersected FDK	$\mathbf{2.10 \pm 0.35}$	$\mathbf{4.53 \pm 0.50}$	$\mathbf{0.86 \pm 0.11}$	$\mathbf{1.68 \pm 0.63}$

tion is performed using xvr [7]. Qualitative results are reported in Fig. 4 showing that the completed mesh recovers distal branches visible in DSA but absent from the CTA, at a spatial DSA resolution of 0.216 ± 0.054 mm, compared to the prior CTA resolution of 0.833 ± 0.144 mm, on average over the three patients.

4 Conclusion

In this paper we proposed to tackle the CTA-DSA imaging resolution complementarity, as an shape completion problem, where 3D vessels not visible in CTA can be reconstructed using bi-plane DSA and CTA priors, by optimizing 3D Gaussians primitives. Our method outperforms related works on both synthetic and clinical data with only AP and LAT projections, and further improves in the presence of an oblique projection, suggesting the potential for applicability in real-world clinical settings. Future work will focus on the transfer of DSA-based hemodynamics to the 3D completed vessels, and on incorporating denoising priors, which could improve the initialization and optimization. Validating the method on a dataset where high- and low-resolution CTA can be acquired, with their corresponding DSA, is also an important direction for future work.

Acknowledgments. The work was supported by the NIH Grants R01EB034223, K25EB035166 and R03EB033910. We would like to thank Danilo Thurber and Taewoong Lee for their assistance with the code implementation.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Biguri, A., Dosanjh, M., Hancock, S., Soleimani, M.: Tigre: a matlab-gpu toolbox for cbct image reconstruction. Biomedical Physics & Engineering Express **2**(5), 055010 (2016). https://doi.org/10.1088/2057-1976/2/5/055010
2. Cafaro, A., Dorent, R., Haouchine, N., Lepetit, V., Paragios, N., III, W.M.W., Frisken, S.F.: Two projections suffice for cerebral vascular reconstruction. In: Linguraru, M.G., Dou, Q., Feragen, A., Giannarou, S., Glocker, B., Lekadir, K., Schnabel, J.A. (eds.) Medical Image Computing and Computer Assisted Intervention - MICCAI 2024 - 27th International Conference, Marrakesh, Morocco, October 6-10, 2024, Proceedings, Part VII. LNCS, vol. 15007, pp. 722–731. Springer (2024). https://doi.org/10.1007/978-3-031-72104-5_69

3. Copeland, A.D., Mangoubi, R.S., Desai, M.N., Mitter, S.K., Malek, A.M.: Spatiotemporal data fusion for 3d+ t image reconstruction in cerebral angiography. IEEE transactions on medical imaging **29**(6), 1238–1251 (2010)
4. Feldkamp, L.A., Davis, L.C., Kress, J.W.: Practical cone-beam algorithm. Journal of the Optical Society of America A **1**(6), 612–619 (1984). https://doi.org/10.1364/JOSAA.1.000612
5. Frisken, S.F., Gopalakrishnan, V., Chlorogiannis, D.D., Haouchine, N., Cafaro, A., Golby, A.J., III, W.M.W., Du, R.: Spatiotemporally constrained 3d reconstruction from biplanar digital subtraction angiography. Int. J. Comput. Assist. Radiol. Surg. **20**(8), 1689–1701 (2025). https://doi.org/10.1007/S11548-025-03427-9, `https://doi.org/10.1007/s11548-025-03427-9`
6. Frisken, S.F., Haouchine, N., Du, R., Golby, A.J.: Using temporal and structural data to reconstruct 3d cerebral vasculature from a pair of 2d digital subtraction angiography sequences. Comput. Medical Imaging Graph. **99**, 102076 (2022). https://doi.org/10.1016/J.COMPMEDIMAG.2022.102076
7. Gopalakrishnan, V., Dey, N., Chlorogiannis, D.D., Abumoussa, A., Larson, A.M., Orbach, D., Frisken, S.F., Golland, P.: Rapid patient-specific neural networks for intraoperative x-ray to volume registration. CoRR **abs/2503.16309** (2025). https://doi.org/10.48550/ARXIV.2503.16309, `https://doi.org/10.48550/arXiv.2503.16309`
8. Kak, A.C., Slaney, M.: Principles of computerized tomographic imaging, Classics in applied mathematics, vol. 33. SIAM (2001)
9. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. **42**(4), 139:1–139:14 (2023). https://doi.org/10.1145/3592433, `https://doi.org/10.1145/3592433`
10. Liu, Z., Zha, R., Zhao, H., Li, H., Cui, Z.: 4drgs: 4d radiative gaussian splatting for efficient 3d vessel reconstruction from sparse-view dynamic DSA images. In: Oguz, I., Zhang, S., Metaxas, D.N. (eds.) Information Processing in Medical Imaging - 29th International Conference, IPMI 2025, Kos, Greece, May 25-30, 2025, Proceedings, Part II. LNCS, vol. 15830, pp. 361–374. Springer (2025). https://doi.org/10.1007/978-3-031-96625-5_24
11. Liu, Z., Zhao, H., Qin, W., Zhou, Z., Wang, X., Wang, W., Lai, X., Zheng, C., Shen, D., Cui, Z.: 3d vessel reconstruction from sparse-view dynamic DSA images via vessel probability guided attenuation learning. CoRR **abs/2405.10705** (2024). https://doi.org/10.48550/ARXIV.2405.10705, `https://doi.org/10.48550/arXiv.2405.10705`
12. Mkhize, N.N., Mngomezulu, V., Buthelezi, T.E.: Accuracy of ct angiography for detecting ruptured intracranial aneurysms. SA Journal of Radiology **27**(1), 2636 (2023)
13. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Trans. Graph. **41**(4) (2022). https://doi.org/10.1145/3528223.3530127
14. TopBrain Challenge Organizers: Topbrain segmentation challenge for whole brain vessel anatomy: Data. `https://topbrain2025.grand-challenge.org/data/` (2025), grand Challenge data page, last updated 2025-11-16, accessed 2026-02-21
15. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. **13**(4), 600–612 (2004). https://doi.org/10.1109/TIP.2003.819861
16. Wu, S., Kaneko, N., Mendoza, S., Liebeskind, D.S., Scalzo, F.: 3d reconstruction from 2d cerebral angiograms as a volumetric denoising problem. In: Bebis, G., Ghiasi, G., Fang, Y., Sharf, A., Dong, Y., Weaver, C.E., Leo, Z.,

Jr., J.J.L., Kohli, L. (eds.) Advances in Visual Computing - 18th International Symposium, ISVC 2023, Lake Tahoe, NV, USA, October 16-18, 2023, Proceedings, Part I. Lecture Notes in Computer Science, vol. 14361, pp. 382–393. Springer (2023). https://doi.org/10.1007/978-3-031-47969-4_30, https://doi.org/10.1007/978-3-031-47969-4_30

17. Xu, Y., Shi, Z., Wang, Y., Chen, H., Yang, C., Peng, S., Shen, Y., Wetzstein, G.: GRM: large gaussian reconstruction model for efficient 3d reconstruction and generation. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XV. Lecture Notes in Computer Science, vol. 15073, pp. 1–20. Springer (2024). https://doi.org/10.1007/978-3-031-72633-0_1, https://doi.org/10.1007/978-3-031-72633-0_1
18. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: R^2-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In: Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J.M., Zhang, C. (eds.) Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024 (2024)
19. Zhao, H., Zhou, Z., Wu, F., Xiang, D., Zhao, H., Zhang, W., Li, L., Li, Z., Huang, J., Hu, H., Liu, C., Wang, T., Liu, W., Ma, J., Yang, F., Wang, X., Zheng, C.: Self-supervised learning enables 3d digital subtraction angiography reconstruction from ultra-sparse 2d projection views: A multicenter study. Cell Reports Medicine **3**(10), 100775 (Oct 2022). https://doi.org/10.1016/j.xcrm.2022.100775
20. Zhou, Z., Zhao, H., Fang, J., Xiang, D., Chen, L., Wu, L., Wu, F., Liu, W., Zheng, C., Wang, X.: Tiavox: Time-aware attenuation voxels for sparse-view 4d DSA reconstruction. CoRR **abs/2309.02318** (2023). https://doi.org/10.48550/ARXIV.2309.02318
21. Zuo, J.: 2d to 3d neurovascular reconstruction from biplane view via deep learning. In: 2021 2nd International Conference on Computing and Data Science (CDS). pp. 383–387 (2021). https://doi.org/10.1109/CDS52072.2021.00071