
4D Vessel Reconstruction for Benchtop Thrombectomy Analysis

Ethan Nguyen^{1,2}, Javier Carmona, Ph.D.³, Arisa Matsuzaki², Naoki Kaneko, M.D., Ph.D.^{4,*}
and Katsushi Arisaka, Ph.D.^{2,*}

¹UCLA Health

²UCLA Physics and Astronomy

³Chan Zuckerberg Biohub Network

⁴Ronald Reagan UCLA Medical Center

*Authors to whom any correspondence should be addressed.

E-mail: ethunguy@ucla.edu, arisaka@physics.ucla.edu; nkaneko@mednet.ucla.edu

Keywords: Mechanical thrombectomy, Gaussian Splatting, Multi-view 3D reconstruction, 4D reconstruction, Cerebrovascular phantom, Surface stress analysis, Synthetic validation, Biomechanical imaging

Abstract

Introduction: Mechanical thrombectomy can cause vessel deformation and procedure-related injury. Benchtop models are widely used for device testing, but time-resolved, full-field 3D vessel-motion measurements remain limited.

Methods: We developed a nine-camera, low-cost multi-view workflow for benchtop thrombectomy in silicone middle cerebral artery phantoms (2160p, 20 fps). Multi-view videos were calibrated, segmented, and reconstructed with 4D Gaussian Splatting. Reconstructed point clouds were converted to fixed-connectivity edge graphs for region-of-interest (ROI) displacement tracking and a relative surface-based stress proxy. Stress-proxy values were derived from edge stretch using a Neo-Hookean mapping and reported as comparative surface metrics. A synthetic Blender pipeline with known deformation provided geometric and temporal validation.

Results: In synthetic bulk translation, the stress proxy remained near zero for most edges (median ≈ 0 MPa; 90th percentile 0.028 MPa), with sparse outliers. In synthetic pulling (1–5 mm), reconstruction showed close geometric and temporal agreement with ground truth, with symmetric Chamfer distance of 1.714–1.815 mm and precision of 0.964–0.972 at $\tau = 1$ mm. In preliminary benchtop comparative trials (one trial per condition), cervical aspiration catheter placement showed higher max-median ROI displacement and stress-proxy values than internal carotid artery terminus placement.

Conclusion: The proposed protocol provides standardized, time-resolved surface kinematics and comparative relative displacement and stress proxy measurements for thrombectomy benchtop studies. The framework supports condition-to-condition comparisons and methods validation, while remaining distinct from absolute wall-stress estimation. Implementation code and example data are available at <https://ethanuser.github.io/vessel4D>.

1 Introduction

Mechanical thrombectomy is an effective treatment for acute ischemic stroke due to large-vessel occlusion, but procedure-related vessel injury and hemorrhagic complications remain clinically important. Recent work has highlighted post-procedural subarachnoid hemorrhage (SAH) as a meaningful complication associated with vessel overextension or displacement during retrieval, while multicenter data show that vessel perforation, although relatively uncommon, is associated with poor functional outcomes and higher mortality [1, 2]. These risks motivate experimental methods that can localize where deformation concentrates during retrieval maneuvers rather than relying only on aggregate procedural outcomes.

Benchtop and in-vitro thrombectomy platforms provide controlled and repeatable environments for device development, training, and mechanistic testing [3, 4]. However, many current benchtop measurements emphasize recanalization, distal embolization, clot–device interaction, or force/friction metrics, rather than time-resolved, full-field 3D vessel-surface kinematics. Recent examples include studies of stent-retriever removal forces in tortuous models, friction-focused analyses, vessel deviation measured at a limited set of landmarks, and objective quantification of thrombus deformation during retrieval [5, 6, 1, 7]. This leaves a gap between global or sparse measurements and spatially resolved

characterization of vessel-surface deformation at anatomically meaningful regions during thrombectomy.

Existing optical and imaging approaches also pose practical constraints for this application. Multi-view 3D digital image correlation (3D-DIC) can provide full-field deformation measurements, but it depends on stable visible surface texture or applied speckle patterns, accurate multi-camera calibration, and adequate view coverage across time [8, 9, 10]. More broadly, dynamic multi-view reconstruction requires explicit treatment of temporal coherence, since temporally unstable geometry is a recognized challenge in dynamic scene reconstruction [11, 12]. Clinical dynamic CT methods such as 4D-CTA address a different imaging niche in neurovascular imaging and are not the focus of the present benchtop optical workflow [13].

This study presents a multi-view dynamic reconstruction protocol for silicone middle cerebral artery (MCA) thrombectomy phantoms, combining multi-view camera acquisition, segmentation, 4D Gaussian Splatting reconstruction, fixed-connectivity edge-graph construction, region-of-interest (ROI) displacement tracking, and a relative surface-based stress proxy [12]. The stress proxy is used as a comparative surface metric and is not interpreted as absolute wall stress.

The study aims are: (1) to define a standardized end-to-end acquisition-to-analysis protocol for dynamic benchtop thrombectomy imaging; (2) to formalize ROI-level displacement and relative stress-proxy metrics on a fixed deforming edge graph; (3) to validate geometric and temporal behavior in synthetic data with known motion; and (4) to demonstrate comparative benchtop application in two aspiration catheter (AC) placement conditions (cervical internal carotid artery (ICA) vs ICA terminus).

2 Materials and Methods

2.1 Pipeline overview

The workflow comprises acquisition, calibration, segmentation, dynamic reconstruction, fixed-connectivity edge-graph generation, ROI metric extraction, and synthetic/benchtop evaluation (Fig. 1). Representative qualitative outputs across these stages are shown in Fig. 2. The benchtop phantom was a patient-specific silicone ICA/MCA model prepared as described previously [14]. Synthetic validation and benchtop experiments used a variant with an additional loop connecting the M2 branches, whereas Fig. 2 uses a different ICA/MCA silicone model without that loop for illustration only. Typical per-experiment end-to-end processing from multi-view video feeds to displacement/stress-proxy graph generation was approximately 20 minutes for automated reconstruction plus 10 minutes for user-guided preprocessing/curation.

2.1.1 Benchtop demonstration scope To clarify interpretation of downstream outputs, benchtop data were used as a methods demonstration, and one representative trial was analyzed per AC placement condition. A Trevo NXT stent retriever was deployed from the inferior M2 branch to the M1/M2 bifurcation together with a Toro 88 aspiration catheter; no clot analog was present. Analysis was restricted to the retraction interval after pullback began. We defined $t = 0$ as linear-actuator pullback onset. Displayed peak-deformation frames were chosen by maximizing whole-vessel median displacement and whole-vessel median stress-proxy values; these maxima coincided in the shown experiments. The same processing and metric definitions were then applied across both conditions.

2.1.2 Imaging hardware and acquisition A low-cost multi-view rig comprising nine Arducam IMX586 cameras was used to record raw 2160p video at 20 fps. Cameras were mounted on a dodecahedron-like frame constructed from PVC pipes, 3D-printed connecting joints, and custom 3D-printed camera mounts, with five cameras on the lower ring and four on the upper ring. The cameras were manually aimed toward a common central target region and fixed after calibration. Preliminary synthetic-data testing indicated that nine views were sufficient for 3D recovery, although additional views would be expected to further improve reconstruction quality. The imaging hardware (PVC structure, printed parts, mounts, fasteners, cameras, and lighting) cost less than \$1500 in total.

For the reported configuration, camera-to-target distances ranged from approximately 230 to 271 mm (mean ≈ 247 mm). The M1 mid segment was positioned near the center of the shared field of view. Views were acquired through a tiled OBS composite workflow and later cropped to 1200×1200 pixels, with crop centers aligned to a common chessboard origin to simplify downstream 4DGS processing. Although inter-camera skew was not independently hardware-validated, no visually apparent offset was observed across shared motion events, and the analyzed deformation evolved over seconds rather than impulsive frame-to-frame motion; residual mismatch was therefore expected to be small relative to the motion timescale studied here. The usable imaged volume after clipping was approximately $70 \times 90 \times 60$ mm (width \times height \times depth), although the effective




Figure 1. Overview of the experimental and computational pipeline. Multi-camera chessboard calibration yields camera parameters for nine-view vessel recordings. After SAM2-based segmentation, cropping, and dataset preparation, the deforming vessel is reconstructed with 4D Gaussian Splatting and filtered to obtain a time-varying point cloud. The point cloud is converted into a fixed-connectivity edge graph using clustering, DBSCAN, Delaunay graph pruning, and one-time manual curation. Downstream outputs are region-of-interest displacement magnitude (mm) and edge-based stress-proxy magnitude $|\sigma_e|$ (MPa). Square boxes denote inputs/outputs and rounded gray boxes denote processing steps.

working volume could be adjusted depending on camera focus and the region of interest. Approximately 20–30 s were recorded per experiment, then temporally downsampled so that roughly 1000 total frames were used for 4DGS processing. The rig layout and representative multi-view raw frames are shown in Fig. 2A–2B. Representative raw benchtop footage is provided in Video S1.

2.1.3 Illumination and surface feature preparation Red fluorescent microspheres (approximately 300 μm) were applied to the vessel exterior. UV illumination (approximately 400 nm) was used to improve feature contrast on the silicone surface, as illustrated by the representative raw frames in Fig. 2B. This preparation improved visible surface texture before segmentation.

2.1.4 Multi-camera calibration and metric scaling Extrinsic parameters were estimated using planar chessboard calibration [15] implemented with OpenCV `calibrateCamera` [16] from multiple board orientations using a 6×8 board, with the camera calibration process shown in Fig. 2A. Additional side-by-side nine-view calibration panels for real and synthetic setups are provided in Supplementary Fig. S3. The mean reprojection error across cameras was approximately 0.134 pixels (range 0.112–0.164 pixels). Metric scaling was obtained by matching reconstructed chessboard geometry to measured physical board dimensions. These calibrated parameters were used directly in 4DGS reconstruction.

2.1.5 Deformation apparatus and pullback protocol A linear actuator imposed controlled pullback at 4 mm s^{-1} . The vessel trunk/base was fixed relative to the camera rig to reduce global drift during acquisition. The analyzed interval was the retraction maneuver after pullback onset. This same temporal reference was used throughout downstream tracking and metric extraction.

2.1.6 Video segmentation and background suppression Each camera stream was segmented independently using SAM2 video prediction (`sam2.1_hiera_large.pt`) [17]. For each view, prompts

were placed once on the first frame and propagated through the sequence; masks were reinitialized only when flicker or incorrect vessel coverage appeared. The same segmentation protocol was used for all sequences. Non-vessel pixels were set to black, and no additional mask post-processing beyond cropping/alignment was applied (representative output in Fig. 2C). The resulting masked frames were then passed to dynamic reconstruction.

2.1.7 Dynamic reconstruction Calibrated multi-view frames were reconstructed with 4D Gaussian Splatting (4DGS) using the implementation of Wu et al. [12]. One model was trained per sequence for 20,000 iterations on an NVIDIA RTX 4070 Ti GPU, yielding the representative dynamic point cloud shown in Fig. 2D. Exports contained Gaussian primitive centers and colors over time; primitives with small radii, low opacity, or no color were filtered out before downstream graph construction. These filtered primitive trajectories were then converted to a fixed graph representation.

2.1.8 Point-cloud post-processing and fixed edge-graph construction To support consistent temporal metric extraction, reconstructed primitives were converted to a fixed vertex set and static edge-graph connectivity. For each sequence, frames are indexed by $t = 0, \dots, T - 1$, and the resulting clustered vertices and graph structure are illustrated in Fig. 2E.

Initial clustering used KMeans in RGB ($K_c = 5$; [18]) followed by DBSCAN ([19]), implemented in scikit-learn [20], within each color group, with $\varepsilon_s = 0.7$ mm (eps) and $m_s = 3$ (minPts). This two-stage strategy was used first to prevent points with different colors from being grouped together, and then to separate spatially distinct local components within each color group for downstream tracking. To formalize this procedure, let $\mathbf{p}_n(t) \in \mathbb{R}^3$ denote the XYZ position of primitive n at frame t , and let $g_n \in \{1, \dots, K_c\}$ denote its KMeans color-group label at $t = 0$. Formally, for color group c :

$$\{\mathbf{p}_n(0) : g_n = c\} \xrightarrow{\text{DBSCAN}(\varepsilon_s, m_s)} \{\mathcal{C}_{c,r}\}_{r=1}^{R_c}, \quad (1)$$

where $\mathcal{C}_{c,r}$ are DBSCAN-connected spatial components (clusters) within color group c . Cluster memberships $\ell_n \in \{1, \dots, K\}$ from $t = 0$ were held fixed, where K is the number of retained clusters, and vertex positions were computed by centroid averaging:

$$\mathbf{x}_k(t) = \frac{1}{|\mathcal{I}_k|} \sum_{n \in \mathcal{I}_k} \mathbf{p}_n(t), \quad \mathcal{I}_k = \{n : \ell_n = k\}, \quad (2)$$

with displacement relative to baseline:

$$\mathbf{u}_k(t) = \mathbf{x}_k(t) - \mathbf{x}_k(0). \quad (3)$$

If $\mathcal{I}_k = \emptyset$ for a frame, the previous vertex position was retained.

Static connectivity was generated on $\{\mathbf{x}_k(0)\}_{k=1}^K$ using a 3D Delaunay neighborhood graph (SciPy spatial routines [21]) and edge-length pruning. Candidate edges (i, j) were retained when:

$$d_{ij} \leq \mu_d + \gamma \sigma_d, \quad d_{ij} = \|\mathbf{x}_i(0) - \mathbf{x}_j(0)\|_2, \quad (4)$$

where μ_d and σ_d are the candidate edge-length mean and standard deviation; $\gamma = 0.25$ in reported analyses. Let \mathcal{E} denote the retained edge set of the pruned Delaunay graph. Neighborhoods $\mathcal{N}(k)$ for spatial coherence were defined by adjacency in this pruned graph at $t = 0$. This fixed graph representation was then used for curation, smoothing, and ROI aggregation.

2.1.9 Manual curation and topology locking To reduce obvious topological artifacts before temporal analysis, manual curation was performed once on the initial frame using a custom PyVista-based 3D editor to remove extraneous clusters, standardize vessel crop length, and remove implausible edges (e.g., edges spanning nearby trunks). The curated cluster set and edge topology were then reused unchanged for all subsequent frames.

2.1.10 Spatial coherence filtering To suppress isolated local tearing in the tracked displacement field before ROI-level aggregation, we applied a mild spatial regularization on the fixed cluster-adjacency graph defined above, using the curated saved edge set when available and otherwise the pruned Delaunay graph. This update can be viewed as relaxed graph/Laplacian smoothing (low-pass fairing), while the residual-dependent weights limit smoothing across locally inconsistent motions, in the spirit of anisotropic diffusion and robust piecewise-smooth motion regularization [22, 23, 24, 25]. Standard reported analyses used one iteration with $\alpha = 0.1$, robust weighting

enabled, $\kappa = 2.5$, and displacement reference $\mathbf{x}_k(0)$ (“ t_0 ” reference mode in code), although the number of smoothing iterations is configurable. For each frame, with residuals and neighbor averages recomputed at each iteration,

$$\mathbf{u}_k \leftarrow (1 - \alpha)\mathbf{u}_k + \alpha \frac{\sum_{j \in \mathcal{N}(k)} w_{kj} \mathbf{u}_j}{\sum_{j \in \mathcal{N}(k)} w_{kj} + \varepsilon}, \quad (5)$$

with residuals $r_{kj} = \|\mathbf{u}_k - \mathbf{u}_j\|_2$ and robust weights

$$w_{kj} = \begin{cases} 1, & r_{kj} \leq \tau, \\ \tau / (r_{kj} + \varepsilon), & r_{kj} > \tau, \end{cases} \quad \tau = \kappa s, \quad (6)$$

where the weight function corresponds to a Huber-like IRLS weighting, so neighbors with similar displacements are averaged more strongly whereas neighbors with large local disagreements are downweighted rather than fully trusted [26]. The robust scale is computed once per frame over edges using a median-absolute-deviation (MAD) estimate,

$$s = \text{median}_{(i,j) \in \mathcal{E}} |r_{ij} - \text{median}_{(a,b) \in \mathcal{E}} r_{ab}| + \varepsilon, \quad (7)$$

which provides an adaptive threshold that is less sensitive to outlying residuals than variance-based scaling [27]. Here $\varepsilon = 10^{-12}$ is a numerical stabilizer used in code for both the MAD scale and denominator safeguards. Clusters with no points in a frame retained their previous positions and were not themselves updated by the smoothing step. The filtered displacements were then used for ROI displacement and stress-proxy computation. To assess the effect of this step, we also repeated the synthetic validation analysis without spatial coherence filtering; the corresponding ablation results are provided in Supplementary Fig. S1, Supplementary Fig. S2, and Supplementary Table S3.

2.1.11 ROI displacement and stress-proxy metrics At this stage, graph trajectories were summarized into ROI-level outputs for comparison. ROIs were defined on anatomical landmarks (Fig. 4); each ROI denotes a selected set of vertex indices on the fixed graph. For benchtop data, ROIs were manually selected by the researcher on the initial fixed graph and reused over time, with placement matched as closely as possible across conditions. For synthetic data, ROIs were defined as sphere-based neighborhoods around fixed centers and reused for both ground truth (GT) and reconstruction because initial synthetic frames were spatially aligned. Per-frame ROI displacement was the vertex-wise median magnitude:

$$d_{\text{ROI}}^{\text{med}}(t) = \text{median}_{k \in \text{ROI}} \|\mathbf{u}_k(t)\|_2. \quad (8)$$

Condition-level displacement summary used the temporal maximum of this median:

$$d_{\text{ROI}}^{\text{max}} = \max_t d_{\text{ROI}}^{\text{med}}(t). \quad (9)$$

To map edge deformation into a comparative stress-like quantity, for each edge $e = (i, j) \in \mathcal{E}$ in the fixed connectivity graph let $\ell_e(t) = \|\mathbf{x}_i(t) - \mathbf{x}_j(t)\|_2$ denote the edge length at frame t . Stretch was then defined by:

$$\lambda_e(t) = \frac{\ell_e(t)}{\ell_e(0)}, \quad (10)$$

and mapped to a Neo-Hookean uniaxial stress-like quantity. For incompressible isotropic hyperelasticity with traction-free lateral contraction, the uniaxial Cauchy stress reduces to the Neo-Hookean form [28]:

$$\sigma_e = \mu (\lambda_e^2 - \lambda_e^{-1}), \quad (11)$$

with μ derived from silicone Young’s modulus E . For near-incompressible material, $\mu = E/[2(1 + \nu)] \approx E/3$ at $\nu \approx 0.5$. Using $E = 1.15$ MPa for silicone [29] gives $\mu \approx 0.383$ MPa.

Per-frame ROI stress summary used the median of $|\sigma_e|$ over edges with both endpoints in the ROI, with condition-level max-median reporting analogous to Eq. 9. Representative stress-proxy output is shown in Fig. 2F. The stress proxy is interpreted as a relative, surface-based comparative metric and not as absolute wall stress. These ROI summaries were then carried into synthetic and benchtop comparisons.




Figure 2. Experimental setup and major stages of the reconstruction and analysis pipeline. (A) Nine-camera acquisition rig used for multi-view imaging, shown during chessboard-based calibration; long arrows (shown in blue) indicate the cameras’ viewing directions toward the working volume. (B) Example multi-view raw frames acquired during benchtop thrombectomy imaging under UV illumination, showing the fluorescent-beaded silicone vessel phantom. (C) Corresponding segmented vessel frames after preprocessing with SAM2, used to isolate vessel for downstream reconstruction. (D) Example time-varying 3D point-cloud reconstruction obtained from 4D Gaussian Splatting. (E) Fixed set of tracked surface vertices obtained from point-cloud clustering and used for fixed edge-graph construction and temporal tracking. (F) Edge-based stress-proxy magnitude $|\sigma_e|$ (MPa) on the resulting edge graph at the representative timepoint.




Figure 3. Synthetic validation workflow (schematic; no single frame/timepoint metric shown). Ground-truth vessel geometries were deformed in Blender using controlled bulk-translation and localized-pulling conditions, rendered with the same camera geometry as the physical setup, and passed through the same camera geometry and downstream reconstruction pipeline, with segmentation bypassed because the renders use transparent backgrounds. The resulting reconstructed point clouds and fixed edge graphs were compared with ground truth using Chamfer distance, temporal consistency metrics, and deformation/stress-proxy agreement analyses.




Figure 4. Regions of the ICA/MCA surface used for regional aggregation and analysis. R1: M1 mid segment; R2: M1/M2 bifurcation; R3: proximal inferior M2; R4: proximal superior M2; R5: ICA terminus/M1 origin.

2.2 Synthetic validation design and metric definitions

Synthetic experiments were used to check geometric and temporal behavior against known motion before interpreting benchtop comparisons.

To enable quantitative validation with ground truth, a known vessel surface mesh was rendered in Blender with matched camera geometry (Fig. 3) under two conditions: bulk translation and localized pulling near the M1/M2 bifurcation (1–5 mm nominal pulls). Because these Blender renders used transparent backgrounds, this validation isolates the reconstruction-and-analysis stages and does not test the SAM 2-based background-filtering/segmentation step required for benchtop videos. The synthetic vessel appearance also only approximates the real phantom’s optical surface characteristics, including vessel-surface contrast and bead size/visibility. Metrics used for comparison followed established dynamic-surface evaluation definitions and protocol [30]. GT stress-proxy maps were computed from the Blender-exported ground-truth edge graph using the same stress-proxy definition (Eqs. 10–11).

Symmetric Chamfer distance (CD) was defined as:

$$\text{CD}(P, G) = \frac{1}{|P|} \sum_{p \in P} \min_{g \in G} \|p - g\|_2 + \frac{1}{|G|} \sum_{g \in G} \min_{p \in P} \|g - p\|_2, \quad (12)$$

where P and G are reconstructed and ground-truth point sets. We also report $\text{CD}_{\text{norm}} = \text{CD}/\tilde{d}_{\text{GT}}$, where \tilde{d}_{GT} is the median nearest-neighbor GT point spacing.

Non-temporal metrics (CD, CD_{norm} , precision, recall, F-score) were evaluated at maximum deformation (last frame, $t = T - 1$). Temporal consistency used frame-to-frame CD disagreement, following the same protocol [30]. For last-frame non-temporal metrics, directed distances were:

$$d_{P \rightarrow G}(p) = \min_{g \in G_{T-1}} \|p - g\|_2, \quad (13)$$

$$d_{G \rightarrow P}(g) = \min_{p \in P_{T-1}} \|g - p\|_2, \quad (14)$$

with temporal terms for $t = 0, \dots, T - 2$:

$$\text{CD}_t^{\text{pred}} = \text{CD}(P_t, P_{t+1}), \quad \text{CD}_t^{\text{gt}} = \text{CD}(G_t, G_{t+1}), \quad (15)$$

and temporal disagreement:

$$\Delta\text{CD} = \frac{1}{T-1} \sum_{t=0}^{T-2} \left| \text{CD}_t^{\text{pred}} - \text{CD}_t^{\text{gt}} \right|, \quad (16)$$

reported also as $\Delta\text{CD}_{\text{rel}}$ after normalization by the median ground-truth frame-to-frame CD.

At tolerance $\tau = 1$ mm by default (evaluated at $t = T - 1$),

$$\text{Precision} = \frac{1}{|P_{T-1}|} \sum_{p \in P_{T-1}} \mathbb{1}[d_{P \rightarrow G}(p) < \tau], \quad (17)$$

$$\text{Recall} = \frac{1}{|G_{T-1}|} \sum_{g \in G_{T-1}} \mathbb{1}[d_{G \rightarrow P}(g) < \tau], \quad (18)$$

and

$$\text{F-score} = \begin{cases} \frac{2 \text{Precision Recall}}{\text{Precision} + \text{Recall}}, & \text{if Precision} + \text{Recall} > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (19)$$

Displacement error was defined as $d_{\text{ROI,ours}}^{\text{max}} - d_{\text{ROI,gt}}^{\text{max}}$, with percent error computed relative to $d_{\text{ROI,gt}}^{\text{max}}$. Stress bias was defined analogously as the Ours–GT difference in max-median ROI stress proxy. Regression and Bland–Altman analyses [31] used 25 paired points (5 ROIs \times 5 pull magnitudes). Together, these metrics define the synthetic benchmarks reported in Results.

3 Results

Results are presented in the same sequence as the workflow: synthetic control, synthetic pulling validation, and benchtop comparison.

Synthetic rigid-motion rejection

The rigid-translation control yielded near-zero stress-proxy values for most edges: median ≈ 0 MPa, 90th percentile 0.028 MPa, and a sparse outlier tail with maximum 1.898 MPa (Fig. 5). Regional displacement error was 0.091 mm (0.64%), $\Delta\text{CD}_{\text{rel}}$ was 0.085, and stress bias was 0.018 MPa (Table 1). Raw/displacement/stress videos are provided in [Video S6](#), [Video S7](#), and [Video S8](#). With this control behavior established, we next evaluated localized pulling.




Figure 5. Rigid-body translation control (synthetic vessel), evaluated at the final frame of the bulk-translation sequence. (A) Displacement magnitude (mm) for ground truth (GT) and reconstruction (Ours). (B) Edge-based stress-proxy magnitude $|\sigma_e|$ (MPa); GT remains near zero under rigid motion, and Ours is near zero for most edges with sparse outliers.

Synthetic deformation fidelity

Under localized pulling, reconstruction preserved condition-dependent deformation patterns with strong task-relevant agreement. Across 1–5 mm pulling conditions, symmetric CD ranged from 1.714 to 1.815 mm, CD_{norm} from 2.872 to 3.041, and $\Delta\text{CD}_{\text{rel}}$ from 0.100 to 0.139 (Table 1 and Fig. 7; expanded directed-distance and overlap metrics are provided in Supplementary Table S1). Precision at $\tau = 1$ mm ranged from 0.964 to 0.972. Regional displacement error ranged from 0.047 mm to 0.190 mm, and stress bias ranged from 0.004 MPa to 0.007 MPa (Fig. 6).

Agreement analysis of max-median ROI summaries showed displacement slope 0.987, intercept -0.112 mm, and $R^2 = 0.992$, with Bland–Altman bias -0.142 mm and limits -0.385 mm to

0.100 mm (Fig. 8). Stress-proxy agreement showed slope 1.068, intercept 0.002 MPa, and $R^2 = 0.969$, with bias 0.004 MPa and limits -0.002 MPa to 0.010 MPa. Without spatial coherence filtering, displacement agreement remained similar, whereas stress-proxy agreement became less favorable, with larger positive bias and wider limits of agreement (Supplementary Fig. S1, Supplementary Fig. S2, and Supplementary Table S3). Full ROI-wise condition summaries are provided in Supplementary Table S2. Raw/displacement/stress videos are provided in Video S9, Video S10, and Video S11. These synthetic findings motivate the same ROI summaries in the benchtop comparison.




Figure 6. Synthetic pulling deformation (1–5 mm), evaluated at each condition’s final frame. Rows compare ground truth (GT) and reconstruction (Ours). (A) Displacement magnitude (mm). (B) Edge-based stress-proxy magnitude $|\sigma_e|$ (MPa).

Benchtop comparative application

Using the same validated metrics, the benchtop comparison (one trial per condition) showed broader and larger displacement/stress-proxy fields for AC cervical placement than ICA terminus placement (Fig. 9).

Max-median ROI values (Fig. 10) were higher in cervical vs terminal placement for displacement: distal M1 segment 3.776 vs 1.151 mm, MCA bifurcation (M1/M2) 3.144 vs 0.846 mm, and inferior M2 division 1.235 vs 0.528 mm. Max-median stress proxy was also higher in cervical vs terminal placement: distal M1 segment 0.109 vs 0.058 MPa, bifurcation 0.062 vs 0.035 MPa, and inferior M2 division 0.056 vs 0.034 MPa. Corresponding benchtop videos are provided in Video S2, Video S3,




Figure 7. GT-versus-reconstruction agreement across synthetic pull magnitudes (condition-level comparison; not a single-frame plot). Pairwise heatmaps report (A) Chamfer distance (CD, mm), (B) temporal Chamfer delta (ΔCD , mm), and (C) F-score; metric definitions follow Sec. 2.2 (Eqs. 12, 16, and 19). Best agreement occurs along the diagonal (matched pull conditions).

Table 1. Synthetic validation metrics by deformation condition. Definitions for CD, CD_{norm} , ΔCD_{rel} , precision, recall, and F-score follow Sec. 2.2 (Eqs. 12, 16, 18, and 19). Displacement error (mm, %) and stress bias (MPa) are reported as Ours – GT using the max-median ROI definitions described in Sec. 2.1.11.

Condition	CD		ΔCD_{rel}	Precision	Displacement Error		Stress Bias
	(mm)	CD_{norm}			(mm)	(%)	
Bulk	1.881	3.152	0.085	0.919	0.091	0.64	0.018
1 mm	1.815	3.041	0.139	0.970	0.047	7.17	0.004
2 mm	1.752	2.937	0.100	0.967	0.094	7.82	0.004
3 mm	1.738	2.912	0.114	0.964	0.087	4.98	0.007
4 mm	1.745	2.924	0.100	0.972	0.164	7.38	0.007
5 mm	1.714	2.872	0.105	0.968	0.190	5.28	0.006

[Video S4](#), and [Video S5](#). Given single-trial sampling, these values are interpreted as comparative directional observations.

Scope of reported quantities

Consistent with the methods scope, stress values are reported as relative surface-based stress-proxy magnitudes ($|\sigma_e|$) from reconstructed surface deformation. They are unsigned comparative metrics and do not represent absolute wall stress or tensile/compressive sign. The rigid-motion control supports this interpretation by showing near-zero stress proxy under bulk translation for most edges.




Figure 8. Agreement analyses for max-median ROI metrics between reconstruction and ground truth (condition-level summaries, not single-frame values): correlation and Bland–Altman for displacement magnitude (A–B, mm) and relative surface-based stress-proxy magnitude (C–D, MPa).




Figure 9. Benchtop application demonstration at peak deformation for a single representative trial per AC placement condition (cervical and terminal). (A) Reconstructed displacement magnitude (mm) relative to the initial frame. (B) Corresponding edge-based stress-proxy magnitude $|\sigma_e|$ (MPa) at the same peak-deformation timepoint.




Figure 10. Regional max-median displacement and stress-proxy during thrombectomy for a single representative trial per AC placement condition. Bars report the maximum over time of the regional median displacement magnitude (A, mm) and edge-based stress-proxy magnitude (B, MPa) across distal M1 segment, MCA bifurcation (M1/M2), and inferior M2 division.

4 Discussion

This methods/validation study defines and tests a multi-view reconstruction protocol for dynamic benchtop thrombectomy experiments, moving from workflow definition to synthetic validation and then benchtop comparison. Three findings support the utility of the approach. First, the rigid-motion control produced near-zero stress proxy for most edges, indicating expected suppression of spurious internal loading under bulk translation. Second, synthetic pulling experiments showed moderate absolute geometric error but strong temporal and task-relevant agreement across conditions, supporting use of the framework for comparative kinematic analysis. Third, the benchtop comparison showed a consistent directional difference between cervical ICA and ICA terminus AC placement in both max-median displacement and max-median stress-proxy outputs.

These trends are broadly consistent with prior work showing that thrombectomy strategy and device configuration influence vessel loading and procedural risk. In vitro evidence has shown that combined stent-retriever and aspiration-catheter techniques can reduce vessel stretching relative to stent retriever alone in selected MCA configurations [32]. Clinical studies likewise suggest tradeoffs between distal access strategy, reperfusion performance, and vascular injury: more distal access catheter tip position can improve reperfusion but has also been associated with higher vascular injury rates [33], and recent cohort analyses have linked device/strategy factors to recanalization performance [34]. Relative to prior benchtop studies that emphasize removal forces, friction, or sparse landmark-based motion surrogates [5, 6, 32], the present framework contributes time-resolved, spatially resolved vessel-surface kinematics and regional comparative loading maps at anatomically defined ROIs.

The benchtop comparison is intended as a comparative demonstration rather than an inferential biomechanical conclusion because each condition includes one representative trial. Likewise, the synthetic validation is designed to test whether the reconstruction-and-analysis pipeline recovers known imposed surface motion, preserves temporal behavior, and suppresses spurious stress-proxy response under rigid translation; it does not validate absolute biomechanical realism, device-clot-vessel interaction mechanics, or the constitutive fidelity of the stress mapping. It also does not test the video-segmentation stage under realistic benchtop imaging conditions, because the Blender renders use transparent backgrounds rather than the cluttered backgrounds present in the physical videos that must be filtered with the SAM 2-based masking workflow. In addition, the synthetic vessel appearance may not fully match the optical surface characteristics of the real phantom, including vessel-surface contrast and the apparent size or visibility of the surface beads, so reconstruction behavior in the synthetic setting may differ from performance on benchtop footage. The stress proxy is derived from reconstructed surface geometry and should therefore be interpreted as a relative surface-based measure rather than an absolute wall-stress estimate. It does not resolve through-wall mechanics, lumen-facing stresses, or stress sign, and the edge-based Neo-Hookean mapping is a simplified constitutive surrogate rather than a full continuum model. More direct estimation of device-, clot-, and vessel-level internal mechanics would require complementary continuum or coupled in-silico modeling approaches [35, 36]. Experiments were performed in silicone phantoms, which are valuable for controlled testing but do not fully reproduce in vivo tissue properties, artery-clot-device interaction, or physiological boundary conditions [3, 4, 35].

Future work should therefore include repeated trials across a broader range of anatomies, clot analogs, and device configurations, together with cross-validation against complementary experimental or computational models. Methodologically, an important next step would be to move from the current fixed post hoc edge-graph representation toward more explicit surface-aware dynamic Gaussian formulations. Static work such as SuGaR motivates surface-aligned mesh extraction from Gaussian representations, while more recent dynamic methods such as DynaSurfGS, DG-Mesh, and GSTAR suggest possible paths toward temporally consistent surface reconstruction and tracking directly within Gaussian-based pipelines [37, 38, 39, 40]. If future extensions move beyond comparative kinematics toward coupled reconstruction-and-simulation workflows, mesh-coupled Gaussian methods such as MaGS may also be relevant [41]. On the segmentation side, beyond the current SAM 2 workflow, future work could evaluate motion-aware and memory-based video segmentation approaches such as SAMURAI, Cutie, and DEVA to reduce manual correction and improve temporal consistency under self-occlusion, distractors, or longer sequences; concept-promptable models such as SAM 3 may also become useful if later datasets require more flexible open-vocabulary prompting [42, 43, 44, 45].

5 Conclusion

This work introduces a standardized and accessible protocol for comparative benchtop thrombectomy analysis by combining low-cost multi-view acquisition, dynamic 4D reconstruction, fixed-connectivity ROI tracking, and relative surface-based stress-proxy mapping. The main novelty is the integration of these components with synthetic ground-truth validation and a clinically relevant two-condition benchtop demonstration in a single workflow, which together provide a coherent acquisition-to-analysis framework.

The framework provides practical, time-resolved regional metrics that can support method development and hypothesis generation for thrombectomy-device and access-strategy testing, while maintaining clear separation from absolute biomechanical wall-stress inference. Future work will focus on repeated trials per condition, expanded aspiration-catheter and pullback-condition sampling, broader anatomical variability, and improved reconstruction robustness through better camera coverage, reduced self-occlusion, and stronger temporal regularization.

Acknowledgments

This work was supported by the Elegant Mind Collaboration, the UCLA Department of Physics and Astronomy, and the Kaneko Lab at UCLA Health.

Author contributions (CRediT).

Ethan Nguyen: Conceptualization, Methodology, Software, Resources, Data Curation, Supervision, Formal Analysis, Investigation, Validation, Visualization, Project Administration, Writing – original draft, Writing – review & editing.

Javier Carmona: Conceptualization, Methodology, Software, Investigation, Supervision, Writing – review & editing.

Arisa Matsuzaki: Investigation, Validation, Resources, Writing – review & editing.

Naoki Kaneko: Conceptualization, Resources, Funding Acquisition, Supervision, Writing – review & editing.

Katsushi Arisaka: Conceptualization, Funding Acquisition, Supervision, Writing – review & editing.

All authors reviewed and approved the final manuscript.

The authors thank members of the Elegant Mind Collaboration for their support. Soumya Bukkapatnam assisted with manual placement of fluorescent beads on the vessel models. Charlie Long designed the initial iterations of the camera mounts. Thomas Leung and Martin Leung constructed the dodecahedron camera rig. Mingda He assisted with code development.

The authors also thank members of the Kaneko Lab. Dr. Eisuke Tsukagoshi and Lea Guo provided the silicone vessel models. Lea Guo and Mahsa Ghovvati also assisted with equipment procurement.

The authors further thank Dr. Shusaku Goto (visiting from Nagoya Hospital), Dr. Natsuhi Sasaki (research scholar at UCLA), and Dr. Eisuke Tsukagoshi for conducting the physical thrombectomy experiments.

Preprint Notes

Version: Version 1 April 2026.

Preprint status: This manuscript is a preprint and has not yet undergone peer review. It is intended for community feedback.

Funding: This work was supported by UCLA Physics & Astronomy / UCLA Health.

Ethics statement: This study did not involve human subjects or animal experiments.

Data availability: Code and example data are publicly available at

<https://ethanuser.github.io/vessel4D>. Additional benchtop data are available from the corresponding authors upon reasonable request.

Conflicts of interest: The authors declare no conflicts of interest.

References

- [1] T. Ishiguro, Y. Omura, Y. Takano, T. Arai, K. Karagiozov, K. Fukuda, Y. Kikuta, N. Momozaki, H. Eguchi, M. Nishitani, M. Sakaguchi, M. Uchida, M. Suzuki, T. Kawamata, Low radial and axial force stent retriever reduces symptomatic subarachnoid hemorrhage after mechanical thrombectomy for acute middle cerebral artery and medium vessel occlusion: A prospective pilot study, *Frontiers in Neurology* 16 (2025) 1723287. doi:10.3389/fneur.2025.1723287. 1
- [2] A. A. Dmytriw, B. Musmar, H. Salim, S. Ghozy, J. E. Siegler, H. Kobeissi, H. Shaikh, J. Khalife, M. Abdalkader, P. Klein, T. N. Nguyen, J. J. Heit, R. W. Regenhardt, N. M. Cancelliere, J. D. Bernstock, K. E. Naamani, A. Amllay, L. Meyer, A. Dusart, F. Bellante, G. Forestier, A. Rouchaud, S. Saleme, C. Mounayer, J. Fiehler, A. L. Kühn, A. S. Puri, C. Dyzmann, P. T. Kan, M. Colasurdo, G. Marnat, J. Berge, X. Barreau, I. Sibon, S. Nedelcu, N. Henninger, T. R. Marotta, C. J. Stapleton, J. D. Rabinov, T. Ota, S. Dofuku, L. L. Yeo, B. Y. Tan, A. Gopinathan, J. C. Martinez-Gutierrez, S. Salazar-Marioni, S. Sheth, L. Renieri, C. Capirossi, A. Mowla, L. Chervak, A. Vagal, N. Adeeb, H. H. Cuellar-Saenz, S. I. Tjoumakaris, P. Jabbour, P. Khandelwal, A. Biswas, F. Clarençon, M. Elhorany, K. Premat, I. Valente, A. Pedicelli, J. P. Filipe, R. Varela, M. Quintero-Consuegra, N. R. Gonzalez, M. A. Möhlenbruch, J. Jesser, V. Costalat, A. Ter Schiphorst, V. Yedavalli, P. Harker, Y. Aziz, B. Gory, C. P. Stracke, C. Hecker, R. Kadirvel, M. Killer-Oberpfalzer, C. J. Griessenauer, A. J. Thomas, C.-Y. Hsieh, D. S. Liebeskind, R. Alexandru Radu, A. M. Alexandre, I. Tancredi, T. D. Faizy, R. Fahed, C. Weyland, B. Lubicz, A. B. Patel, V. M. Pereira, A. Guenego, MAD-MT Consortium, Incidence and clinical outcomes of perforations during mechanical thrombectomy for medium vessel occlusion in acute ischemic stroke: A retrospective, multicenter, and multinational study, *European Stroke Journal* 9 (2024) 328–337. doi:10.1177/23969873231219412. 1
- [3] Y. Liu, M. Abbasi, J. L. Arturo Larco, R. Kadirvel, D. F. Kallmes, W. Brinjikji, L. Savastano, Preclinical testing platforms for mechanical thrombectomy in stroke: A review on phantoms, in-vivo animal, and cadaveric models, *Journal of Neurointerventional Surgery* 13 (2021) 816–822. doi:10.1136/neurintsurg-2020-017133. 1, 12
- [4] S. Johnson, A. Dwivedi, M. Mirza, R. McCarthy, M. Gilvarry, A Review of the Advancements in the in-vitro Modelling of Acute Ischemic Stroke and Its Treatment, *Frontiers in Medical Technology* 4 (Jun. 2022). doi:10.3389/fmedt.2022.879074. 1, 12
- [5] D. A. Poulos, M. T. Froehler, B. C. Good, Investigation of stent retriever removal forces in an experimental model of acute ischemic stroke, *Frontiers in Neurology* 15 (Oct. 2024). doi:10.3389/fneur.2024.1486738. 1, 12
- [6] M. Nagargoje, The role of friction forces in arterial mechanical thrombectomy: A review, *Journal of Biomechanics* 192 (2025) 112966. doi:10.1016/j.jbiomech.2025.112966. 1, 12
- [7] M. Ernst, F. Sommer, M. Bartl, C. H. Riedel, P. Langer, Making the invisible visible: Integrated visualization and automated quantification of thrombus deformation during mechanical thrombectomy, *Frontiers in Medical Technology* 8 (Jan. 2026). doi:10.3389/fmedt.2026.1704010. 1
- [8] D. Solav, K. M. Moerman, A. M. Jaeger, K. Genovese, H. M. Herr, MultiDIC: An Open-Source Toolbox for Multi-View 3D Digital Image Correlation, *IEEE Access* 6 (2018) 30520–30535. doi:10.1109/ACCESS.2018.2843725. 2
- [9] M. Palanca, G. Tozzi, L. Cristofolini, The use of digital image correlation in the biomechanical area: A review, *International Biomechanics* 3 (2016) 1–21. doi:10.1080/23335432.2015.1117395. 2
- [10] A. C. Sanchez, D.-K. Kim, Assessment of the Film-Free Water Decal Method for Speckle Pattern Application in Digital Image Correlation, *Sensors (Basel, Switzerland)* 24 (2024) 5657. doi:10.3390/s24175657. 2
- [11] A. Mustafa, M. Volino, H. Kim, J.-Y. Guillemaut, A. Hilton, Temporally Coherent General Dynamic Scene Reconstruction, *International Journal of Computer Vision* 129 (2021) 123–141. doi:10.1007/s11263-020-01367-2. 2
- [12] G. Wu, T. Yi, J. Fang, L. Xie, X. Zhang, W. Wei, W. Liu, Q. Tian, X. Wang, 4D Gaussian Splatting for Real-Time Dynamic Scene Rendering (Jul. 2024). doi:10.48550/arXiv.2310.08528. 2, 4
- [13] H. Kortman, E. Smit, M. Oei, R. Manniesing, M. Prokop, F. Meijer, 4D-CTA in Neurovascular Disease: A Review, *AJNR: American Journal of Neuroradiology* 36 (2015) 1026–1033. doi:10.3174/ajnr.A4162. 2

-
- [14] N. Kaneko, T. Mashiko, T. Ohnishi, M. Ohta, K. Namba, E. Watanabe, K. Kawai, Manufacture of patient-specific vascular replicas for endovascular simulation using fast, low-cost method, *Scientific Reports* 6 (2016) 39168. doi:10.1038/srep39168. 2
- [15] Z. Zhang, A flexible new technique for camera calibration, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (2000) 1330–1334. doi:10.1109/34.888718. 3
- [16] G. Bradski, The OpenCV library, *Dr. Dobb's Journal of Software Tools* (2000). 3
- [17] N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. Ryali, T. Ma, H. Khedr, R. Rädle, C. Rolland, L. Gustafson, E. Mintun, J. Pan, K. V. Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. Dollár, C. Feichtenhofer, SAM 2: Segment Anything in Images and Videos (Oct. 2024). doi:10.48550/arXiv.2408.00714. 3
- [18] J. MacQueen, Some methods for classification and analysis of multivariate observations, in: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, Vol. 5.1, University of California Press, 1967, pp. 281–298. 4
- [19] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in: *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96*, AAAI Press, Portland, Oregon, 1996, pp. 226–231. 4
- [20] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, A. Müller, J. Nothman, G. Louppe, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, É. Duchesnay, *Scikit-learn: Machine Learning in Python*, https://arxiv.org/abs/1201.0490v4 (Jan. 2012). 4
- [21] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, *SciPy 1.0: Fundamental algorithms for scientific computing in Python*, *Nature Methods* 17 (2020) 261–272. doi:10.1038/s41592-019-0686-2. 4
- [22] G. Taubin, A signal processing approach to fair surface design, in: *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '95*, Association for Computing Machinery, New York, NY, USA, 1995, pp. 351–358. doi:10.1145/218380.218473. 4
- [23] P. Perona, J. Malik, Scale-space and edge detection using anisotropic diffusion, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12 (1990) 629–639. doi:10.1109/34.56205. 4
- [24] M. Black, G. Sapiro, D. Marimont, D. Heeger, Robust anisotropic diffusion, *IEEE Transactions on Image Processing* 7 (1998) 421–432. doi:10.1109/83.661192. 4
- [25] M. J. Black, P. Anandan, The Robust Estimation of Multiple Motions: Parametric and Piecewise-Smooth Flow Fields, *Computer Vision and Image Understanding* 63 (1996) 75–104. doi:10.1006/cviu.1996.0006. 4
- [26] P. J. Huber, Robust Estimation of a Location Parameter, *The Annals of Mathematical Statistics* 35 (1964) 73–101. doi:10.1214/aoms/1177703732. 5
- [27] P. J. Rousseeuw, C. Croux, Alternatives to the Median Absolute Deviation, *Journal of the American Statistical Association* 88 (1993) 1273–1283. doi:10.1080/01621459.1993.10476408. 5
- [28] G. A. Holzapfel, *Nonlinear Solid Mechanics: A Continuum Approach for Engineering*, Kluwer Academic Publishers Dordrecht, 2000. 5
- [29] J. Y. Chueh, A. K. Wakhloo, M. J. Gounis, Neurovascular Modeling: Small-Batch Manufacturing of Silicone Vascular Replicas, *American Journal of Neuroradiology* 30 (2009) 1159–1164. doi:10.3174/ajnr.A1543. 5
- [30] K. Gong, Z. Wen, M. Xu, W. He, Q. Wang, N. Zhang, Z. Li, C. Li, D. Lian, W. Zhao, X. He, M. Zhang, SWiT-4D: Sliding-Window Transformer for Lossless and Parameter-Free Temporal 4D Generation (Dec. 2025). doi:10.48550/arXiv.2512.10860. 7
- [31] J. M. Bland, D. G. Altman, Statistical methods for assessing agreement between two methods of clinical measurement, *Lancet* 1 (1986) 307–310. 8
- [32] Y. Nariai, T. Takigawa, A. Hyodo, K. Suzuki, Modification by an aspiration catheter for vessel stretching in thrombectomy using a stent retriever in vitro, *Journal of Stroke and Cerebrovascular Diseases* 32 (Mar. 2023). doi:10.1016/j.jstrokecerebrovasdis.2022.106948. 12
- [33] S. H. Baek, S. Kim, M. Kang, J.-H. Choi, H. J. Kwon, D. W. Kim, Effect of distal access catheter tip position on angiographic and clinical outcomes following thrombectomy using the combined stent-retriever and aspiration approach, *PLoS ONE* 16 (2021) e0252641. doi:10.1371/journal.pone.0252641. 12
- [34] K. Kawamoto, Y. Nagao, M. Naganuma, Y. Inatomi, Y. Hashimoto, T. Yonehara, M. Nakajima, Stent-retriever characteristics and strategies associated with recanalization in thrombectomy for acute ischemic stroke, *Clinical Neurology and Neurosurgery* 242 (2024) 108332. doi:10.1016/j.clineuro.2024.108332. 12
- [35] G. Luraghi, J. F. Rodriguez Matas, G. Dubini, F. Berti, S. Bridio, S. Duffy, A. Dwivedi, R. McCarthy, B. Fereidoonzhad, P. McGarry, C. B. L. M. Majoie, F. Miglia vacca, INSIST investigators,
-

-
- Applicability assessment of a stent-retriever thrombectomy finite-element model, *Interface Focus* 11 (2021) 20190123. [doi:10.1098/rsfs.2019.0123](https://doi.org/10.1098/rsfs.2019.0123). 12
- [36] S. M. Mousavi J S, D. Faghihi, K. Sommer, M. M. S. Bhurwani, T. R. Patel, B. Santo, M. Waqas, C. Ionita, E. I. Levy, A. H. Siddiqui, V. M. Tutino, Realistic computer modelling of stent retriever thrombectomy: A hybrid finite-element analysis-smoothed particle hydrodynamics model, *Journal of the Royal Society, Interface* 18 (2021) 20210583. [doi:10.1098/rsif.2021.0583](https://doi.org/10.1098/rsif.2021.0583). 12
- [37] A. Guédon, V. Lepetit, SuGaR: Surface-Aligned Gaussian Splatting for Efficient 3D Mesh Reconstruction and High-Quality Mesh Rendering (Dec. 2023). [doi:10.48550/arXiv.2311.12775](https://doi.org/10.48550/arXiv.2311.12775). 13
- [38] W. Cai, W. Ye, P. Ye, T. He, T. Chen, DynaSurfGS: Dynamic Surface Reconstruction with Planar-based Gaussian Splatting (Aug. 2024). [doi:10.48550/arXiv.2408.13972](https://doi.org/10.48550/arXiv.2408.13972). 13
- [39] I. Liu, H. Su, X. Wang, Dynamic Gaussians Mesh: Consistent Mesh Reconstruction from Dynamic Scenes (Mar. 2025). [doi:10.48550/arXiv.2404.12379](https://doi.org/10.48550/arXiv.2404.12379). 13
- [40] C. Zheng, L. Xue, J. Zarate, J. Song, GauSTAR: Gaussian Surface Tracking and Reconstruction, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025, pp. 16543–16553. 13
- [41] S. Ma, Y. Luo, W. Yang, Y. Yang, MaGS: Reconstructing and Simulating Dynamic 3D Objects with Mesh-adsorbed Gaussian Splatting (Nov. 2024). [doi:10.48550/arXiv.2406.01593](https://doi.org/10.48550/arXiv.2406.01593). 13
- [42] C.-Y. Yang, H.-W. Huang, W. Chai, Z. Jiang, J.-N. Hwang, SAMURAI: Adapting Segment Anything Model for Zero-Shot Visual Tracking with Motion-Aware Memory (Nov. 2024). [doi:10.48550/arXiv.2411.11922](https://doi.org/10.48550/arXiv.2411.11922). 13
- [43] H. K. Cheng, S. W. Oh, B. Price, J.-Y. Lee, A. Schwing, Putting the Object Back into Video Object Segmentation (Apr. 2024). [doi:10.48550/arXiv.2310.12982](https://doi.org/10.48550/arXiv.2310.12982). 13
- [44] H. K. Cheng, S. W. Oh, B. Price, A. Schwing, J.-Y. Lee, Tracking Anything with Decoupled Video Segmentation (Sep. 2023). [doi:10.48550/arXiv.2309.03903](https://doi.org/10.48550/arXiv.2309.03903). 13
- [45] N. Carion, L. Gustafson, Y.-T. Hu, S. Debnath, R. Hu, D. Suris, C. Ryali, K. V. Alwala, H. Khedr, A. Huang, J. Lei, T. Ma, B. Guo, A. Kalla, M. Marks, J. Greer, M. Wang, P. Sun, R. Rädle, T. Afouras, E. Mavroudi, K. Xu, T.-H. Wu, Y. Zhou, L. Momeni, R. Hazra, S. Ding, S. Vaze, F. Porcher, F. Li, S. Li, A. Kamath, H. K. Cheng, P. Dollár, N. Ravi, K. Saenko, P. Zhang, C. Feichtenhofer, SAM 3: Segment Anything with Concepts (Nov. 2025). [doi:10.48550/arXiv.2511.16719](https://doi.org/10.48550/arXiv.2511.16719). 13

Supplementary Information

This supplement provides overflow material that supports, but is not required to understand, reproduce, or review the main methods/validation framework.

S1 Supplementary Methods

This section retains only implementation details and supporting tables/media references that are not needed in the main manuscript.

S1.1 Acquisition implementation details

For reproducibility, the incoming OBS composite was recorded at 4080×4080 resolution using the NVENC H.264 encoder in .mkv format with YUY2 color, “High Quality, Medium File Size” recording quality, white balance set to approximately 2498, maximum exposure set to −11, and an OBS focus setting of approximately 500–525. A single UV lamp was positioned beneath the apparatus, and the camera/lighting settings were held fixed across experiments. No formal temporal synchronization measurement beyond visual inspection was performed; in practice, the light source was toggled before recording to confirm that all cameras were awake and actively responding.

S1.2 Segmentation and graph-construction implementation details

The 4080×4080 OBS composite was first split into nine 1360×1360 view videos, and SAM 2 was run independently on each view before the segmented outputs were merged back into the multi-view reconstruction workflow. Typical segmentation used approximately eight additive/subtractive prompt points per vessel. Common failure modes were view-specific occlusion of the vessel by parts of the 3D-printed support structure and reduced fluorescent-bead visibility near the vessel ends; manual re-prompting was uncommon (less than 20% of sequences), and no mask cleanup beyond background suppression was applied.

Before clustering and graph construction, primitives were removed if their estimated radius was below 0.07 mm, if opacity was below 0.05, or if their RGB standard deviation was below 0.05 (used to reject near-gray/background-like points). Reconstruction otherwise followed the default 4DGS settings summarized in the main manuscript. Typical runs retained on the order of ~70,000 filtered primitives, which were reduced to roughly 2,000–3,000 clusters after clustering and to roughly ~2,000 curated clusters after manual cleanup; the corresponding graph typically contained on the order of

50,000–60,000 edges. Formal clustering, graph, spatial-coherence, and metric definitions remain in the main manuscript.

S2 Synthetic Validation

This section adds condition-level metrics and brief implementation details that complement the main-text synthetic summary. Table S1 reports the expanded condition-level metrics. In Blender, the bulk control used a rigid translation, whereas the pulling conditions used a lattice deformation in which the main control point was displaced by n mm for the nominal “ n mm” pull.

Metric definitions follow the main manuscript: CD (Eq. 12); directed distances $d_{P \rightarrow G}$ and $d_{G \rightarrow P}$ (Eqs. 13 and 14); temporal disagreement ΔCD (Eq. 16); precision/recall (Eq. 18); and F-score (Eq. 19).

Table S1. Synthetic validation metrics across deformation conditions, including directed and symmetric Chamfer terms, temporal disagreement, and overlap statistics.

Condition	CD				ΔCD		τ (mm)	P	R	F
	$d_{P \rightarrow G}$ (mm)	$d_{G \rightarrow P}$ (mm)	Sym (mm)	CD_{norm}	mm	rel				
Bulk	0.535	1.345	1.881	3.152	1.52×10^{-2}	0.085	1	0.919	0.313	0.467
							2	0.983	0.874	0.925
							3	0.998	0.981	0.989
1 mm	0.465	1.349	1.815	3.041	4.75×10^{-4}	0.139	1	0.970	0.309	0.468
							2	0.998	0.863	0.926
							3	0.999	0.981	0.990
2 mm	0.474	1.279	1.752	2.937	6.82×10^{-4}	0.100	1	0.967	0.344	0.507
							2	0.995	0.890	0.940
							3	0.999	0.991	0.995
3 mm	0.462	1.276	1.738	2.912	1.16×10^{-3}	0.114	1	0.964	0.348	0.512
							2	0.994	0.891	0.940
							3	0.998	0.989	0.994
4 mm	0.465	1.279	1.745	2.924	1.37×10^{-3}	0.100	1	0.972	0.346	0.510
							2	0.997	0.889	0.940
							3	0.998	0.990	0.994
5 mm	0.465	1.249	1.714	2.872	1.79×10^{-3}	0.105	1	0.968	0.356	0.521
							2	0.995	0.905	0.948
							3	0.999	0.993	0.996

At the default $\tau = 1$ mm threshold, recall is lower than precision because the reconstruction is summarized as a clustered point cloud that is intentionally sparser than the dense GT sampling. In other words, most reconstructed points still lie near the GT surface, but many GT points do not have a reconstructed point within the same tolerance because clustering reduces point density to improve temporal stability and suppress local Gaussian drift/noise. This asymmetry is why precision remains high while recall and F-score are more conservative in Table S1.

S3 Additional Quantitative Results

Table S2 reports regional median displacement and stress values across sweep settings for ground truth and reconstructed estimates, using the ROI displacement/stress definitions in Sec. 2.1.11 of the main manuscript (including Eqs. 8–11). Percent error is intentionally omitted here, because small ground-truth magnitudes can inflate relative percentage error and obscure practical agreement. Agreement is assessed primarily through Bland–Altman analysis and correlation trends in the main manuscript.

S3.1 Spatial coherence filtering ablation

To assess the effect of spatial coherence filtering (SCF), we repeated the synthetic ROI agreement analysis without SCF using the same 25 paired points (5 ROIs \times 5 pull magnitudes). Figure S1 shows the corresponding correlation and Bland–Altman plots, and Fig. S2 provides a qualitative stress-map comparison for the last frame (maximum deformation) of the 5 mm pulling condition. Without SCF, stray high-stress points are more visibly apparent in localized regions, whereas SCF suppresses these artifacts while preserving the broader stress pattern.

Relative to the SCF-enabled analysis, displacement agreement metrics remain similar, while stress-proxy agreement shows larger positive bias and wider limits of agreement without SCF. Specifically, stress-proxy Bland–Altman bias increases from 0.004 MPa to 0.010 MPa, and the upper

Table S2. Regional summary statistics (median) for displacement and stress across sweep settings. Agreement is evaluated primarily via Bland–Altman and correlation in the main manuscript.

Region	Method	Displacement (mm)					Stress (MPa)				
		1 mm	2 mm	3 mm	4 mm	5 mm	1 mm	2 mm	3 mm	4 mm	5 mm
R1	GT	1.023	2.046	3.068	4.091	5.114	0.005	0.009	0.014	0.018	0.023
	Ours	0.999	2.040	3.074	4.116	5.079	0.006	0.010	0.016	0.017	0.022
R2	GT	0.956	1.912	2.868	3.823	4.779	0.007	0.015	0.022	0.029	0.036
	Ours	0.922	1.925	2.791	3.731	4.750	0.009	0.015	0.025	0.037	0.034
R3	GT	0.302	0.604	0.906	1.208	1.510	0.010	0.021	0.031	0.042	0.052
	Ours	0.275	0.507	0.789	1.004	1.458	0.015	0.023	0.038	0.045	0.058
R4	GT	0.596	1.193	1.789	2.385	2.981	0.010	0.022	0.034	0.046	0.059
	Ours	0.639	1.114	1.847	2.416	2.999	0.016	0.026	0.040	0.052	0.064
R5	GT	0.821	1.642	2.463	3.284	4.105	0.008	0.016	0.024	0.033	0.041
	Ours	0.724	1.400	2.250	3.021	3.507	0.012	0.021	0.033	0.042	0.052




Figure S1. Agreement analyses without spatial coherence filtering (SCF) for max-median ROI metrics between reconstruction and ground truth: correlation and Bland–Altman for displacement magnitude (A–B, mm) and relative surface-based stress-proxy magnitude (C–D, MPa).

limit of agreement widens from 0.010 MPa to 0.016 MPa; qualitatively, the no-SCF maps also show more localized stray high-stress points (Fig. S2).

S4 Supplementary Figures

Table S3. Agreement summary with and without spatial coherence filtering (SCF) for max-median ROI synthetic validation metrics.

Metric	Condition	Slope	Intercept	R^2	BA Bias	Lower LoA	Upper LoA
Displacement (mm)	w/ SCF	0.987	-0.112	0.992	-0.142	-0.385	0.100
	w/o SCF	0.986	-0.108	0.991	-0.140	-0.388	0.107
Stress proxy (MPa)	w/ SCF	1.068	0.002	0.969	0.004	-0.002	0.010
	w/o SCF	1.083	0.008	0.964	0.010	0.003	0.016




Figure S2. Qualitative stress-proxy comparison with and without spatial coherence filtering (SCF) at the last frame (maximum deformation) of the 5 mm synthetic pulling condition. GT, reconstruction with SCF, and reconstruction without SCF are shown with a shared color scale. The zoomed view shows more visible stray high-stress points without SCF, whereas SCF suppresses these localized artifacts while preserving the overall stress pattern.




Figure S3. Nine-view chessboard calibration comparison between physical and synthetic setups. (A) Real-camera views from the nine-camera rig using the 6×8 chessboard used to estimate camera extrinsic parameters. (B) Corresponding views from the nine synthetic Blender cameras using the same pattern.

S5 Supplementary Videos and Data

Table S4 indexes all supplementary videos using labels S1–S11 and provides a short description for each item.

Table S4. Supplementary video index and short descriptions.

Label	Type	Condition	Short description	Link
S1	Benchtop	Raw	Multi-view raw benchtop footage of the silicone vessel experiment before downstream reconstruction and metric analysis.	Link
S2	Benchtop	Cervical displacement	Displacement visualization for the benchtop cervical aspiration-catheter placement condition.	Link
S3	Benchtop	Terminal displacement	Displacement visualization for the benchtop ICA-terminus aspiration-catheter placement condition.	Link
S4	Benchtop	Cervical stress proxy	Stress-proxy visualization for the benchtop cervical aspiration-catheter placement condition.	Link
S5	Benchtop	Terminal stress proxy	Stress-proxy visualization for the benchtop ICA-terminus aspiration-catheter placement condition.	Link
S6	Synthetic	Bulk raw	Raw rendered sequence for the synthetic bulk-translation control experiment.	Link
S7	Synthetic	Bulk displacement	Displacement visualization for the synthetic bulk-translation control experiment.	Link
S8	Synthetic	Bulk stress proxy	Stress-proxy visualization for the synthetic bulk-translation control experiment.	Link
S9	Synthetic	Pulling raw	Raw rendered sequence for the synthetic localized-pulling validation experiment.	Link
S10	Synthetic	Pulling displacement	Displacement visualization for the synthetic localized-pulling validation experiment.	Link
S11	Synthetic	Pulling stress proxy	Stress-proxy visualization for the synthetic localized-pulling validation experiment.	Link
