Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Geometry-Aware Abdominal Aortic Aneurysm Digital Twin for Patient-Specific Wall Stress Mapping

Julian Carvajal Rico
The University of Texas at San Antonio
Victor De Oliveira
The University of Texas at San Antonio
Satish C. Muluk
Allegheny Health Network
Mark K. Eskandari
Northwestern University
Vikram S. Kashyap
Corewell Health
Ender A. Finol
ender.finol@utsa.edu

The University of Texas at San Antonio

Article

Keywords: Abdominal Aortic Aneurysm, Wall Stress, Deep Learning, Graph Neural Networks, Transformers

Posted Date: July 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9971735/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: Competing interest reported. M.K.E. has a paid consulting role with W.L. Gore and Boston Scientific. The other authors declare that they have no competing interests.

Geometry-Aware Abdominal Aortic Aneurysm Digital Twin for Patient-Specific Wall Stress Mapping

Julian Carvajal Rico[1], Victor De Oliveira[2], Satish C. Muluk[3], Mark K. Eskandari[4], Vikram S. Kashyap[5], and Ender A. Finol[1,*]

[1]Department of Mechanical, Aerospace, and Industrial Engineering, The University of Texas at San Antonio, San Antonio, TX 78249, USA
[2]Department of Statistics and Data Science, The University of Texas at San Antonio, San Antonio, TX 78249, USA
[3]Department of Thoracic and Cardiovascular Surgery, Allegheny Health Network, Allegheny General Hospital, Pittsburgh, PA 15212, USA
[4]Feinberg School of Medicine, Northwestern University, Chicago, IL 60611, USA
[5]Cardiovascular Health, Corewell Health, Grand Rapids, MI 49503, USA
[*]Corresponding Author: ender.finol@utsa.edu

ABSTRACT

Accurate estimation of wall stress distributions in abdominal aortic aneurysms (AAAs) is critical for improving rupture risk prediction beyond the traditional maximum diameter criterion. While high-fidelity finite element analysis (FEA) provides precise patient-specific wall stress estimates, its computational cost limits its use in real-time clinical decision making. We introduce a graph-based deep learning framework that leverages three types of graph neural networks (GNNs), namely a Gated Graph Convolutional Network (GGCN), an Equivariant Graph Neural Network (EGNN), and a Graph Transformer (GT) to rapidly predict FEA-derived wall stress distributions directly from AAA outer wall surface meshes. Computed tomography angiography (CTA) images from 202 AAA patients treated at three clinical centers, were segmented using an in-house U-Net pipeline and refined manually to generate high quality AAA volume and surface meshes. From these, several node-specific geometric and biomechanical features were recorded, including wall thickness (WallTHK), intraluminal thrombus thickness (ILTTHK), wall strength (σ_{ult}), the principal minor and major curvatures (k_1 and k_2) and their derived curvatures (Gaussian K and mean M), distance to the lumen centerline (r_L), normalized local diameter (NORD), and the ILTTHK-to-WallTHK ratio (ρ^{ILT}). Local information was incorporated via neighborhood averaged features from the six nearest nodes. All GNNs achieved high fidelity in reproducing patient-specific FEA wall stress fields, with the GT providing the best overall agreement, while the GGCN and EGNN also delivered strong performances. The best performing GT captured spatial stress patterns, preserved biomechanical trends, and enabled near real-time inference. It reproduced FEA-derived wall stress with high accuracy (node-specific $R^2 > 0.97$, graph-level $R^2 > 0.98$) while reducing per-case inference time from hours for a complete FEA simulation to a few seconds on a single GPU. With further refinement and larger cohorts, this graph-based framework may strengthen rupture risk assessment and facilitate routine, patient-specific biomechanical analyses in clinical workflows.

Keywords: Abdominal Aortic Aneurysm; Wall Stress; Deep Learning; Graph Neural Networks; Transformers

1 Introduction

Abdominal aortic aneurysm (AAA) is a life-threatening vascular disease characterized by a permanent, localized dilation of the abdominal aorta. Although screening and surveillance have reduced the incidence of fatal ruptures, clinical decision making is still mainly based on the size of the aneurysm, in particular the maximum diameter and growth rate, to recommend elective repair[1]. In line with this practice, current medical guidelines recommend elective intervention at diameter thresholds of approximately ≥ 5.5 cm in men and 5.0 - 5.4 cm in women, with additional considerations for rapid expansion or symptoms[2,3]. However, small aneurysms can rupture, for example, at a rate of $\sim 23\%$ for AAAs between 4.1–5.0 cm[4]. The clinical management criteria are convenient, but only indirectly reflect the underlying biomechanics that precipitate rupture. The risk of rupture is governed by the balance between local wall stress and wall strength; regions of complex geometry, heterogeneous wall thickness, and intraluminal thrombus (ILT) can create highly nonuniform stress fields that diameter alone cannot identify[1,5,6].

FEA enables patient-specific estimation of wall stress, which has been repeatedly associated with improved risk stratification compared to diameter thresholds[7–9]. However, traditional FEA requires laborious preprocessing (segmentation, mesh generation, material model and wall thickness assumptions, specification of boundary conditions) and substantial computational time, which limits its use in routine practice[10]. In addition, uncertainty about material properties, the constitutive behavior of ILT, and

boundary conditions can propagate to stress predictions, complicating clinical translation[6,11–16]. Despite decades of research progress, biomechanical stress metrics have not been adopted for routine decision making in current clinical guidelines, and recent evidence highlights limited and heterogeneous data supporting their incremental value[3,17].

Advancements in machine learning (ML) and computer vision offer a complementary path: learn predictive mappings from image-derived anatomy to biomechanical stress directly from data. Robust automated segmentation with U-Net–style architectures has improved performance in reconstructing the AAA wall, lumen, and ILT[18–20]. Surface geometry encodes rich biomechanical information; features such as principal curvatures (minor k_1 and major k_2), mean (M) and Gaussian (K) curvatures, wall thickness (WallTHK), ILT thickness (ILTTHK), and centerline distance (r_L) correlate with wall stress concentrations and rupture-prone regions[21–25]. Because these descriptors are calculated on irregular vessel surfaces, graph neural networks (GNNs) provide a natural modeling framework by operating directly on the surface meshes as graphs, where message passing aggregates local and global geometric information[26–29].

Faithfully capturing local interactions on complex aneurysm geometries remains a challenge. Standard convolutional or purely local GNNs may struggle to model distant but biomechanically coupled regions (e.g., across the sac or near the neck), while models that rely heavily on absolute coordinates can be sensitive to orientation and scale. In addition, clinically significant high stress nodes are rare relative to low-stress nodes; therefore, loss functions and readouts that emphasize tail behavior are needed to avoid systematic underestimation of extremes. Recent GNN architectures such as gated graph convolutional networks (GGCNs), equivariant graph neural networks (EGNNs), and Graph Transformers (GTs) provide complementary strategies to address these limitations by strengthening message passing, enforcing geometric symmetries, and enabling global and local attention over complex vessel wall meshes. Within this family of GNNs, we first consider a non-equivariant GGCN inspired by Gated Graph Sequence Neural Networks[30]. In this architecture, information is propagated along the wall mesh via gated message passing, where recurrent units (e.g., GRU-style updates) regulate how newly aggregated neighborhood information is integrated with existing node states. This gating mechanism helps stabilize optimization on large, irregular graphs and mitigates the over-smoothing and vanishing-gradient issues that can arise in deeper message-passing networks. For AAA modeling, the GGCN provides an expressive yet relatively lightweight baseline for modeling local and mid-range wall stress patterns, with a graph-based formulation that can be applied inductively to patient-specific surface meshes of varying resolution.

EGNNs complement the baseline by explicitly encoding the underlying Euclidean symmetries of the geometry. By constraining the network to be equivariant to translations and rotations, EGNN-type architectures can reuse information more effectively for differently oriented aneurysms (e.g. due to CT scanner orientation or patient positioning) and reduce the reliance on arbitrary coordinate frames[31–33]. This is particularly desirable when wall stress depends on shape and relative configuration rather than global positioning and when training data are limited or heterogeneous between imaging centers. In practice, this symmetry-aware inductive bias allows EGNNs to achieve strong performance with relatively modest capacity, making them a computationally efficient alternative to more complex attention-based architectures in 3D medical imaging settings. GTs extend self-attention to irregular domains by conditioning attention on graph structure and geometry, enabling information flow across distant but biomechanically coupled regions. Relative and positional encodings are key to robustness and long-range reasoning in attention[34], while graph-specialized Transformer designs (e.g., spectral/Laplacian positional encodings and edge-aware keys/values) have shown strong performance in non-Euclidean domains[35,36]. Moreover, recent GTs that incorporate distance and shortest-path signals achieve state-of-the-art accuracy on structural benchmarks[37]. For applications where strict orientation consistency is critical beyond the present work, equivariant attention architectures provide an alternative route, preserving prediction behavior under rigid motion while retaining the benefits of attention[38]. For AAA modeling, this attention-based formulation provides global receptive fields for modeling sac-to-neck and sac-to-branch interactions, enables orientation-insensitive reasoning through relative encodings, and remains compatible with large graphs via sparse attention over the mesh connectivity.

Informed by the aforementioned literature gaps, we view each patient-specific reconstruction as a geometry-aware *AAA digital twin* and seek fast, physically informed surrogates that map from surface meshes generated from computed tomography angiography (CTA) images to wall stress endpoints. Rather than treating peak, mean, and full-field stresses as separate modeling stages, we train node-specific surrogates that recover the spatial stress field on the aneurysm wall surface and then derive clinically salient global summaries, such as the mean wall stress ($\bar{\sigma}$) and the 99$^{\text{th}}$ percentile wall stress (σ_{99})[7,8]. With this formulation, an accurate reproduction of every nodal value is not required. Rather, the models are optimized to adequately reproduce the general stress distribution and high stress tail so that $\bar{\sigma}$ and σ_{99} calculated from the predicted field closely match their FEA-estimated counterparts, supporting the assessment of individual rupture risk.

The present work describes a reproducible AAA wall stress *digital twin* pipeline and systematically evaluates three GNN architectures for graph- and node-specific stress targets by contributing:

1. A geometry *digital twin* pipeline that converts CTA images into standardized graph representations of the AAA outer wall with node-specific geometric and biomechanical descriptors, and spatial location.

2. A systematic benchmark of GGCN, EGNN, and GT models in an AAA cohort and preprocessing, clarifying which graph paradigms are most suited to estimate $\bar{\sigma}$, σ_{99}, and full-field wall stress endpoints in this cohort.

3. Tail-aware training and readout strategies that enhance sensitivity to clinically important high stress regions without sacrificing overall accuracy.

4. A comprehensive evaluation on a multi-institution cohort of AAA CTA–derived geometries, reporting graph-level errors for $\bar{\sigma}$ and σ_{99}, node-specific accuracy and spatial correlations, and agreement of predicted vs. ground-truth FEA wall stress maps.

Framing wall stress estimation as a geometry-aware *digital twin* enables fast patient-specific wall stress mapping once the models are trained, improving scalability for large cohorts and longitudinal AAA surveillance. The same framework could be applied to other vascular pathologies that require biomechanical estimates and where the vasculature can be represented by a surface mesh with rapid, data-driven surrogates needed to complement traditional finite element modeling workflows. The remainder of the paper is organized as follows. Section 2 describes the generation of the data set and the computation of the features. Section 3 introduces graph-based surrogate models, including graph representation, model families, and training procedures. Section 4 presents quantitative and qualitative results for node-specific and graph-level wall stress endpoints. Section 5 discusses the clinical implications, limitations, and future directions for integrating *digital twin* wall stress mapping into clinical workflows, and Section 6 presents the conclusions.

2 Methods

We propose a geometry-aware *digital twin* workflow that links CTA-derived anatomy to patient-specific wall stress endpoints. As shown schematically in Figure 1, the pipeline integrates: (i) image segmentation to obtain binary masks for the wall, ILT, and lumen; (ii) standardized surface reconstruction for feature computations; (iii) volumetric meshing and FEA for biomechanical ground truth; and (iv) supervised graph learning for wall stress prediction. From each CTA exam, an in-house U-Net code with expert refinement produces masks that serve two functions. First, to reconstruct the outer wall surface, which we standardize its resolution for all patients, and compute node-specific geometric/biomechanical descriptors, including WallTHK, ILTTHK, k_1, k_2, M, K, r_L, the ILTTHK-to-WallTHK ratio (ρ^{ILT}), wall strength (σ_{ult}), and mesh coordinates (x,y,z) relative to the centroid (0,0,0) of the distal end of each AAA surface mesh. Second, to generate a volume mesh and compute first principal wall stresses in Abaqus (Dassault Systèmes, Waltham, MA), projecting the results back to the outer wall surface as the ground truth. We represent every AAA in the cohort as a surface graph whose nodes carry geometric/biomechanical features and wall stress targets, and we train three complementary GNN families: a GGCN, an EGNN, and a GT to estimate $\bar{\sigma}$, σ_{99}, and the full wall stress field. Training uses standardized features, early stopping, and tail-aware objectives while evaluation reports graph- and node-specific errors and spatial agreement with the ground truth FEA output as the reference.

2.1 Cohort, Imaging, Segmentation, and Surface Standardization

We retrospectively acquired CTA exams from 202 AAA patients under surveillance following IRB-approved protocols at three clinical centers (Allegheny General Hospital in Pittsburgh, PA; Northwestern Memorial Hospital in Chicago, IL; Corewell Health in Grand Rapids, MI). From each exam, we used images of the abdominal aorta starting immediately below the left renal artery and ending at the aorto-iliac bifurcation. Data set acquisition was conducted in accordance with the Declaration of Helsinki, and informed consent was waived because of the retrospective design and use of existing de-identified clinical data.

The CTA images were processed with an in-house U-Net based algorithm[20] that identified the outer wall, inner wall, and lumen boundaries, followed by expert refinement to correct artifacts and ensure anatomical fidelity, yielding four-region binary masks per CTA exam. (Figure 1.a). From the refined masks, we reconstructed triangulated surfaces for the outer wall and lumen using standard smoothing routines, preserving geometric detail relevant for wall stress calculation. The binary masks also served as input to the meshing pipeline that generated the surface and volume meshes used in the subsequent modeling calculations.

To make graph statistics comparable between patients and to allow for consistent batching, we standardized the outer wall surface resolution to a target of $\approx 150,000$ nodes. The target size was derived by averaging node counts among AAAs with maximum diameter ≤ 7.7 cm, which avoids anchoring to extremes. We performed isotropic, curvature-preserving remeshing using a PyMeshLab-based wrapper, with a short calibration pass to estimate the relation between face count and edge length, followed by a final pass with up to two bounded refinements to reach a tolerance of $\pm 2\%$ in node count relative to the target. For each AAA we registered pre- and post-remeshing node and face counts, calibration settings, final error to target, number of refinement iterations, and key clinical metadata such as $\bar{\sigma}$ and σ_{99}, serving as quality control records. Visual inspection confirmed preservation of sac morphology, neck geometry, and local bulges, while removing redundant facets that can destabilize feature estimates and message passing.

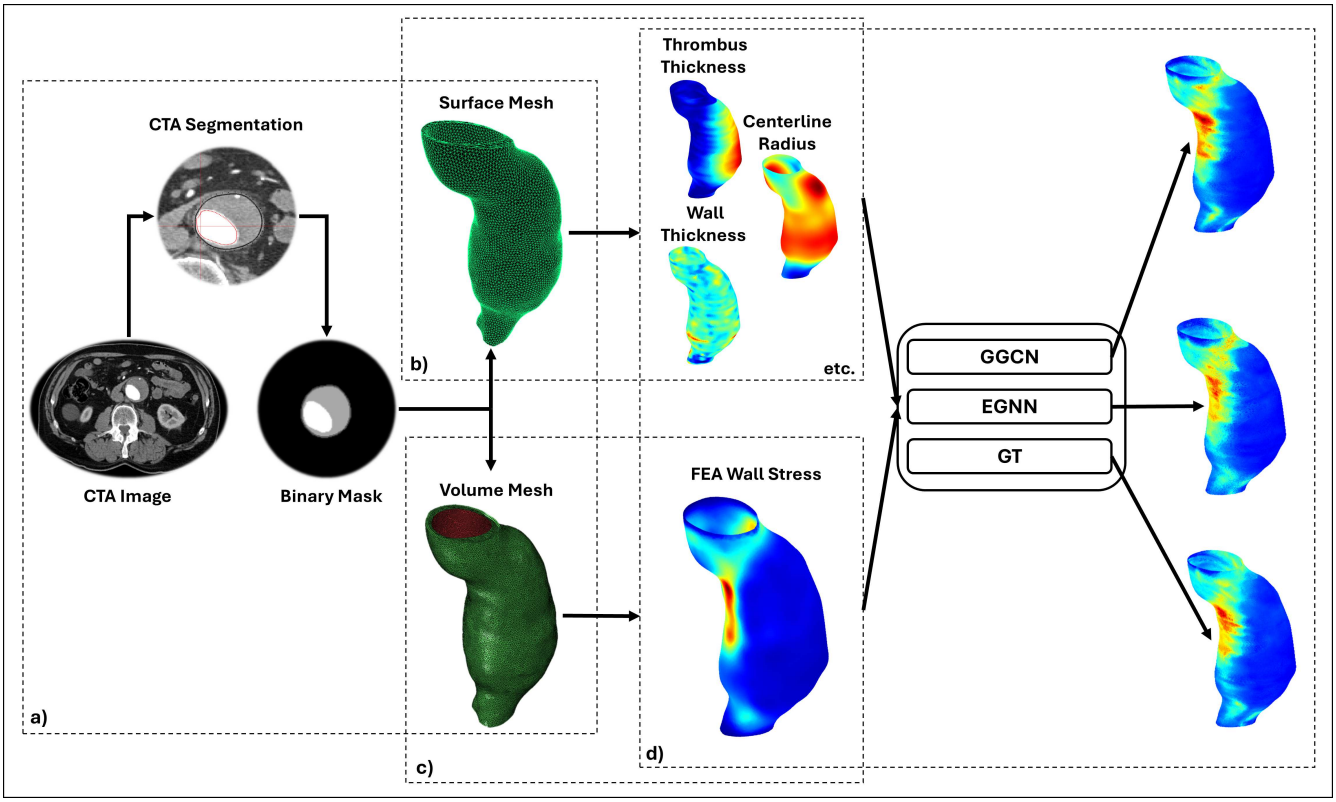


Figure 1. Geometry-aware *digital twin* framework for AAA wall stress mapping. (a) Segmentation from CTA imaging to binary masks, followed by reconstruction of surface and volume meshes. (b) Feature engineering on the standardized outer wall surface (e.g., WallTHK, ILTTHK, and r_L). (c) Biomechanical ground truth: first principal wall stress computed with the volume mesh via FEA and projected onto the outer wall surface. (d) Graph-based prediction of wall stress using GGCN, EGNN, and GT.

2.2 Geometric and Biomechanical Descriptors

From the standardized outer wall, inner wall, and lumen surfaces, we computed node-specific descriptors that represent thicknesses, thrombus burden, surface shape, and position. These quantities summarize geometric measures that modulate wall stress and the likelihood of focal stress amplification. For each outer wall vertex at position $\mathbf{x}_i$, its wall thickness is calculated as the minimum Euclidean distance to the inner wall surface vertex at position y_i, as given by Eq. (1),

$$\mathrm{WallTHK}_i = \min_{\mathbf{y}\in S_{\mathrm{inner}}} \left\|\mathbf{x}_i - \mathbf{y}\right\|_2. \tag{1}$$

ILTTHK is calculated as the minimum Euclidean distance from the outer wall vertex at position x_i to the lumen surface vertex at position z_i, calculated using Eq. (2), providing a surrogate for local thrombus burden.

$$\mathrm{ILTTHK}_i = \min_{\mathbf{z}\in S_{\mathrm{lumen}}} \left\|\mathbf{x}_i - \mathbf{z}\right\|_2. \tag{2}$$

The principal curvatures $k_{1,i} \geq k_{2,i}$ are the eigenvalues of the shape operator at the surface point $\mathbf{x}_i$, in the tangent plane, quantifying maximal and minimal normal bending along orthogonal directions[39,40]. From these, K and M are defined by Eqs. (3) and (4). A positive K indicates dome-like regions; a negative K indicates saddle-shaped regions. Curvatures were estimated with a modified biquintic finite element (BQFE)[41] patch fit to reduce discretization noise.

$$K_i = k_{1,i} * k_{2,i} \tag{3}$$

$$M_i = \tfrac{1}{2}\left(k_{1,i} + k_{2,i}\right). \tag{4}$$

r_L encodes local caliber and sac-to-neck variation. It is calculated as the Euclidean distance from each outer wall vertex at position $\mathbf{x}_i$ to the anatomically derived aortic centerline $\mathbf{c}(s)$, obtained by lumen skeletonization and smoothing[42,43], using Eq.

(5).

$$r_{L,i} = \min_{s} \left\| \mathbf{x}_i - \mathbf{c}(s) \right\|_2. \tag{5}$$

ρ^{ILT} represents the local thrombus burden relative to wall thickness and is used as a dimensionless indicator with a small ε for numerical stability, calculated using Eq. (6).

$$\rho_i^{\mathrm{ILT}} = \frac{\mathrm{ILTTHK}_i}{\mathrm{WallTHK}_i + \varepsilon}. \tag{6}$$

A wall strength surrogate was calculated at each node using a regression model[44] that relates local thrombus burden, normalized diameter, aneurysm familial history, and sex to the ultimate wall strength. For node i, the wall strength is calculated in units of $\mathrm{N/cm^2}$ using Eq. (7).

$$\sigma_{\mathrm{ult},i} = 72.9 - 33.5\left(\sqrt{\mathrm{ILT}_i} - 0.79\right) - 12.3\left(\mathrm{NORD}_i - 2.31\right) - 24\,\mathrm{HIST} + 15\,\mathrm{SEX}, \tag{7}$$

where ILT_i is the local ILTTHK (in cm), NORD_i is the normalized local diameter (in cm), HIST encodes the AAA family history (+0.5 for positive history and -0.5 otherwise), and SEX encodes sex as +0.5 for male or -0.5 for female patients.

The spatial coordinates (x, y, z) are centered and scaled in a common patient space. During learning, models use standardized coordinates or relative encodings such as $\Delta\mathbf{x}$, pairwise distances, or low-band Fourier features to promote invariance to translation and rotation.

2.3 Biomechanical Ground Truth (FEA)

Patient-specific wall stresses were computed from a lumen–thrombus–wall assembly using a standardized finite-element workflow that constructs a continuous ILT between the luminal and outer wall boundaries. From the refined binary masks, the outer wall and lumen surfaces were first meshed independently in `AAAMesh`[45] to generate NASTRAN shell meshes with nominal constant wall thickness and case-specific in-plane resolution; the outer wall was also exported as an STL surface for further processing. The lumen and inner wall boundaries were then used to derive a continuous ILT volume, smoothing and remeshing the resulting volume.

The lumen–thrombus–wall assembly was created with a multi-tool pipeline to ensure geometric continuity and mesh quality (Figure 2). The lumen and inner wall surfaces extracted from `AAAMesh` were subjected to a series of meshing and smoothing operations through Ansys ICEM CFD (Ansys Inc., Canonsburg, PA), MeshLab (Visual Computing Lab, ISTI-CNR, Italy) SolidWorks (Dassault Systèmes, Waltham, MA), and HyperMesh (Altair Engineering Inc., Frisco, TX). With the latter application, surface triangles were isotropically refined and converted to high-quality tetrahedrals (target $\approx 250,000 - 500,000$ elements).

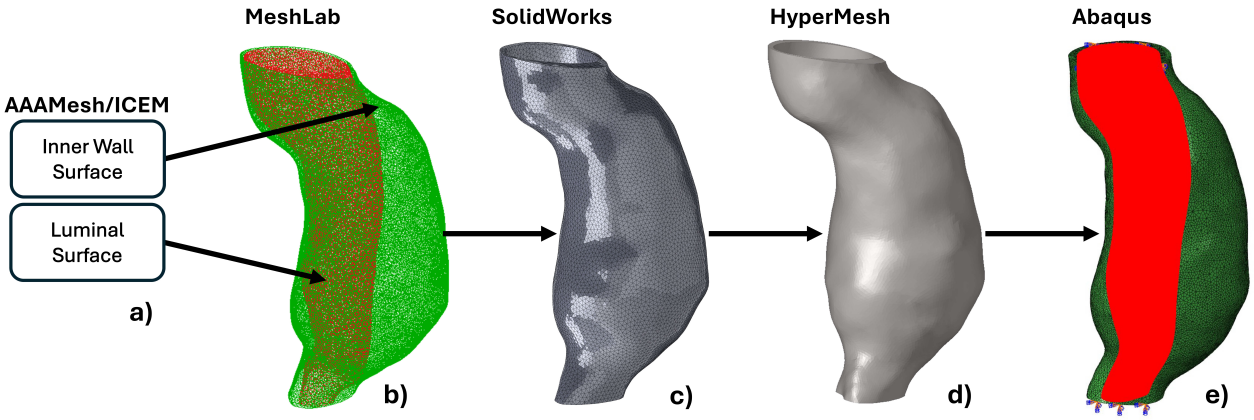


Figure 2. Schematic of the pipeline used to generate wall-ILT volume meshes for biomechanical analyses. (a) Extraction of the lumen and inner wall surfaces from AAAMesh[45] and Ansys ICEM CFD; (b) Curvature-preserving remeshing and smoothing in MeshLab; (c) Conversion to a closed solid in SolidWorks; (d) Tetrahedral volume meshing in HyperMesh; and (e) Wall–ILT assembly with luminal surface imported into Abaqus for FEA simulation.

For the FEA simulations, the AAA wall was modeled as an isotropic hyperelastic Mooney–Rivlin material (reduced polynomial, order 2) with coefficients $C_{10} = 17.4\ \mathrm{N/cm^2}$ and $C_{20} = 188.1\ \mathrm{N/cm^2}$ [6], and the ILT was modeled as an isotropic

Mooney–Rivlin material with $C_{10} = 7.98$ N/cm^2 and $C_{20} = 8.71$ N/cm^2 [46]. A surface tie constraint was applied to the ILT-wall interface to prevent relative sliding. An mean intraluminal pressure of 93.33 mmHg was uniformly applied to the inner ILT and inner wall surfaces, while the proximal and distal ends were constrained to suppress rigid-body motion. Hybrid solid elements (C3D4H) were used for the ILT and 20-node hybrid brick elements (C3D20H) for the wall. Geometric nonlinearity was enabled and automatic time increments were used for quasi-static equilibrium.

After convergence, the first principal stress at the wall integration points was exported and mapped to the standardized outer wall surface mesh used for graph construction. Node-specific stresses $\{\sigma_i\}_{i=1}^{N}$ were obtained by nearest-point projection from surface nodes to the underlying volumetric elements, with a distance criterion to exclude ill-mapped nodes from subsequent statistics. From this mapped field, we retained node-specific stresses as supervision targets for the GNNs and derived graph-level scalars used in the evaluation, in particular the empirical 99$^{\text{th}}$ percentile wall stress as given by Eq. (8),

$$\sigma_{99} \;=\; \inf\left\{ s \in \mathbb{R} \;:\; \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{\sigma_i \le s\} \ge 0.99 \right\}, \tag{8}$$

which serves as a robust surrogate for peak wall loading. $\bar{\sigma}$ and σ_{99} provide the biomechanical ground truth for training and evaluating node-specific wall stress maps and their clinically relevant global summaries.

3 Graph-Based Surrogate Models

To leverage GNNs for wall stress prediction, we view each AAA surface as a graph, so that local geometry and tissue properties become node-specific features, and the model can learn wall stress directly on the native surface mesh. This representation preserves the mesh connectivity while enabling message passing between neighboring wall regions.

3.1 Graph Representation

Each AAA geometry is represented as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of vertices (graph nodes) of the standardized outer wall surface (one node per mesh vertex) and $\mathcal{E}$ is the set of undirected edges inherited from the surface mesh connectivity. The node feature vector $\mathbf{h}_i$ for vertex i comprises the geometric descriptors detailed in Section 2.2, i.e. WallTHK, ILTTHK, k_1, k_2, M, K, r_L, ρ^{ILT}, and the spatial coordinates (x, y, z). To model short-range information strongly correlated with local wall stress, we augment these features with statistics calculated with the six nearest neighbors (by Euclidean distance). Specifically, for ILTTHK, ρ^{ILT}, k_1, k_2, σ_{ult}, and r_L, we compute the mean, standard deviation, maximum, and delta (max–min), generating descriptors based on the six closest spatial neighbors (kNN-6), e.g. $\text{ILTTHK}_{\text{mean}}^{\text{knn6}}$, $k_{1\,\text{std}}^{\text{kNN-6}}$, and $\text{r}_{\text{Ldelta}}^{\text{kNN-6}}$. Together with a categorical ILT class and simple products (e.g. $k_1 \times$ WallTHK), the augmentation yields 43 node-specific features per $\mathcal{V}$ in the processed data set, some of which are illustrated for an exemplary AAA in Figure 3.

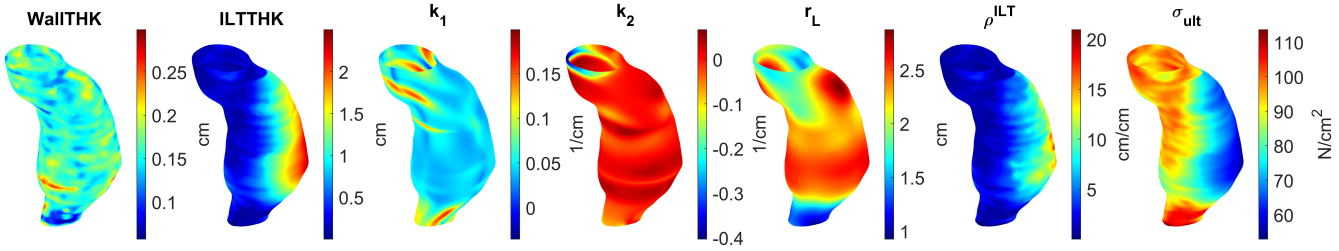


Figure 3. Exemplary graph representation of a patient-specific AAA. Color maps show selected node-specific geometric or biomechanical descriptors on the outer wall surface, illustrating the spatial variability of the features used as input to the GNNs.

Because raw meshes exhibited large size variability (from 78,000 to 900,000 nodes), we first remeshed the outer wall surface of each AAA to a common target size of approximately 150,000 nodes, as described in Section 2.1. For the final cohort, the graphs contained 147,662–153,564 nodes (a mean $\approx 150,000$ nodes) and 221,303–230,147 edges, ensuring comparable graph sizes for batching and training. In addition, we attached FEA-derived wall stresses as supervision signals. Each node i carries its mapped wall stress σ_i as a scalar regression target, and we computed graph-level summaries directly from the predicted and reference fields, e.g., $\bar{\sigma}$ and σ_{99}. Global descriptors were stored as graph attributes and broadcast to nodes when used as inputs; these include maximum diameter (*MaxDia*, in cm), encoded SEX and ILT burden metrics ($\bar{\sigma}$, σ_{99}, maximum ρ^{ILT}) together with an ILT class label. The three types of models (GGCN, EGNN, and GT) are based on the same node-specific feature vectors and mesh connectivity; EGNN also uses standardized coordinates for distance-based messages, while GT can be augmented with relative encodings derived from (x, y, z) to condition attention on geometric relationships throughout the aneurysm surface.

3.2 Model Types and Learning Tasks

We benchmarked three GNN architectures that share a common input–output formulation but differ in how they aggregate geometric information on the AAA wall. All models operate on the standardized outer wall graphs, use node-specific geometric descriptors, graph-level clinical attributes, and spatial coordinates as input features. They also predict the full node-specific wall stress field, from which $\bar{\sigma}$ and σ_{99} are derived. The first architecture is a residual GGCN that serves as a strong, scalable baseline for irregular meshes. The second is an EGNN that explicitly encodes invariance to rigid motions via distance-based messages and coordinate updates. The third is a local GT that replaces conventional message passing with multi-head self-attention over the mesh connectivity to better replicate anisotropic and long-range geometric interactions.

Shared training setup. For all architectures, node-specific inputs are formed by z-scoring and concatenating geometric descriptors, graph-level clinical attributes, and standardized coordinates, while normalized FEA wall stresses serve as regression targets. Training uses a composite, tail-aware loss that combines (i) a node-specific, stress-weighted mean squared error (up-weighting nodes in the upper stress quantiles), (ii) penalties on the discrepancy between predicted and true $\bar{\sigma}$ and between predicted and true σ_{99}, and (iii) a pinball (quantile) term to emphasize the high stress tail. To improve robustness and efficiency, we applied stress-stratified node subsampling within graphs, graph-level data augmentation on geometric and clinical descriptors, and a focal-style reweighting of "complex" aneurysms identified by high batch losses. All models were optimized with Adam and early stopping was applied as needed.

Gated graph convolution network (GGCN). The GGCN operates directly on the standardized outer wall graphs for node-specific wall stress prediction. A linear projection first maps the input features to a hidden dimension, followed by a stack of gated graph convolution layers with layer normalization, rectified linear unit (ReLU) activations, dropout and residual connections, and a final linear readout outputs one scalar stress value per node (see section A.1 in the Supplementary Material).

Equivariant graph neural network (EGNN). To explicitly encode geometric invariances, the EGNN jointly updates node features and coordinates using distance-based messages. Each layer computes edge messages from pairs of node features and squared inter-node distances, aggregates them at receiving nodes, and then applies (i) a node multilayer perceptron (MLP) to update features and (ii) a coordinate MLP to produce scalar edge-wise coefficients that drive equivariant coordinate updates along relative directions. Feature updates are followed by layer normalization, ReLU activations, dropout and residual connections, and a final linear readout maps hidden features to a single stress value per node (see section A.2 in the Supplementary Material).

Graph Transformer (GT). The GT is built from stacked multi-head self-attention layers operating on the AAA wall mesh. Each GT layer applies attention restricted to the mesh adjacency, producing edge-aware updates that reweigh information from neighboring nodes before layer normalization, ReLU activations, dropout and residual connections, and a final linear layer maps hidden features to a scalar stress prediction per node. EGNN and GT thus provide complementary alternatives to GGCN by enforcing geometric symmetries and enabling flexible attention-based aggregation of wall geometry, respectively (see section A.3 in the Supplementary Material).

3.3 Training and Evaluation

All models were trained and evaluated on the same cohort of 202 patient-specific AAA graphs using a consistent sampling, augmentation, and cross-validation strategy. To avoid information leakage and preserve the distribution of thrombus burden across splits, we performed patient-level stratified splitting based on ILT class. Thus, for each fold, 80% of the cases were assigned to the training set and 20% to the test set, with low-, moderate-, and high-ILT aneurysms, defined according to the ratio of ILT volume to AAA sac volume, proportionally represented in both subsets. This procedure was repeated for $N_{\text{fold}} = 5$ different random seeds, yielding five independent 80/20 train–test partitions per architecture; all reported metrics were averaged over the five folds.

Within each training split, the graphs were standardized using feature scalers computed from the training cases only. Node-specific geometric/biomechanical descriptors, graph-level attributes, coordinates, and wall stresses were z-scored. To improve robustness and reduce overfitting in this cohort, we applied the same data augmentation and sampling scheme for the three model types. At the graph level, we perturbed selected global features (e.g., ILT burden measures) with small multiplicative Gaussian noise and clipped them to physiologically plausible ranges. At the node level, we injected feature-wise noise into key ILTTHK, curvatures, and centerline descriptors and enforced non-negativity or upper bounds where appropriate. In addition, we performed stress-stratified subsampling within each graph during training: nodes were partitioned into low, medium, and high stress bins based on the FEA-estimated values, and a fixed fraction of each bin was retained so that high stress regions were consistently represented despite being infrequent on the surface. Node-specific loss terms were weighted with stress-based factors to further emphasize the upper tail of the wall stress distribution. For each fold and model, the training optimized a composite objective on normalized wall stresses $\{\hat{\sigma}_i\}$ using Adam optimizer with weight decay and a

ReduceLROnPlateau scheduler. The loss combined a wall stress-weighted node-specific mean squared error with penalties on discrepancies between predicted and true $\bar{\sigma}$ and between predicted and true σ_{99}, plus a pinball (quantile) term at the 0.99 level to explicitly constrain the high wall stress tail. Training proceeded for up to 1500 epochs with early stopping based on the training loss, while the best-performing weights (per fold) were checkpointed together with the corresponding feature scalers.

The evaluation began at the node level on the test data set graphs for each fold. We computed the root mean squared error (RMSE), relative error (RE), and coefficient of determination R^2 between the denormalized predictions $\hat{\sigma}_i$ and the FEA reference σ_i. In addition, we calculated the spatial agreement measured by the Pearson correlation between predicted and simulated stress fields on each AAA wall surface. Graph-level quantities were then derived directly from these node-specific fields: for every AAA we computed the $\bar{\sigma}$ and the σ_{99} using the predicted and ground truth FEA distributions, and reported RMSE, RE and R^2 for the test cohort. This formulation tested whether the models preserve clinically significant global load descriptors, particularly the high stress tail represented by σ_{99}, even when perfect agreement at every node is not feasible. All random seeds, stratified splits, feature-scaling parameters, and best-performing model weights (per fold and architecture) were archived, and training hyperparameters and hardware details are summarized in Table B1 of the Supplementary Material to allow exact replication of the numerical experiments.

4 Results

For the best-performing stratified split, which is used for all exemplar case visualizations and qualitative analyses, the three architectures accurately reconstructed the node-specific wall stress field ($\hat{\sigma}_i$) for the test data set, with node-specific R^2 values exceeding ~ 0.97 and graph-level R^2 for σ_{99} exceeding ~ 0.98. For nodes with $\sigma_i \geq 5$ N/cm^2, typical relative errors are in the $\sim$ 12–20% range, with the GT consistently achieving the lowest median RE and lower upper-tail errors than GGCN and EGNN. The highest absolute and relative errors are concentrated in the highest stress nodes, where the true stresses are substantially greater than the mean; thus, increases in absolute error reflect the greater magnitude of σ_i rather than a loss of relative accuracy of the model. Among the three, GT achieved the highest node-specific stress and σ_{99} accuracies, followed by GGCN and EGNN, indicating that attention-based aggregation on the aneurysm mesh is particularly effective in this setting.

4.1 GNN-predicted Wall Stress and Comparison with Ground Truth

Graph-level stress metrics derived from node-specific predictions also showed excellent agreement with the FEA ground truth on the best-fold test set. Figure 4 summarizes the comparisons of GNN-predicted vs. FEA-derived $\bar{\sigma}$ and σ_{99} for each GNN architecture per AAA in the test data set. All three models lie close to the identity line with regression slopes near unity and R^2 greater than ~ 0.97, but the GT consistently yields the smallest scatter and deviations from the identity line, particularly for the high stress tail quantified by σ_{99}. For individual AAAs in this fold, relative errors in graph-level metrics remained modest: for most cases, RE for $\bar{\sigma}$ was less than 15% and less than 10% for σ_{99}, with GT yielding the lowest median and upper-tail relative errors. GGCN closely follows the GT performance, while EGNN exhibits a slightly larger spread for σ_{99}, which is consistent with its lower node-specific R^2.

To examine spatial fidelity in detail, we selected two exemplary AAAs from the test set with high FEA σ_{99} and complex geometry. Figure 5 shows the predicted and reference stress maps for a representative case, and Figure 6 displays the corresponding empirical cumulative distribution functions (ECDFs) of node-specific stresses. For all GNNs, high stress regions near the sac shoulder and other regions of high curvature are accurately localized; the GT provides the closest match to the FEA maps and distributions, followed closely by the GGCN, while the EGNN tends to slightly smooth the highest peaks while preserving their locations and relative magnitudes.

4.2 Cross-Validation Results and Performance Evaluation

To rigorously evaluate model generalization and prevent overfitting, we employed a 5-fold stratified cross-validation approach in our data set of 202 patient-specific AAA geometries. The stratification was performed according to the ILT classification to ensure a balanced representation of different anatomical variants between training and test sets. The metrics reported for each fold were obtained solely from the corresponding 41 test cases of that fold, ensuring that none of these test cases were used in the training set of the same fold.

4.2.1 Cross-validation performance and variability

Table 1 summarizes the performance of the models for the 5-fold cross-validation in terms of RMSE, RE, and R^2. Although the three architectures achieve high R^2 values for both $\bar{\sigma}$ and σ_{99}, GT consistently achieves the lowest RMSE and RE, particularly for σ_{99}, indicating a more accurate recovery of the high stress tail in physical units (N/cm^2). GGCN and EGNN exhibit similar R^2 values compared to GT, but their higher RMSE and RE reveal systematically larger discrepancies in predicted stress magnitudes, even when the variance explained is comparable.

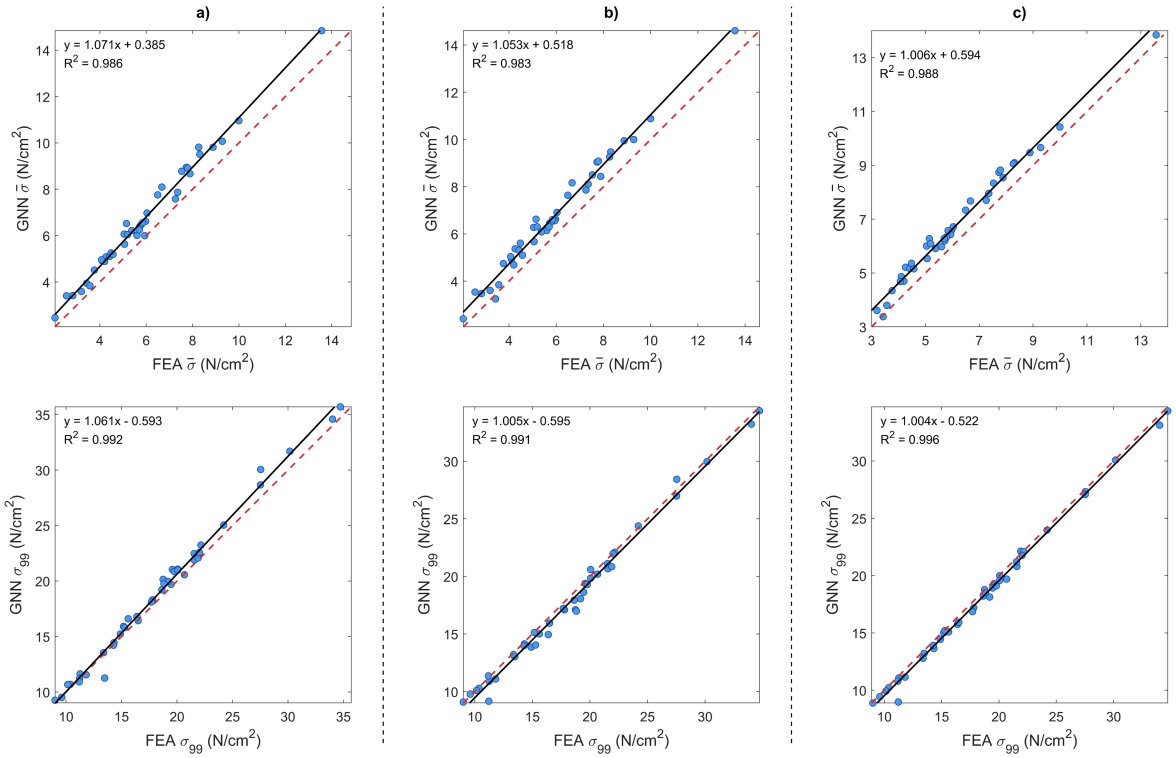


Figure 4. Test data set agreement between GNN-predicted and FEA-derived $\bar{\sigma}$ and σ_{99} for the best stratified fold for a) GGCN, b) EGNN, and c) GT. Solid black lines indicate least-squares fits; dashed red lines denote the identity line.

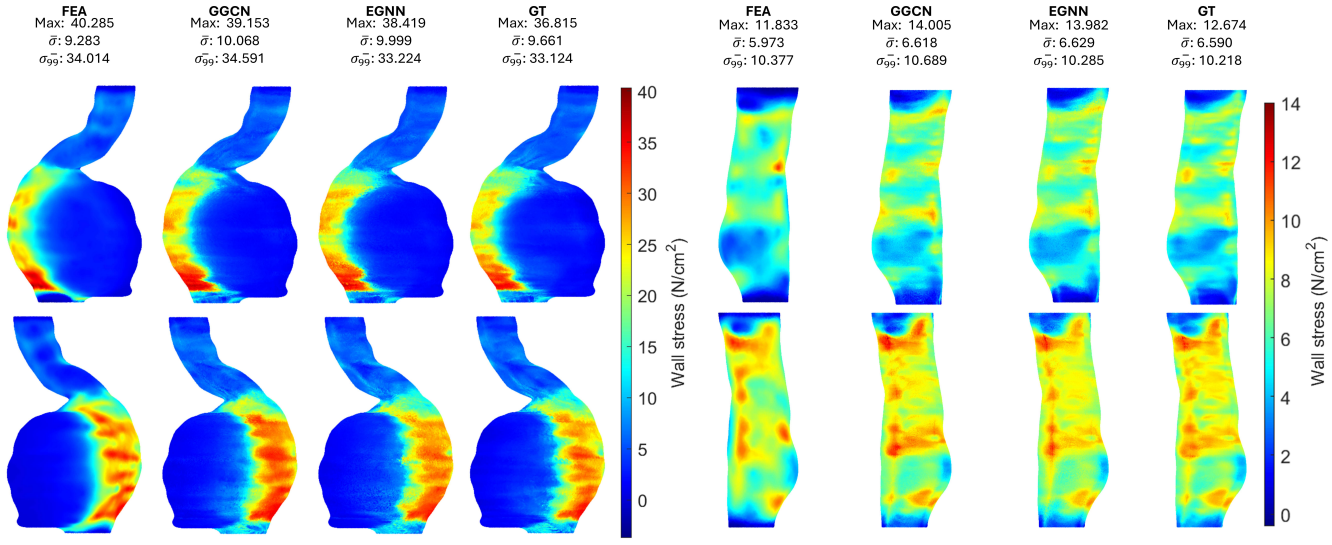


Figure 5. Spatial agreement of wall stress fields for two exemplary AAAs from the best-fold test set, one with relatively high (left frame) and one with relatively low (right frame) wall stresses. Rows show two opposite views of the same AAA, and columns show FEA ground truth, GGCN, EGNN, and GT predictions; color maps correspond to the first principal wall stress distributions projected onto the outer wall surface.

4.2.2 Individual fold analysis

Table 2 complements the aggregate metrics in Table 1 by reporting RMSE for each cross-validation fold. For $\bar{\sigma}$, GT consistently achieves the lowest RMSE in four out of the five folds, with a mean RMSE of 0.671 N/cm^2 compared to 0.741 N/cm^2 for GGCN and 0.775 N/cm^2 for EGNN. For σ_{99}, all models incur larger errors due to the higher variance of this tail metric.

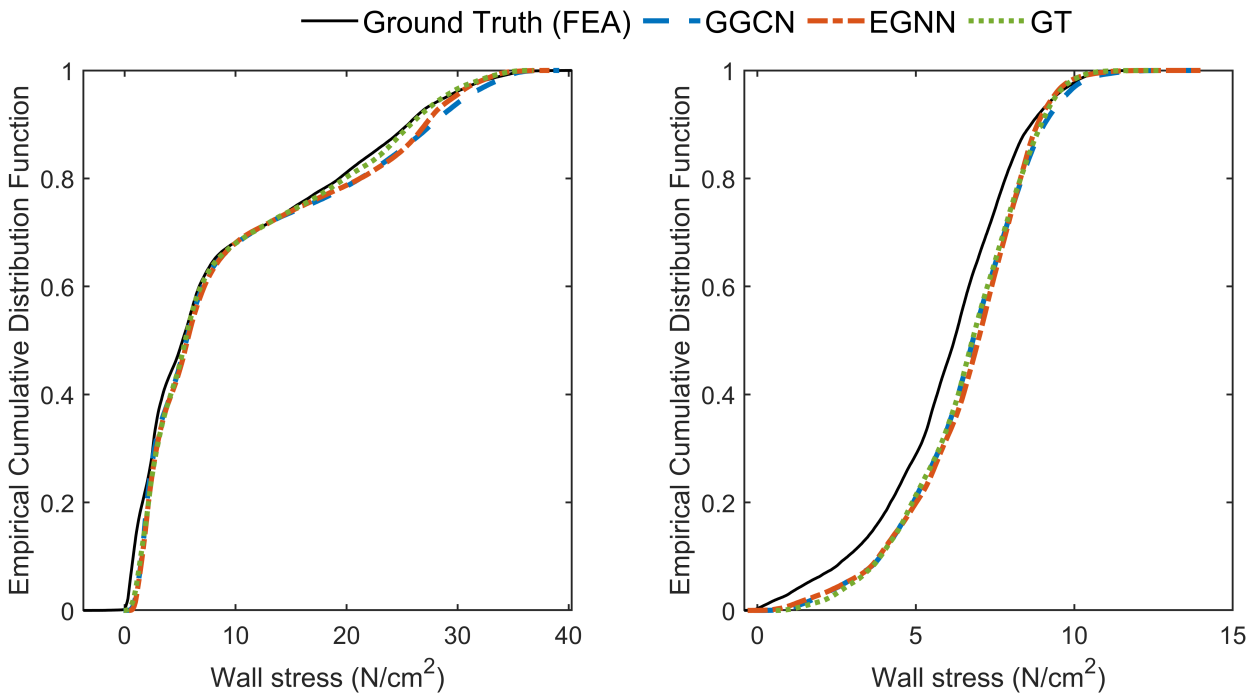


Figure 6. Empirical cumulative distribution functions (ECDFs) for the same AAAs shown in Figure 5. The left plot corresponds to the AAA illustrated on the left frame of Figure 5, while the right plot corresponds to the AAA illustrated on the right frame of Figure 5.

Table 1. Summary of model performance across 5-fold cross-validation (mean ± standard deviation). RMSE is reported in N/cm^2; RE is the mean relative error calculated for all the test cases.

Model	Metric	RMSE (N/cm^2)	RE	R^2
GGCN	$\bar{\sigma}$	0.741 ± 0.086	0.093 ± 0.021	0.977 ± 0.013
	σ_{99}	1.054 ± 0.345	0.082 ± 0.034	0.979 ± 0.015
EGNN	$\bar{\sigma}$	0.775 ± 0.066	0.098 ± 0.020	0.974 ± 0.016
	σ_{99}	1.019 ± 0.270	0.087 ± 0.030	0.984 ± 0.010
GT	$\bar{\sigma}$	0.671 ± 0.059	0.082 ± 0.016	0.978 ± 0.016
	σ_{99}	0.926 ± 0.371	0.073 ± 0.035	0.984 ± 0.011

However, GT exhibits the lowest mean RMSE averaged for all folds.

5 Discussion

The present work demonstrates the feasibility of using GNN architectures to predict patient-specific AAA wall stress distributions directly from geometric descriptors extracted from CTA-derived surface meshes of the outer wall. For the three model types, namely GGCN, EGNN, and a local GT, the node-specific wall stress field and graph-level 99th percentile wall stress were reconstructed with high fidelity relative to the FEA-derived ground truth. Evaluated on a test data set, the proposed GNN surrogates reproduced wall stresses calculated with FEA with high accuracy while reducing the inference time from hours to seconds, transforming patient-specific wall stress mapping from an offline research tool into a potentially real-time component of clinical decision support.

5.1 Key Findings and Clinical Implications

All three GNN architectures achieved on average $R^2 > 0.970$ for σ_{99}, with GT providing the best overall trade-off between accuracy and stability and the lowest σ_{99} RE (0.073). It should be noted that a σ_{99} RE in the interval 0.07–0.10 represents a small fraction of the typical wall stress for an asymptomatic AAA ($\sim$ 5–30 N/cm^2) and therefore is likely to be clinically acceptable for risk stratification. Performance was stable across the folds: the standard deviation of R^2 for $\bar{\sigma}$ and σ_{99} was in the range 0.01–0.02 for all models, with GT showing particularly low variability ($\bar{\sigma}$: $R^2 = 0.979 \pm 0.018$; σ_{99}: $R^2 = 0.985 \pm 0.012$).

Table 2. RMSE (in N/cm^2) for $\bar{\sigma}$ and σ_{99} for the five folds of the cross-validation test for each GNN model.

Model	Metric	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5
GGCN	$\bar{\sigma}$	0.673	0.794	0.747	0.625	0.867
	σ_{99}	1.448	1.079	1.368	0.498	0.879
EGNN	$\bar{\sigma}$	0.704	0.787	0.764	0.729	0.893
	σ_{99}	1.224	1.053	1.383	0.660	0.776
GT	$\bar{\sigma}$	0.636	0.784	0.642	0.618	0.675
	σ_{99}	1.296	0.998	1.334	0.406	0.594

Even in the most challenging folds, all models maintained high R^2 values for $\bar{\sigma}$ and σ_{99}, suggesting that wall stress for complex AAA geometries can be reproduced with good fidelity. Finally, by matching the FEA ground truth while reducing the wall stress inference times to seconds, these GNN-based surrogates have the potential to support near real-time clinical decision making in settings where FEA modeling may be impractical or where existing biomechanical tools (e.g., A4 Clinics) are limited by computation time and workflow complexity[16,17,47].

5.2 Model Performance, Spatial Agreement, and Error Estimates

Qualitative and quantitative analyses showed that all GNNs achieved high node-specific predictive accuracy and strong agreement for the estimates of $\bar{\sigma}$and σ_{99}. Wall stress maps for exemplary AAAs demonstrated accurate localization of high wall stress regions near the sac shoulder and other areas of high curvature, with GT providing the closest match to the FEA ground truth. These findings indicate that incorporating curvature-based features (e.g., k_1, k_2, K, M), WALLTHK, and ILTTHK provide a biomechanically meaningful context that the models can exploit to reconstruct spatial stress patterns. This is consistent with previous work associating curvature, wall thickness, and ILT burden with AAA wall stress and rupture risk[24,48].

The structure of node-specific errors further supports the reliability of the surrogates. The error distributions were narrowly centered around zero with only mild skewness, and most of the nodes exhibited stresses $\pm$ 2 N/cm^2 of the FEA ground truth even in high stress regions. At the graph level, the three architectures tend to slightly overestimate $\bar{\sigma}$; GGCN shows a small positive bias for σ_{99}, while EGNN and GT are nearly unbiased for σ_{99}. Although variance increased concomitantly with wall stress and all models were less accurate at the highest stresses, the errors were modest: near σ_{99}, typical absolute deviations (RMSE) were in the range of 1–2 N/cm^2. These absolute deviations correspond to RE less than 0.10, in line with the cross-validation results in Table 1. GT consistently exhibited the smallest error distribution and smallest high magnitude tails, followed by GGCN and EGNN, underscoring its robustness in the clinically critical high stress regime. We infer from this that GT's attention-based aggregation on the aneurysm wall surface provides a flexible mechanism to capture long-range geometric interactions that are not easily encoded by purely local message passing. In contrast, EGNN's explicit equivariance to rigid motions, while attractive from a geometric learning perspective, did not translate into superior accuracy. This is likely due to the fact that key geometric information is adequately represented by the curvature and thickness descriptors and because strict equivariance may limit representational flexibility on noisy, patient-specific surface meshes.

5.3 Importance of High Stress Region Sensitivity

Accurate identification of high stress regions is critical for rupture risk assessment, as AAA failure is expected to initiate where wall stress exceeds the local wall strength rather than where the mean stress is highest. Pure MSE training tends to underestimate the high stress regions due to the dominance of low- and moderate-stress nodes in the loss. The composite objective used in this work, which combines a stress-weighted node-specific MSE with penalties on $\bar{\sigma}$, σ_{99}, and a high-quantile pinball loss, improved accuracy in the stress tail and reduced systematic underestimation of σ_{99}, as reflected in its narrow error distributions. This behavior is consistent with previous work in medical image regression[49] and quantile regression[50], where targeted loss terms improve performance on clinically important extremes and support the use of tailored objectives when high-risk regions drive clinical decisions. It is important to note that σ_{99} alone is an incomplete surrogate of rupture risk, as σ_{ult} and its spatial heterogeneity are also necessary to adequately estimate risk[17,51]. Several studies have reported that peak wall rupture index, which combines local stress and wall strength, provides better discrimination between ruptured and asymptomatic AAAs than peak wall stress alone. In this context, the present framework should be viewed as a fast geometry-to-stress module that can be coupled with statistical or image-based wall strength models to estimate the peak wall rupture index.

5.4 Advantages over Finite Element Modeling

Compared with conventional FEA modeling pipelines, the GNN-based surrogates offer several practical advantages. Once trained, they can predict the full AAA wall stress field in seconds, while FEA simulations typically require hours and expert supervision for mesh preparation, boundary-condition specification, and material calibration[47]. The proposed approach requires

the segmented outer wall, inner wall, and lumen surfaces with derived geometric descriptors, and can therefore be more readily integrated into semi-automated imaging workflows. This computational efficiency and reduced pre-processing burden would make patient-specific wall stress assessment feasible in a clinical environment, echoing emerging efforts to employ deep neural networks as fast surrogates for vascular FEA modeling[52]. In addition, the proposed GNN framework is suitable for repeated or longitudinal analyses, where rapid recomputation of wall stress fields for evolving aneurysm geometries could support closer surveillance of biomechanical risk than what is currently possible with FEA. Fast surrogates may also enable probabilistic analyses, such as Monte Carlo sampling of material properties or blood pressure, which are often prohibitively expensive with traditional simulations but are relevant for calculating uncertainty in rupture risk estimates.

5.5 Limitations, Generalizability, and Future Work

This work is subject to several important limitations. The GNN models were trained and evaluated on a data set of 202 patient-specific geometries generated from CTA exams obtained at three clinical centers. A broader validation on larger, multi-center cohorts with more heterogeneous imaging protocols will be needed to confirm generalizability of the GNN predictions. In addition, the FEA ground truth assumed homogeneous isotropic wall and ILT material properties, whereas in vivo tissue properties are believed to be spatially heterogeneous. In addition, patient-specific material properties are never known *a priori*, which should be relevant for FEA modeling considering that ILT can substantially modulate local wall stress[1,16]. These modeling simplifications are shared with most prior biomechanical AAA studies, but may limit the absolute accuracy of the predicted wall stresses and hence of any subsequent rupture risk metrics. Future work should focus on extending the data set and systematically exploring alternative constitutive material model assumptions for the FEA ground truth, including anisotropic or spatially heterogeneous wall and ILT material models. Furthermore, augmenting the GNN with uncertainty quantification and calibration could provide case-wise confidence intervals for $\bar{\sigma}$ and σ_{99}, which would further support clinical translation and shared decision-making for elective AAA repair. Another promising direction is to integrate the current geometry-based surrogate into a broader *digital twin* framework that combines wall stress with patient-specific hemodynamics, inflammatory biomarkers, and peak wall rupture index to provide a more comprehensive biomechanical risk stratification strategy[51].

6 Conclusion

We developed, trained and tested geometry-aware graph neural network surrogates to predict patient-specific AAA wall stress distributions directly from CTA-derived surface meshes. Using a cohort of 202 AAA geometries, the three GNN architectures (GGCN, EGNN, and a local GT) reproduced FEA-estimated wall stress fields, $\bar{\sigma}$ and σ_{99} with high fidelity while reducing inference times from hours to seconds per AAA. The local GT consistently delivered the most accurate and stable predictions, particularly in the clinically critical high stress regions, suggesting that attention-based aggregation on curvature- and thickness-enriched meshes is well suited for the overall goal of the framework. These findings support the use of fast GNN-based surrogates as a key building block for AAA *digital twins* and motivate future integration with patient-specific material modeling, uncertainty quantification, and rupture risk assessment as part of a collective strategy for the clinical management of AAAs under surveillance.

References

1. Vorp, D. A. Biomechanics of abdominal aortic aneurysm. *J. Biomech.* **40**, 1887–1902, DOI: 10.1016/j.jbiomech.2006.09.003 (2007).
2. Chaikof, E. L. *et al.* The Society for Vascular Surgery practice guidelines on the care of patients with an abdominal aortic aneurysm. *J. Vasc. Surg.* **67**, 2–77.e2, DOI: 10.1016/j.jvs.2017.10.044 (2018).
3. Wanhainen, A. *et al.* Editor's Choice – European Society for Vascular Surgery (ESVS) 2024 Clinical Practice Guidelines on the Management of Abdominal Aorto-Iliac Artery Aneurysms. *Eur. J. Vasc. Endovascular Surgery: The Off. J. Eur. Soc. for Vasc. Surg.* **67**, 192–331, DOI: 10.1016/j.ejvs.2023.11.002 (2024).
4. Nicholls, S. C., Gardner, J. B., Meissner, M. H. & Johansen, H. K. Rupture in small abdominal aortic aneurysms. *J. Vasc. Surg.* **28**, 884–888, DOI: 10.1016/s0741-5214(98)70065-5 (1998).
5. Fillinger, M. F., Raghavan, M. L., Marra, S. P., Cronenwett, J. L. & Kennedy, F. E. In vivo analysis of mechanical wall stress and abdominal aortic aneurysm rupture risk. *J. Vasc. Surg.* **36**, 589–597, DOI: 10.1067/mva.2002.125478 (2002).
6. Raghavan, M. L. & Vorp, D. A. Toward a biomechanical tool to evaluate rupture potential of abdominal aortic aneurysm: identification of a finite strain constitutive model and evaluation of its applicability. *J. Biomech.* **33**, 475–482, DOI: 10.1016/S0021-9290(99)00201-8 (2000).
7. Fillinger, M. F., Marra, S. P., Raghavan, M. & Kennedy, F. E. Prediction of rupture risk in abdominal aortic aneurysm during observation: Wall stress versus diameter. *J. Vasc. Surg.* **37**, 724–732, DOI: 10.1067/mva.2003.213 (2003).

8. Wu, W. *et al.* Wall Stress and Geometry Measures in Electively Repaired Abdominal Aortic Aneurysms. *Annals Biomed. Eng.* **47**, 1611–1625, DOI: 10.1007/s10439-019-02261-w (2019).

9. Chung, T. K., da Silva, E. S. & Raghavan, S. M. L. Does elevated wall tension cause aortic aneurysm rupture? Investigation using a subject-specific heterogeneous model. *J. Biomech.* **64**, 164–171, DOI: 10.1016/j.jbiomech.2017.09.041 (2017).

10. Raghavan, M. L., Fillinger, M. F., Marra, S. P., Naegelein, B. P. & Kennedy, F. E. Automated methodology for determination of stress distribution in human abdominal aortic aneurysm. *J. Biomech. Eng.* **127**, 868–871, DOI: 10.1115/1.1992530 (2005).

11. Holzapfel, G. A. Determination of material models for arterial walls from uniaxial extension tests and histological structure. *J. Theor. Biol.* **238**, 290–302, DOI: 10.1016/j.jtbi.2005.05.006 (2006).

12. Vande Geest, J. P., Sacks, M. S. & Vorp, D. A. The effects of aneurysm on the biaxial mechanical behavior of human abdominal aorta. *J. Biomech.* **39**, 1324–1334, DOI: 10.1016/j.jbiomech.2005.03.003 (2006).

13. Sacks, M. S. & Sun, W. Multiaxial Mechanical Behavior of Biological Materials. *Annu. Rev. Biomed. Eng.* **5**, 251–284, DOI: 10.1146/annurev.bioeng.5.011303.120714 (2003). Publisher: Annual Reviews.

14. Vande Geest, J. P., Schmidt, D. E., Sacks, M. S. & Vorp, D. A. The Effects of Anisotropy on the Stress Analyses of Patient-Specific Abdominal Aortic Aneurysms. *Annals Biomed. Eng.* **36**, 921–932, DOI: 10.1007/s10439-008-9490-3 (2008).

15. Maier, A. *et al.* A Comparison of Diameter, Wall Stress, and Rupture Potential Index for Abdominal Aortic Aneurysm Rupture Risk Prediction. *Annals Biomed. Eng.* **38**, 3124–3134, DOI: 10.1007/s10439-010-0067-6 (2010).

16. Reeps, C. *et al.* Measuring and modeling patient-specific distributions of material properties in abdominal aortic aneurysm wall. *Biomech. Model. Mechanobiol.* **12**, 717–733, DOI: 10.1007/s10237-012-0436-1 (2013).

17. Singh, T. P., Moxon, J. V., Gasser, T. C. & Golledge, J. Systematic Review and Meta-Analysis of Peak Wall Stress and Peak Wall Rupture Index in Ruptured and Asymptomatic Intact Abdominal Aortic Aneurysms. *J. Am. Hear. Assoc.* **10**, e019772, DOI: 10.1161/JAHA.120.019772 (2021). Publisher: Wiley.

18. Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Navab, N., Hornegger, J., Wells, W. M. & Frangi, A. F. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, 234–241, DOI: 10.1007/978-3-319-24574-4_28 (Springer International Publishing, Cham, 2015).

19. Wu, W. *et al.* Rapid automated lumen segmentation of coronary optical coherence tomography images followed by 3D reconstruction of coronary arteries. *J. Med. Imaging* **11**, 014004, DOI: 10.1117/1.JMI.11.1.014004 (2024). Publisher: SPIE.

20. Roby, M. *et al.* Automatic Segmentation of Abdominal Aortic Aneurysm From Computed Tomography Angiography Using a Patch-Based Dilated UNet Model. *IEEE Access* **13**, 24544–24554, DOI: 10.1109/ACCESS.2025.3533417 (2025).

21. Sacks, M. S., Vorp, D. A., Raghavan, M. L., Federle, M. P. & Webster, M. W. In Vivo Three-Dimensional Surface Geometry of Abdominal Aortic Aneurysms. *Annals Biomed. Eng.* **27**, 469–479, DOI: 10.1114/1.202 (1999).

22. Martufi, G., Di Martino, E. S., Amon, C. H., Muluk, S. C. & Finol, E. A. Three-Dimensional Geometrical Characterization of Abdominal Aortic Aneurysms: Image-Based Wall Thickness Distribution. *J. Biomech. Eng.* **131**, DOI: 10.1115/1.3127256 (2009).

23. Georgakarakos, E. *et al.* The Role of Geometric Parameters in the Prediction of Abdominal Aortic Aneurysm Wall Stress. *Eur. J. Vasc. Endovascular Surg.* **39**, 42–48, DOI: 10.1016/j.ejvs.2009.09.026 (2010).

24. Shum, J. *et al.* Quantitative Assessment of Abdominal Aortic Aneurysm Geometry. *Annals Biomed. Eng.* **39**, 277–286, DOI: 10.1007/s10439-010-0175-3 (2011).

25. Vorp, D. A., Raghavan, M. L. & Webster, M. W. Mechanical wall stress in abdominal aortic aneurysm: Influence of diameter and asymmetry. *J. Vasc. Surg.* **27**, 632–639, DOI: 10.1016/S0741-5214(98)70227-7 (1998).

26. Hamilton, W. L., Ying, R. & Leskovec, J. Inductive representation learning on large graphs, DOI: 10.48550/arXiv.1706.02216 (2018). ArXiv:1706.02216 [cs].

27. Hamilton, W. L. Graph representation learning. *Synth. Lect. on Artif. Intell. Mach. Learn.* **14**, 1–159 (2020).

28. Zhou, J. *et al.* Graph neural networks: A review of methods and applications. *AI Open* **1**, 57–81, DOI: 10.1016/j.aiopen.2021.01.001 (2020).

29. Bhagat, S., Cormode, G. & Muthukrishnan, S. Node classification in social networks. In Aggarwal, C. C. (ed.) *Social Network Data Analytics*, 115–148, DOI: 10.1007/978-1-4419-8462-3_5 (Springer US, Boston, MA, 2011).

30. Li, Y., Tarlow, D., Brockschmidt, M. & Zemel, R. Gated graph sequence neural networks. In *International Conference on Learning Representations (ICLR)* (2016). 1511.05493.

31. Batzner, S. *et al.* E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **13**, 2453, DOI: 10.1038/s41467-022-29939-5 (2022). Publisher: Nature Publishing Group.

32. Musaelian, A. *et al.* Learning local equivariant representations for large-scale atomistic dynamics. *Nat. Commun.* **14**, 579, DOI: 10.1038/s41467-023-36329-y (2023). Publisher: Nature Publishing Group.

33. Satorras, V. G., Hoogeboom, E. & Welling, M. E(n) Equivariant Graph Neural Networks (2021).

34. Shaw, P., Uszkoreit, J. & Vaswani, A. Self-Attention with Relative Position Representations. In Walker, M., Ji, H. & Stent, A. (eds.) *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 464–468, DOI: 10.18653/v1/N18-2074 (Association for Computational Linguistics, New Orleans, Louisiana, 2018).

35. Dwivedi, V. P. & Bresson, X. A Generalization of Transformer Networks to Graphs, DOI: 10.48550/arXiv.2012.09699 (2021). ArXiv:2012.09699 [cs].

36. Yun, S., Jeong, M., Kim, R., Kang, J. & Kim, H. J. Graph transformer networks. In Wallach, H. *et al.* (eds.) *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32 (Curran Associates, Inc., 2019).

37. Ying, C. *et al.* Do transformers really perform bad for graph representation? In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS '21, 28877–28888 (Curran Associates Inc., Red Hook, NY, USA, 2021).

38. Thölke, P. & Fabritiis, G. D. Equivariant transformers for neural network based molecular potentials. In *International Conference on Learning Representations* (2022).

39. Carmo, M. P. d. *Differential Geometry of Curves and Surfaces: Revised and Updated Second Edition* (Courier Dover Publications, 2016). Google-Books-ID: uXF6DQAAQBAJ.

40. Meyer, M., Desbrun, M., Schröder, P. & Barr, A. H. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds. In Hege, H.-C. & Polthier, K. (eds.) *Visualization and Mathematics III*, 35–57 (Springer Berlin Heidelberg, Berlin, Heidelberg, 2003).

41. Smith, D. B., Sacks, M. S., Vorp, D. A. & Thornton, M. Surface Geometric Analysis of Anatomic Structures Using Biquintic Finite Element Interpolation. *Annals Biomed. Eng.* **28**, 598–611, DOI: 10.1114/1.1306342 (2000).

42. Antiga, L. *et al.* An image-based modeling framework for patient-specific computational hemodynamics. *Med. & Biol. Eng. & Comput.* **46**, 1097–1112, DOI: 10.1007/s11517-008-0420-1 (2008).

43. Pappu, S., Dardik, A., Tagare, H. & Gusberg, R. J. Beyond Fusiform and Saccular: A Novel Quantitative Tortuosity Index May Help Classify Aneurysm Shape and Predict Aneurysm Rupture Potential. *Annals Vasc. Surg.* **22**, 88–97, DOI: 10.1016/j.avsg.2007.09.004 (2008).

44. Vande Geest, J. P., Di Martino, E. S., Bohra, A., Makaroun, M. S. & Vorp, D. A. A biomechanics-based rupture potential index for abdominal aortic aneurysm risk assessment: demonstrative application. *Annals New York Acad. Sci.* **1085**, 11–21 (2006).

45. Raut, S. S., Liu, P. & Finol, E. A. An approach for patient-specific multi-domain vascular mesh generation featuring spatially varying wall thickness modeling. *J. Biomech.* **48**, 1972–1981, DOI: https://doi.org/10.1016/j.jbiomech.2015.04.006 (2015).

46. Vande Geest, J. P., Sacks, M. S. & Vorp, D. A. A planar biaxial constitutive relation for the luminal layer of intra-luminal thrombus in abdominal aortic aneurysms. *J. Biomech.* **39**, 2347–2354, DOI: https://doi.org/10.1016/j.jbiomech.2006.05.011 (2006).

47. Grassl, K. *et al.* Early Prediction of Abdominal Aortic Aneurysm Rupture Risk Using Numerical Biomechanical Analysis. *Diagnostics* **15**, 25, DOI: 10.3390/diagnostics15010025 (2024).

48. Azar, D. *et al.* Mechanical and geometrical determinants of wall stress in abdominal aortic aneurysms: A computational study. *PLOS ONE* **13**, e0192032, DOI: 10.1371/journal.pone.0192032 (2018).

49. Chao, H., Zhang, J. & Yan, P. Regression Metric Loss: Learning a Semantic Representation Space for Medical Images. In Wang, L., Dou, Q., Fletcher, P. T., Speidel, S. & Li, S. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, 427–436, DOI: 10.1007/978-3-031-16452-1_41 (Springer Nature Switzerland, Cham, 2022).

50. Koenker, R. & Hallock, K. F. Quantile Regression. *J. Econ. Perspectives* **15**, 143–156, DOI: 10.1257/jep.15.4.143 (2001).

51. Siika, A. *et al.* Peak wall rupture index is associated with risk of rupture of abdominal aortic aneurysms, independent of size and sex. *The Br. J. Surg.* **111**, znae125, DOI: 10.1093/bjs/znae125 (2024).

52. Liang, L., Liu, M., Elefteriades, J. & Sun, W. Synergistic integration of deep neural networks and finite element method with applications of nonlinear large deformation biomechanics. *Comput. Methods Appl. Mech. Eng.* **416**, 116347, DOI: 10.1016/j.cma.2023.116347 (2023).

Funding Declaration

This work was supported in part by the National Institutes of Health award R01HL159300 and the Zachry Department Chair Endowment in Mechanical Engineering.

Author contributions statement

J.C.R. designed the study, processed the data, and wrote the manuscript. S.C.M., M.K.E. and V.S.K supervised the data collection. V.D.O. oversaw and revised the statistical analysis. E.A.F. acquired the funding and supervised the study. All authors read, edited, and agreed with the manuscript.

Additional information

Competing interests. M.K.E. has a paid consulting role with W.L. Gore and Boston Scientific. The other authors declare that they have no competing interests. **Data availability.** The data set used in this study is not publicly accessible due to IRB regulations in the respective clinical centers. However, it can be obtained by making a reasonable request to the corresponding author.

Acknowledgements

E.A.F. thanks the Zachry Mechanical Engineering Department Chair Endowment and the National Institutes of Health (Award No. R01HL159300) for providing financial support for the execution of this work. The authors also acknowledge the use of the UT San Antonio ARC high-performance computing cluster, which provided the computational resources required for the numerical experiments.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryMaterialCarvajaletal.pdf