







Robust automated calcification meshing for personalized cardiovascular biomechanics



Daniel H. Pak¹  , Minliang Liu², Theodore Kim¹, Caglar Ozturk^{3,4} , Raymond McKay⁵, Ellen T. Roche³, Rudolph Gleason⁶ & James S. Duncan¹

Calcification has significant influence over cardiovascular diseases and interventions. Detailed characterization of calcification is thus desired for predictive modeling, but calcium deposits on cardiovascular structures are still often manually reconstructed for physics-driven simulations. This poses a major bottleneck for large-scale adoption of computational simulations for research or clinical use. To address this, we propose an end-to-end automated image-to-mesh algorithm that enables robust incorporation of patient-specific calcification onto a given cardiovascular tissue mesh. The algorithm provides a substantial speed-up from several hours of manual meshing to ~1 min of automated computation, and it solves an important problem that cannot be addressed with recent template-based meshing techniques. We validated our final calcified tissue meshes with extensive simulations, demonstrating our ability to accurately model patient-specific aortic stenosis and Transcatheter Aortic Valve Replacement. Our method may serve as an important tool for accelerating the development and usage of personalized cardiovascular biomechanics.

Calcification plays an important role in cardiovascular diseases. High levels of calcification in the coronaries, aorta, and heart valves have all been shown to be effective predictors of cardiovascular disease incidence and death^{1–4}, and vascular calcification has been correlated with atherosclerosis^{5,6}. For aortic stenosis (AS), degenerative calcification and the subsequent over-stiffening of the aortic valve leaflets is the most common etiology in industrialized countries⁷, and calcification is a diagnostic biomarker for inconclusive AS⁸.

Beyond simple quantification and statistical analyses, physics-driven biomechanical simulations can help elevate our understanding of the effects of calcification on personalized cardiovascular physiology. Aortic stenosis caused by heavy calcification often increases the ventricular afterload, resulting in mitral regurgitation, atrial fibrillation, and secondary cardiomyopathy^{9–11}. The morphometry of calcification can significantly alter the hemodynamics^{12,13}, leading to increased risks of aortic dissection¹⁴. The irregular stress patterns induced by calcification may stimulate further progression of calcification^{15–17}. Calcification heavily influences the outcome of Transcatheter Aortic Valve Replacements (TAVR), potentially causing various complications, such as aortic rupture, paravalvular leak, and conduction abnormalities^{18,19}. Simulation-based digital twins have the potential to help predict related disease progressions and treatment outcomes, and

they have been suggested as a tool for patient-specific risk assessment and procedural planning^{20,21}.

Accurate modeling of calcification has thus been a long-standing clinical interest, leading to notable progress in automated algorithms for calcification segmentation. For non-contrast computed tomography (CT), global thresholding and filtering via aorta segmentation has been suggested for calcification segmentation²². For CT angiography (CTA), where the lumen intensity is much higher and may vary across different patients, adaptive thresholding based on luminal attenuation is the most common approach^{23–26}. Learning-based methods have also been proposed using hand-crafted features^{27,28} and deep learning²⁹.

The caveat is that incorporating calcification into simulations often requires converting the predicted segmentation into a finite element mesh. The conversion generally consists of extensive manual work, primarily due to the complexity of the final mesh and the unreliability of automated meshing operations^{20,27}. Especially when accurate multi-part interaction is a key contributor of simulation accuracy, the raw segmentation must often be further processed in voxelgrid, spline surface, and/or mesh representations to accurately reflect the true physical relations between different components. Furthermore, mesh boolean operations are notoriously error-prone and often necessitates manual mesh cleaning. Manual geometry



¹Yale University, 300 Cedar St, New Haven, CT, 06511, USA. ²Texas Tech University, 805 Boston Avenue, Lubbock, TX, 79409, USA. ³Massachusetts Institute of Technology, 45 Carleton St, Cambridge, MA, 02142, USA. ⁴University of Southampton, University Road, Southampton, SO17 1BJ, UK. ⁵Hartford Hospital, 85 Seymour St, Hartford, CT, 06106, USA. ⁶Georgia Institute of Technology, 315 Ferst Dr NW, Atlanta, GA, 30332, USA.  e-mail: daniel.pak@yale.edu

Fig. 1 | An overview schematic of the C-MAC algorithm and its representative results. a Starting from a 3D CTA, a patient-specific volumetric mesh of cardiovascular tissues is first generated using DeepCarve. The tissue mesh is used to aid the voxelgrid segmentation of calcification, as well as to generate the background mesh for DM Tet. The SDF values are first initialized by sampling the voxelgrid segmentation at each node of the background grid, and subsequently the node coordinates and nodal SDF values are optimized for better element quality. The resulting DM Tet output mesh is further processed via node-based remeshing to generate the final input for tetrahedralization. **b** C-MAC robustly and automatically incorporates a voxelgrid segmentation onto an existing mesh without changing the original mesh topology. Here, we demonstrate its performance using two vastly different calcification segmentations (orange) and complex aortic valve meshes (gray).



adjustments and mesh cleaning could take several hours per model for a trained expert, which introduces a key bottleneck in the simulation workflow.

Alternatives exist, such as assigning different element properties on the tissue geometry^{30,31}, defining surface-based tie constraints³², and remeshing calcification with a nearby tissue³³. However, existing workarounds are suboptimal due to various reasons, such as missing geometry, instability in complex simulations, laborious simulation setup for each patient, and loss of tissue mesh correspondence.

Clinically, severe aortic stenosis is often characterized by large amounts of calcification that occupies a significant portion of the aortic root and the left ventricular outflow tract. The calcification geometry itself can thus be important for modeling fused leaflets, detailed hemodynamics, and the deformed configuration of valve prostheses (Supplementary Fig. 2). Simplifying the calcification to stiffer tissue properties would fail to capture such conditions without manual patient-specific modifications involving advanced simulation techniques. Similarly, establishing surface tie constraints with suboptimal calcification meshes requires significant manual adjustments of the constraint parameters, potentially for each island of segmented calcification. Incorporating calcification via remeshing presents more nuanced disadvantages associated with the loss of mesh correspondence, such as the inability to use a fixed tissue mesh topology to perform motion tracking and large cohort analyses. If the original tissue mesh was generated by deforming an optimized template mesh with high-quality elements, the remeshing operation would also negate the advantages of the template-based operation.

In contrast, we propose a robust fully automated calcification meshing algorithm that satisfies the following challenging constraints:

- Accurate calcification geometry
- Clean manifold surface for tetrahedralization
- Preservation of the input tissue mesh topology
- Complete surface matching (i.e. node-to-node and edge-to-edge) along the tissue-calcification contact surfaces

A key benefit of our approach is that it enables large-scale physics-driven simulations with complex multi-body interactions. The initial geometry definition as well as the simulation setup are both massively simplified due to our final mesh characteristics, resulting in a speed-up of up to a few

orders of magnitude for the simulation setup. Together with the recent advances in template-based meshing algorithms, our method could serve as an important tool for accelerating the development and usage of physics-driven simulations for personalized cardiovascular biomechanics.

Our approach is summarized in Fig. 1a and Algorithm 1. First, we utilize deep learning models to obtain the initial cardiovascular tissue mesh and calcification segmentation. We convert them into inputs for Deep Marching Tetrahedra (DM Tet), namely the background tetrahedral mesh and the post-processed segmentation. Then, we obtain an optimized calcification surface mesh using DM Tet optimization followed by constrained remeshing, and finally perform constrained tetrahedralization to obtain the simulation-ready calcification mesh.

Our technical contributions are spread across the entire image-to-mesh pipeline:

- Improved deep learning model for tissue mesh registration
- A deep learning segmentation model for the initial calcification segmentation
- Segmentation post-processing algorithm with adaptive morphological operations
- A novel DM Tet optimization scheme for exact conforming surfaces with high spatial accuracy and element quality
- A novel constrained remeshing algorithm that maintains the contact surface conformity and further improves the element quality

Results

Cardiovascular tissue mesh reconstruction

For automated meshing of the calcified cardiovascular structures, we must consider two main components with substantially different geometrical characteristics: the native tissue and the calcification. For the native tissue, we modeled the aortic valve, ascending aorta, and partial left ventricular myocardium (LV). This allowed us to test our algorithm on a sizable portion of the ascending aorta with the complex aortic valve geometry included in the modeling domain.

From the numerous recently proposed algorithms^{34–37}, we decided to use a slightly modified version of DeepCarve for tissue mesh reconstruction³⁸, mainly to take advantage of its demonstrated performance on aortic valve meshing. Here, we briefly describe the method for completeness.

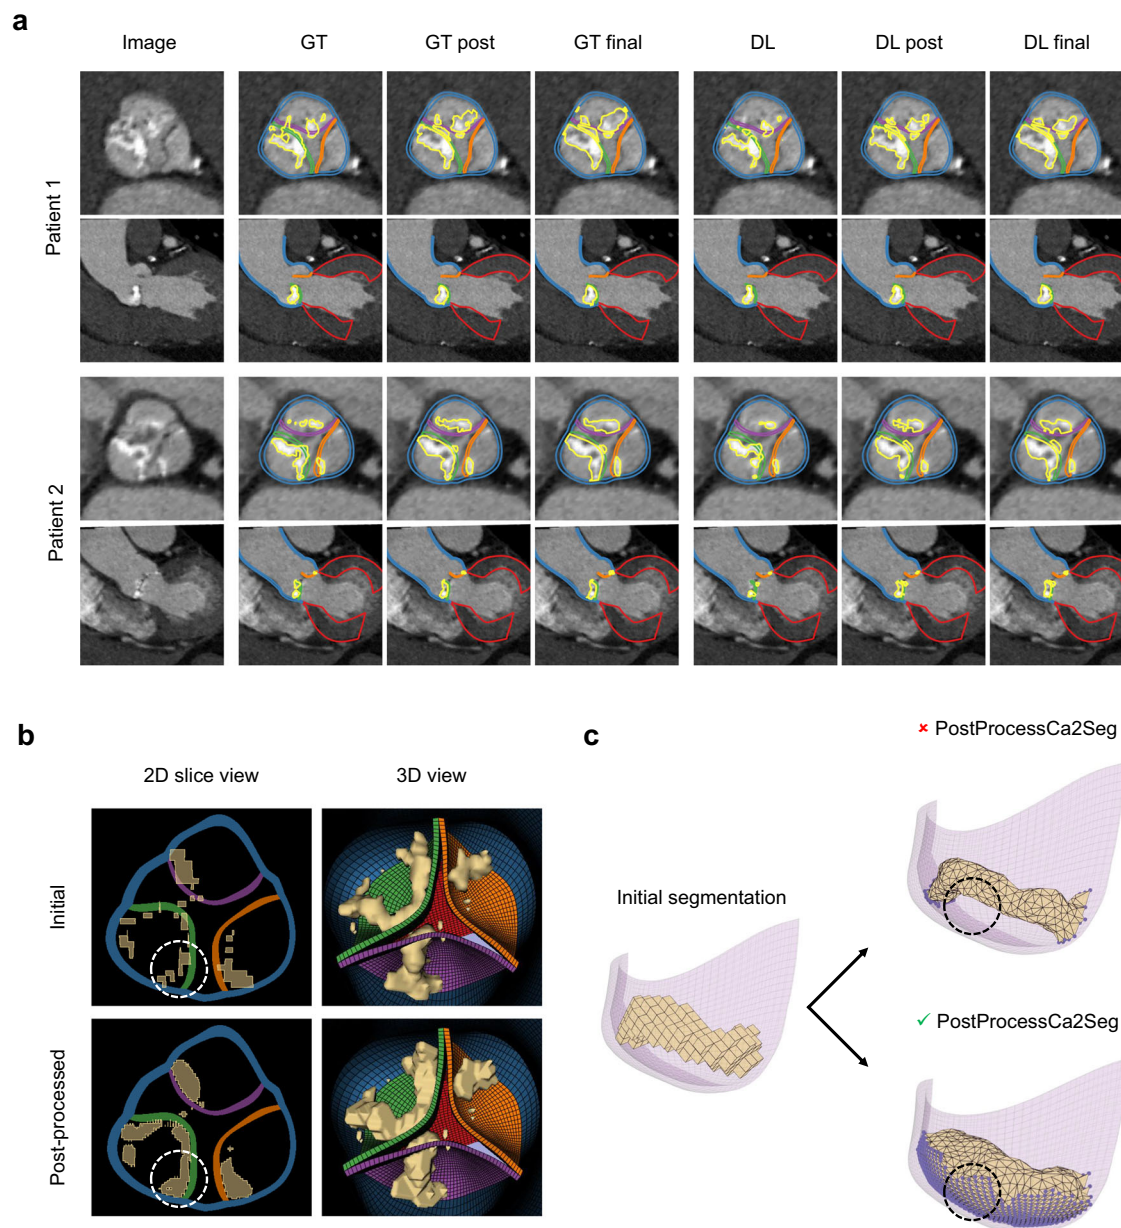


Fig. 2 | Qualitative evaluations of the initial segmentation and post-processing algorithms. **a** Assessing the segmentation accuracy in two image slice views from two test-set patients. GT ground-truth segmentation, DL deep learning segmentation using “GDL (ours)”, post: result of `PostProcessCa2Seg` using either GT or DL as the initial segmentation, final: result of the full C-MAC. Yellow: calcification, red: partial LV myocardium, blue: aorta, (green, orange, purple): aortic valve leaflets. **b** Before and after `PostProcessCa2Seg` on one test-set patient, where the white

dashed circles highlight the regions with clear gap-closing effects on the voxelgrid segmentation. **c** Visualizing the effects of `PostProcessCa2Seg` on the final C-MAC mesh. Purple dots indicate the merged nodes between the calcification and the leaflet mesh, and the black dashed circles indicate the regions with large changes in the merged nodes. This illustrates both the benefit and drawback of our post-processing algorithm. Benefit: improved anatomical consistency with the surrounding tissue. Drawback: some overestimation of calcified regions.

The main objective for DeepCarve is to train a deep learning model that will take an input image and generate a patient-specific mesh by deforming a template mesh. The required training label is a set of component-specific pointclouds that can help measure the spatial accuracy of the prediction. The predicted mesh is penalized for large deformations from the original template elements via deformation gradient-based isotropic and anisotropic energies. Dense displacement fields are obtained via b-spline interpolation of the control point displacements predicted by the deep learning model, followed by the scaling-and-squaring method to enforce approximate diffeomorphism. True diffeomorphism is not guaranteed, but combined with the 3 mm control point spacing, this method leads to zero degenerate elements in practice³⁸.

The three minor modifications for this work are (1) a better LV template with fewer distorted elements, and (2) the addition of $0.1 \cdot \text{AMIPS}^{39}$ to the training loss, and (3) random translation augmentation. These changes mainly allowed us to adapt to more realistic meshing environments with the whole image volume instead of pre-cropped images, while maintaining similar model performance.

Once DeepCarve is trained, a single forward-pass of the 3D patient scan directly generates high-quality volumetric meshes of the cardiovascular tissue geometry in <1 s. We confirmed the accuracy of the reconstructed geometry (Figs. 1a and 2a) and focused on robustly incorporating calcification to the resulting output. For additional detail, we refer the readers to the original publication³⁸.

Calcification meshing with anatomical consistency

Our calcification meshing algorithm requires two inputs: the original patient CTA scan and the reconstructed cardiovascular tissue mesh. The main goal is to incorporate calcification into the tissue mesh for more accurate and realistic structural simulations. The major steps are outlined in Fig. 1a and Algorithm 1, and samples of input/output pairs are shown in Fig. 1b.

A deep learning model is used to perform the initial calcification segmentation, and a series of model-based post-processing techniques are used to ensure the anatomical consistency between the predicted segmentation and the reconstructed tissue mesh. The segmentation and the tissue mesh are then processed intricately together to generate the final calcification mesh. The meshing steps are centered around Deep Marching Tetrahedra (DMTet) and controlled mesh simplification via constrained remeshing. The details and the performance of each step will be discussed in the following sections. We will henceforth refer to our algorithm as C-MAC (Calcification Meshing with Anatomical Consistency).

Algorithm 1. C-MAC overview

```

1: function C-MAC( $I, M_{tissue}$ )
2:    $y_0 \leftarrow \text{DLCa2Seg}(I)$ 
3:    $y_{ca2} \leftarrow \text{PostProcessCa2Seg}(y_0, M_{tissue})$ 
4:    $M_{bg} \leftarrow \text{GenBGMesh}(M_{tissue})$ 
5:    $S_{ca2} \leftarrow \text{DMTetOpt}(y_{ca2}, M_{bg})$ 
6:   if not satisfied then
7:      $S_{ca2} \leftarrow \text{ConstrainedRemesh}(S_{ca2})$ 
8:      $M_{ca2} \leftarrow \text{TetGen}(S_{ca2})$ 
9:   return  $M_{ca2}$ 
I: image,  $y$ : segmentation,  $M$ : volume mesh,  $S$ : surface mesh

```

Initial segmentation via deep learning (DLCa2Seg)

Given the inherent voxelgrid structure of 3D medical images, the most common intermediate representation for meshing is voxelgrid segmentation. Following this trend, we first focused on accurately extracting the initial segmentation of the calcified regions within the CTA.

Thresholding techniques are easy to implement and generally perform well for calcification segmentation. However, they are particularly vulnerable to imaging artifacts due to their simplistic approach, and they are usually supplemented with manual post-processing to refine the segmentation. Deep learning (DL) models, on the other hand, utilize learned feature aggregation to make the final prediction, freeing themselves from the simple errors of thresholding approaches. Therefore, for our task, we aimed to train a DL model that can handle the following conditions:

- Highly unbalanced target (i.e. calcification is very sparse compared to the input image volume)
- Empty target (i.e. no calcification in the input image)
- Accurate segmentation boundary

To meet these criteria, we used a modified 3D U-net⁴⁰ and evaluated four different training strategies involving the following losses:

- Dice + Cross Entropy (DiceCE)
- Generalized Dice Loss (GDL)
- GDL + boundary loss (Gbd)
- GDL + chamfer loss (Gch)

The evaluations are summarized in Table 1 (top). We first observed that nnU-net⁴¹, a popular self-configurable segmentation algorithm, suffers from a high rate of false positive segments outside of the aorta. All of our methods had better overall performance than the self-configured nnU-net. The only metric that nnU-net performed relatively well was the number of false positives for empty targets (i.e. when there is no aortic calcification in the CTA scan).

Among our own models, we first compared the effects of different image intensity ranges for the min-max normalization, while using

DiceCE or GDL. The metrics generally indicate that (1) using the right intensity range can help boost the performance and (2) GDL with I minmax $[-200, 1500]$ achieves the best overall performance.

Then, we used the best performing training setup from our initial evaluation and tested the modified GDL losses: Gbd and Gch. Gbd was effective at reducing false positives, but otherwise both Gbd and Gch showed worse overall performance than the standard GDL loss. For our task, false positives for empty targets were less critical than accurate prediction of non-empty targets because they can be easily identified and omitted using the source image. Therefore, we used the GDL with I minmax $[-200, 1500]$ as our final segmentation model for downstream tasks.

Figure 2a shows a few representative examples of ground-truth segmentation vs. DL segmentation using “GDL (ours)” (column 2 vs. column 5). The DL segmentation qualitatively matches the ground-truth very well.

Lastly, we performed one additional evaluation of nnU-net using our chosen image normalization scheme in place of the self-configured CT normalization. The overall performance was generally similar to the self-configured nnU-net, but the quantitative metrics indicated visibly worse performance due to a few egregious errors that were later filtered by the post-processing steps.

Segmentation post-processing (PostProcessCa2Seg)

Although our DL model delivered promising initial results, it still falls short of the high standards required for simulations. In particular, the predicted segmentation may have slight spatial gaps (2–3 voxels) to the input cardiovascular tissue mesh, which are generally negligible errors but not for establishing contacts between the predicted calcification and the surrounding tissue. These small errors could drastically alter the results of the simulations, so we aimed to explicitly improve the anatomical consistency of the DL predictions by removing any obvious gaps between the calcification segmentation and the tissue mesh.

Our main method for filling in the gaps is morphological closing (dilation \rightarrow erosion). Unfortunately, a naive application of this approach with differently sized isotropic filters leads to overly connected components, such as undesired calcification fusing and excessive contact with the neighboring tissue structures. We incorporated several additional operations, such as voxelgrid upsampling and distance-based combination of identity and anisotropic filters to minimize the deviation from the initial segmentation while still being able to fill in the gaps.

For segmentation post-processing, we first evaluated its effects on the final segmentation’s spatial accuracy and anatomical consistency (Table 1 bottom). For isolated evaluation of the post-processing steps, we also included an evaluation of the post-processed ground-truth segmentation. The overall improvements in the mean surface accuracy metric (CD) combined with our qualitative evaluations (Fig. 2a) indicate that our algorithm reasonably preserves the initial segmentation’s spatial characteristics. Significant improvements in the calcification-to-tissue distance (CD_{tissue}) indicates reduced spatial gaps and better anatomical consistency between calcification and its surrounding tissue.

The improved surface metrics HD and CD demonstrate our algorithm’s effectiveness in filtering out easily avoidable false-positive calcification segments, such as those outside the aorta and those with tiny volumes. Interestingly, this filtering process reduced the performance gap between different segmentation models (Table 1 top \rightarrow Table 1 bottom), which suggests that our post-processing algorithm was effective at removing segmentation outliers for various models.

The worse Dice and better CD_{tissue} demonstrates the trade-off between volume overlap and anatomical consistency. This is expected, since the current strategy for removing the calcification-to-tissue spatial gaps is to extend the calcification segments while keeping the tissue mesh fixed. For the purposes of calcification meshing, the benefit of anatomical consistency far outweighs the slight reduction in volume overlap (Fig. 2b, c). An interesting future direction is to jointly optimize both the tissue mesh and the calcification segmentation to encourage anatomical consistency while minimizing this trade-off.

Table 1 | Deep learning model performances on the initial calcification segmentation (DLCa2Seg) with and without post-processing (PostProcessCa2Seg)

Experimental condition			Non-empty target segmentation				Empty target
Train strategy	<i>l</i> minmax	Post	Dice ↑	HD 95% (mm) ↓	CD (mm) ↓	CD _{tissue} (mm) ↓	False pos (vx) ↓
nnU-net	Self-configured	✗	0.656 ± 0.222	13.26 ± 18.06	2.337 ± 2.994	2.501 ± 2.955	5.4 ± 10.0
nnU-net	[-200, 1500]	✗	0.604 ± 0.241	22.35 ± 39.90	4.271 ± 11.143	6.508 ± 19.769	0.2 ± 0.7
DiceCE	[-158, 864]	✗	0.685 ± 0.180	8.25 ± 11.61	1.462 ± 2.421	1.345 ± 0.319	25.5 ± 25.2
DiceCE	[-200, 1500]	✗	0.690 ± 0.144	8.36 ± 13.14	1.320 ± 1.653	1.653 ± 0.922	28.9 ± 30.1
DiceCE	[-1000, 2500]	✗	0.680 ± 0.134	5.68 ± 8.17	1.155 ± 1.529	1.409 ± 0.301	151.9 ± 249.2
GDL	[-158, 864]	✗	0.685 ± 0.184	6.50 ± 9.07	1.380 ± 2.102	1.381 ± 0.419	48.9 ± 81.0
GDL (ours)	[-200, 1500]	✗	0.709 ± 0.111	4.67 ± 6.98*	0.842 ± 1.081*	1.314 ± 0.280*	31.5 ± 44.2
GDL	[-1000, 2500]	✗	0.679 ± 0.152	7.34 ± 9.75	1.419 ± 1.932	1.711 ± 1.320	129.2 ± 171.6
Gbd	[-200, 1500]	✗	0.675 ± 0.142	5.80 ± 7.18	1.000 ± 1.002	1.319 ± 0.295	4.4 ± 4.8
Gch	[-200, 1500]	✗	0.672 ± 0.150	6.31 ± 8.28	1.098 ± 1.089	1.365 ± 0.306	16.2 ± 14.3
Ground truth	N/A	✓	0.805 ± 0.030	1.46 ± 1.09	0.373 ± 0.106	1.167 ± 0.217	0.0 ± 0.0
nnU-net	Self-configured	✓	0.665 ± 0.160	8.91 ± 10.65	1.479 ± 1.916	1.200 ± 0.269	3.0 ± 8.5
nnU-net	[-200,1500]	✓	0.621 ± 0.193	6.99 ± 8.33	1.308 ± 1.631	1.184 ± 0.260	0.2 ± 0.7
DiceCE	[-200, 1500]	✓	0.670 ± 0.082	4.94 ± 4.99	0.879 ± 0.551	1.274 ± 0.278	14.4 ± 29.9
GDL (ours)	[-200, 1500]	✓	0.673 ± 0.078	4.45 ± 4.59*	0.769 ± 0.448	1.195 ± 0.235	9.5 ± 13.3
Gbd	[-200, 1500]	✓	0.644 ± 0.116	5.13 ± 5.09	0.928 ± 0.680	1.278 ± 0.287	3.1 ± 3.2
Gch	[-200, 1500]	✓	0.630 ± 0.117	5.80 ± 5.71	1.045 ± 0.775	1.251 ± 0.255	2.6 ± 2.4

All metrics are mean ± stdev across patients. Bold indicates best in each column.

l minmax range for image intensity min-max normalization, *Post* post-processing, *HD* Hausdorff distance, *CD* mean symmetric chamfer distance, *CD_{tissue}* one-sided chamfer distance from calcification to cardiovascular tissues, *False pos* false positives, *mm* millimeters, *vx* voxels.

**p* < 0.05 between GDL (ours) and nnU-net.

Background mesh generation (GenBGMesh)

The post-processed segmentation and the input tissue mesh are combined together to perform calcification meshing, and Deep Marching Tetrahedra (DMTet)⁴² is at the core of our meshing algorithm. Similar to regular marching tetrahedra, DMTet generates an isosurface triangular mesh given a background tetrahedral mesh and its nodal SDF. Unlike regular marching tetrahedra, DMTet additionally performs simultaneous optimization of the background mesh node positions and the nodal SDF to refine the final mesh.

DMTet defines the typology of the isosurface based on the sign changes of the SDF and the location of the isosurface based on the magnitude of the SDF (Eq. (1)).

$$v'_{ab} = \frac{v_a \cdot SDF(v_b) - v_b \cdot SDF(v_a)}{SDF(v_b) - SDF(v_a)} \quad (1)$$

Here, v_i is the nodal position and $SDF(v_i)$ is the nodal SDF. Note that the node interpolation is only performed along edges where there is a sign change, i.e. between a positive SDF node and a non-positive SDF node ($SDF(v_a) \leq 0$ and $SDF(v_b) > 0$). This prevents division by 0.

Based on the DMTet definition, the background mesh is the main driver of the final output mesh topology. To enable contact surface matching, we designed our background mesh to be constrained by the original tissue surfaces, with a few minor adjustments to enable closed non-overlapping calcification surfaces (Fig. 3a). We first extracted the surface from the volumetric tissue mesh, and then combined it with a clean manifold offset layer to generate the input to TetGen⁴³. Following tetrahedralization, the elements inside the original tissue surface were removed, and fake tetrahedral elements were added to the boundary surfaces to ensure closed calcification surfaces at the end of the DMTet tessellation.

The success rate for background mesh tetrahedralization was 100%, mainly due to the nice tissue mesh quality from DeepCarve and our method of generating a clean manifold offset surface. The density of the tetrahedral elements was tuned to provide enough resolution, but the element quality

was not critical, as we additionally refined the DMTet output surfaces in the following steps.

DMTet optimization (DMTetOpt)

We evaluated the remaining meshing steps of C-MAC using the following metrics: (1) the success rate of tetrahedralization, (2) the spatial accuracy of the final mesh surface, and (3) the element quality of the final tetrahedra (Table 2). We also computed the same metrics for existing mesh editing approaches to showcase our improvements. For all methods, the general strategy was to first construct a triangular surface mesh using the post-processed segmentation and the tissue mesh, and then apply TetGen to tetrahedralize the inner volume.

For baselines, we applied mesh boolean operations from a wide variety of existing meshing libraries. For 5 of the 8 libraries, the success rate for the final tetrahedralization was below 10% (Table 2). For some, even the initial boolean operation failed completely. This is consistent with the notion that mesh boolean operations are generally unreliable, especially when coupled with complex input geometry and complex downstream tasks such as constrained tetrahedralization.

For mesh boolean operations with relatively high success rates, the issue still remained that the element qualities were heavily compromised along the mesh intersections. This is shown both qualitatively as thin triangular elements (Fig. 3c), and quantitatively as low Jacobian determinants (Table 2). Mesh boolean approaches had worse performance in essentially every category compared to our method.

As an additional baseline, we performed DeepCarve's node stitching operations as outlined in³⁸. The only metric that performed well for this method was the Jacobian determinant, which makes sense given the generally smooth and uniform elements (Fig. 3c). However, the method notably suffered from irregular node-stitching along the contact surfaces, which is further confirmed by the low numbers of merged nodes between the calcification and tissue meshes (Table 2). The irregular node-stitching also causes various other issues, such as the occasional undesired mesh

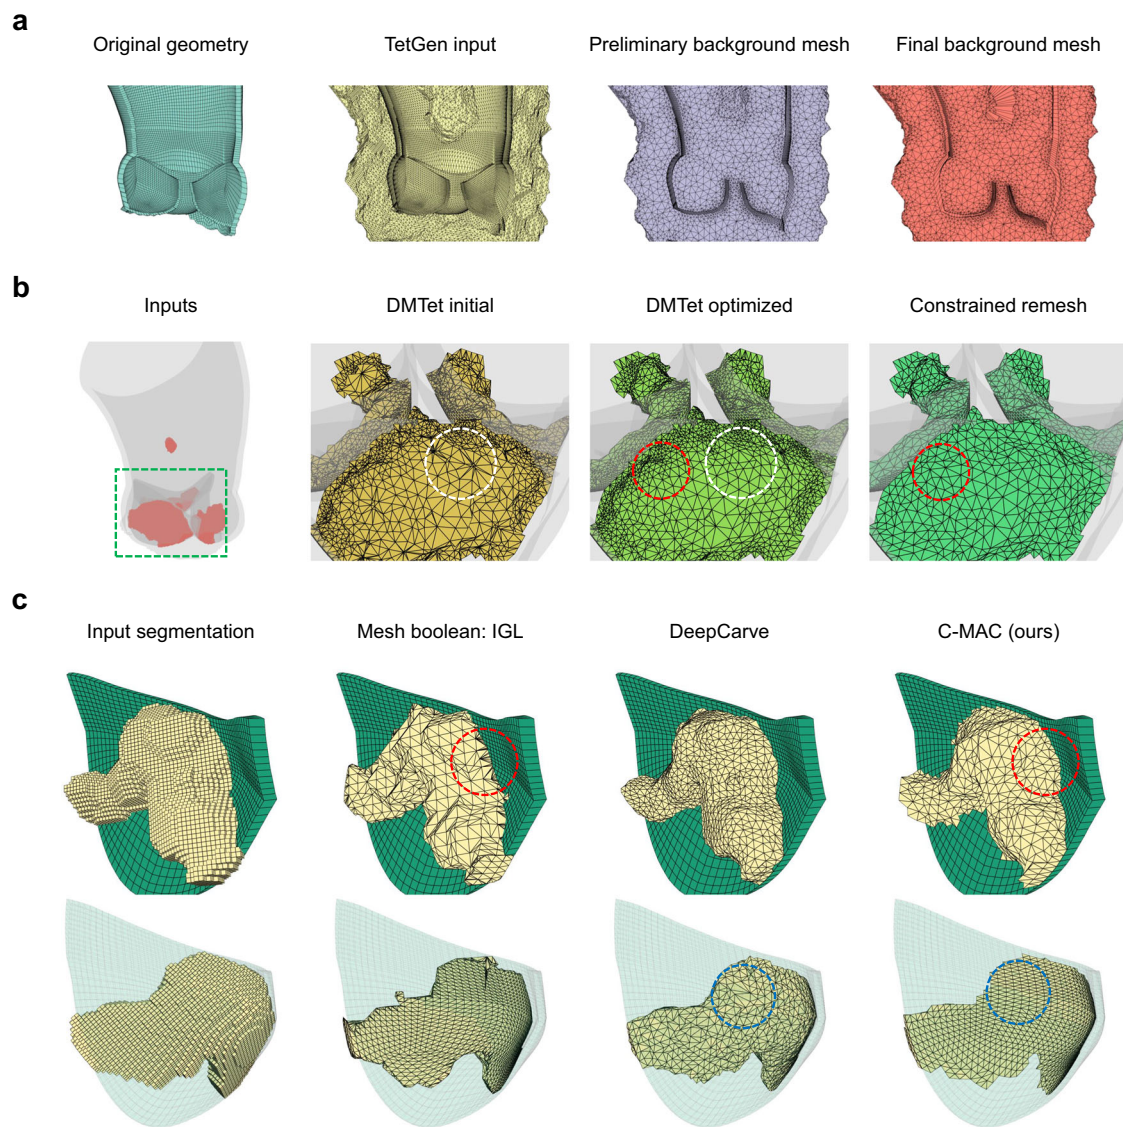


Fig. 3 | Qualitative evaluations of the meshing steps and baseline comparisons. **a** Illustration of the background mesh generation process. From left to right: patient-specific mesh of the aorta + aortic valve leaflets, TetGen input generated by combining the exterior surface and an offset surface from the original geometry, preliminary background mesh generated by TetGen and hollowing, and final background mesh after adding a “fake” vertex to the exterior surface elements of the preliminary mesh. All meshes are clipped at a viewing plane for visualization purposes. **b** The three main sequential steps for anatomically consistent surface

meshing. From left to right: initial inputs of patient-specific aorta + aortic valve leaflets (gray) and voxelgrid segmentation of calcification (red), initial DMTet mesh with raw sampled SDF, optimized DMTet mesh, and final remeshed surface. Green box indicates the viewing region, and colored circles indicate noticeable regions of improvement after each step. **c** Baseline comparisons for the final mesh quality. Yellow: calcification, green: aortic valve leaflet. Top: front view, bottom: back view. Colored circles indicate the noticeable regions of improvement from each baseline to C-MAC.

intersections and failure of TetGen due to the ill-shaped elements. Since it is difficult to automatically and robustly fix ill-formed meshes, our approach aims to generate high-quality meshes from the beginning.

For marching tetrahedra, a well-known shortcoming is the irregular shape and distribution of the triangular elements along the output mesh surfaces^{44,45}. This can be easily verified by applying DMTet without any optimization (Fig. 3b). Such irregularity noticeably degrades the final tetrahedra quality.

Thus, we applied DMTet optimization on both the background mesh node locations and the nodal SDF values to increase the output mesh quality while minimizing the deviation from the input segmentation. In addition, we modified the nodal SDF values for two edge cases to ensure clean transition elements along the boundary and contact surfaces.

DMTet optimization provided significant improvements for all metrics compared to DMTet without optimization (Table 2). Qualitatively,

it especially helped address the aggregation of sliver elements produced by the initial marching tetrahedra (Fig. 3b). In some cases, it was acceptable to stop the meshing algorithm after DMTet optimization, given its high-quality mesh outputs.

Constrained remeshing (ConstrainedRemesh)

Although DMTet optimization already provides high-quality outputs with robust performance, we can further improve the mesh quality and speed up the finite element simulations by additionally simplifying the output mesh with constrained remeshing.

The main remeshing operation is voronoi diagram-based vertex clustering⁴⁶. By our way of constructing the background mesh and DMTet optimization, the contact surfaces already have complete surface matching to the tissue mesh. To preserve this, we first split the output mesh into contact and non-contact surfaces, and only remesh the latter. For the non-

Table 2 | The effects of the calcification surface meshing algorithms on the final tetrahedral mesh

Mesh technique	Success (%) ↑	Dice ↑	HD 95% (mm) ↓	CD (mm) ↓	Merged nodes ↑	J ↑
Bool: pymadcad	0.0	N/A	N/A	N/A	N/A	N/A
Bool: OpenSCAD	0.0	N/A	N/A	N/A	N/A	N/A
Bool: vtk	0.0	N/A	N/A	N/A	N/A	N/A
Bool: blender	3.7	N/A	N/A	N/A	N/A	N/A
Bool: cgal	7.4	0.753 ± 0.140	1.366 ± 0.518	0.316 ± 0.126	333 ± 360	0.332 ± 0.004
Bool: igl	70.4	0.811 ± 0.061	1.507 ± 0.446	0.364 ± 0.074	2480 ± 1552	0.449 ± 0.043
Bool: corefinement	74.1	0.811 ± 0.059	1.503 ± 0.435	0.365 ± 0.073	2553 ± 1545	0.453 ± 0.043
Bool: cork	75.2	0.812 ± 0.056	1.480 ± 0.374	0.367 ± 0.067	2827 ± 1762	0.456 ± 0.039
DeepCarve	96.3	0.763 ± 0.060	2.405 ± 2.778	0.559 ± 0.251	630 ± 231	0.606 ± 0.020 [†]
DMTet (no opt)	100.0	0.830 ± 0.045	1.000 ± 0.000	0.247 ± 0.020	2908 ± 1777	0.549 ± 0.013
DMTet (no opt) → ConstrainedRemesh	100.0	0.838 ± 0.043	1.000 ± 0.000	0.245 ± 0.021	2898 ± 1773	0.585 ± 0.011
DMTetOpt	100.0	0.849 ± 0.034	1.000 ± 0.000	0.238 ± 0.021	2907 ± 1777	0.584 ± 0.007
DMTetOpt → ConstrainedRemesh (C-MAC)	100.0	0.853 ± 0.033 [†]	1.000 ± 0.000 [†]	0.239 ± 0.022 [†]	2898 ± 1774 [†]	0.593 ± 0.010*

"Bool" refers to the mesh boolean difference operation between the original calcification surface and the aorta + aortic valve surface. All metrics were obtained using 10 repeated trials to account for the occasional failure of TetGen. We used the post-processed ground truth segmentation as the initial segmentation. Spatial accuracy metrics were measured against the initial segmentation, so we are strictly measuring the error induced by the meshing steps. Merged nodes indicate the number of the final calcification mesh nodes within 1e–3 mm from the CV tissue mesh nodes. All metrics are calculated using samples with non-empty targets only. All metrics are mean ± stdev across patients. Bold indicates best in each column.

**p* < 0.05 between C-MAC and "Bool: corefinement"; [†]*p* < 0.05 between C-MAC and DeepCarve.

contact surfaces, we further ensure that the remeshed surfaces can be combined with the original contact surfaces by preserving the border vertices during the clustering steps.

Quantitatively, constrained remeshing by itself markedly improves the mesh quality, similar to DMTet optimization (Table 2). However, the output of DMTet optimization can be further improved using constrained remeshing, whereas constrained remeshing breaks the DMTet surface correspondence with the background mesh and disables further DMTet optimization. Constrained remeshing can also occasionally generate non-manifold surfaces due to the independent processing of contact and non-contact surfaces. We are still able to reach 100% tetrahedralization success rate because we can easily store and select the optimally processed mesh at lower iterations of constrained remeshing by simply performing tetrahedralization on all intermediate meshes. This is possible because both constrained remeshing and calcification tetrahedralization take a fraction of a second to perform, and the outputs are not memory intensive.

As shown in Table 2, the condition that includes both DMTetOpt and ConstrainedRemesh results in the best overall meshing performance, both in terms of spatial accuracy and element quality. The final version of C-MAC is able to combine the best of both worlds of the baseline meshing methods, i.e. (1) establish complete contact surface matching, similar to mesh boolean operations, and (2) provide great element quality along the non-contact surfaces, similar to DeepCarve’s approach. (Fig. 3b, c). In addition, C-MAC ensures that any node residing on any of the tissue surface’s planes will be coincident with an existing tissue mesh node, which prevents any potential issues with minute mesh intersections that can introduce cumbersome errors during simulations.

Combining all of the steps, we performed the final end-to-end evaluation of C-MAC (Table 3). Even with the ground-truth segmentation as the initial input, the Dice score is relatively low for the final output. This can be partially attributed to accumulated errors from the post-processing and meshing steps, but the small volume of ground-truth calcification also amplifies small errors in segmentation to large errors in Dice. Using the “GDL (ours)” model for the initial segmentation, all performance metrics are slightly worse than Table 1 as expected, due to the additional meshing operations. Nonetheless, the final calcification meshes have a mean surface error of <1 mm from the ground-truth segmentation, and are qualitatively faithful to the input image (Fig. 2).

Run-time

The approximate run-time to generate the entire calcified tissue geometry (i.e. DeepCarve + C-MAC) is ~1 min per 3D image. The three most time-consuming tasks are (1) closing operations for segmentation post-processing (~10 s), (2) TetGen for background mesh generation (~15 s), and (3) DMTet optimization (~20 s). Most other individual processes typically take <1 s. All times were measured on a single NVIDIA RTX3080Ti laptop workstation. The algorithms are fully automated with no human annotations for any of the meshing steps.

Solid mechanics simulations

We performed two sets of solid mechanics simulations to assess different aspects of our final C-MAC meshes. First is valve opening simulations, which demonstrate our method’s unique ability to set up large-scale simulations. Second is TAVR stent deployment simulations, which demonstrate our meshes’ robustness to complex simulation environments (Fig. 4). Simulations were performed using Abaqus (3DS, Dassault Systèmes, Paris, France).

For valve opening, a uniform pressure of 8 mmHg was applied to the ventricular surface of the leaflets. The pressure was increased compared to previous work⁴⁷ to represent elevated aortic valve pressure gradient in aortic stenosis. Valve opening simulations were performed across all 35 patients with *fully automated* simulation setups. This was possible because we maintained the mesh correspondence of tissue meshes, and all boundary and loading conditions were defined with respect to the tissue mesh elements. This allowed us to both qualitatively confirm the consistency of our results across a large number of samples, and also quantify our observations across the entire test-set population.

For TAVR stent deployment, we employed a validated simulation setup previously developed using manually reconstructed patient models and a 26 mm, self-expandable, first-generation Medtronic CoreValve device (Medtronic, Minneapolis, MN, USA)^{48,49}. Briefly, the crimped TAVR stent with an exterior diameter of 6 mm⁴⁷ was aligned coaxially within the aortic root and centered into the aortic annulus, following the manufacturer’s recommendations⁵⁰. The stent was deployed inside the aortic root by axially moving the cylindrical sheath towards the ascending aorta. We selected 10 patient-specific anatomies that are suitable for this device size according to the manufacturer’s recommendations⁵⁰, and performed the deployment simulations while only altering the patient-specific anatomical meshes.

Table 3 | Final C-MAC mesh metrics evaluated against the ground-truth calcification segmentation

Experiment condition		Non-empty target segmentation				Empty target
Initial segment	Post-process	Dice ↑	HD 95% (mm) ↓	CD (mm) ↓	Merged nodes ↑	False pos (vx) ↓
Ground truth	✓	0.683 ± 0.036	1.703 ± 1.023	0.525 ± 0.110	2899 ± 1774	N/A
GDL (ours)	✓	0.630 ± 0.074	4.560 ± 4.519	0.852 ± 0.423	2295 ± 998	7.5 ± 10.8

These metrics are measuring the error induced by the entire C-MAC pipeline including the initial segmentation, segmentation post-processing, and meshing steps. Similar to Table 2, we used 10 repeated trials to obtain these metrics. All metrics are mean ± stdev across patients.

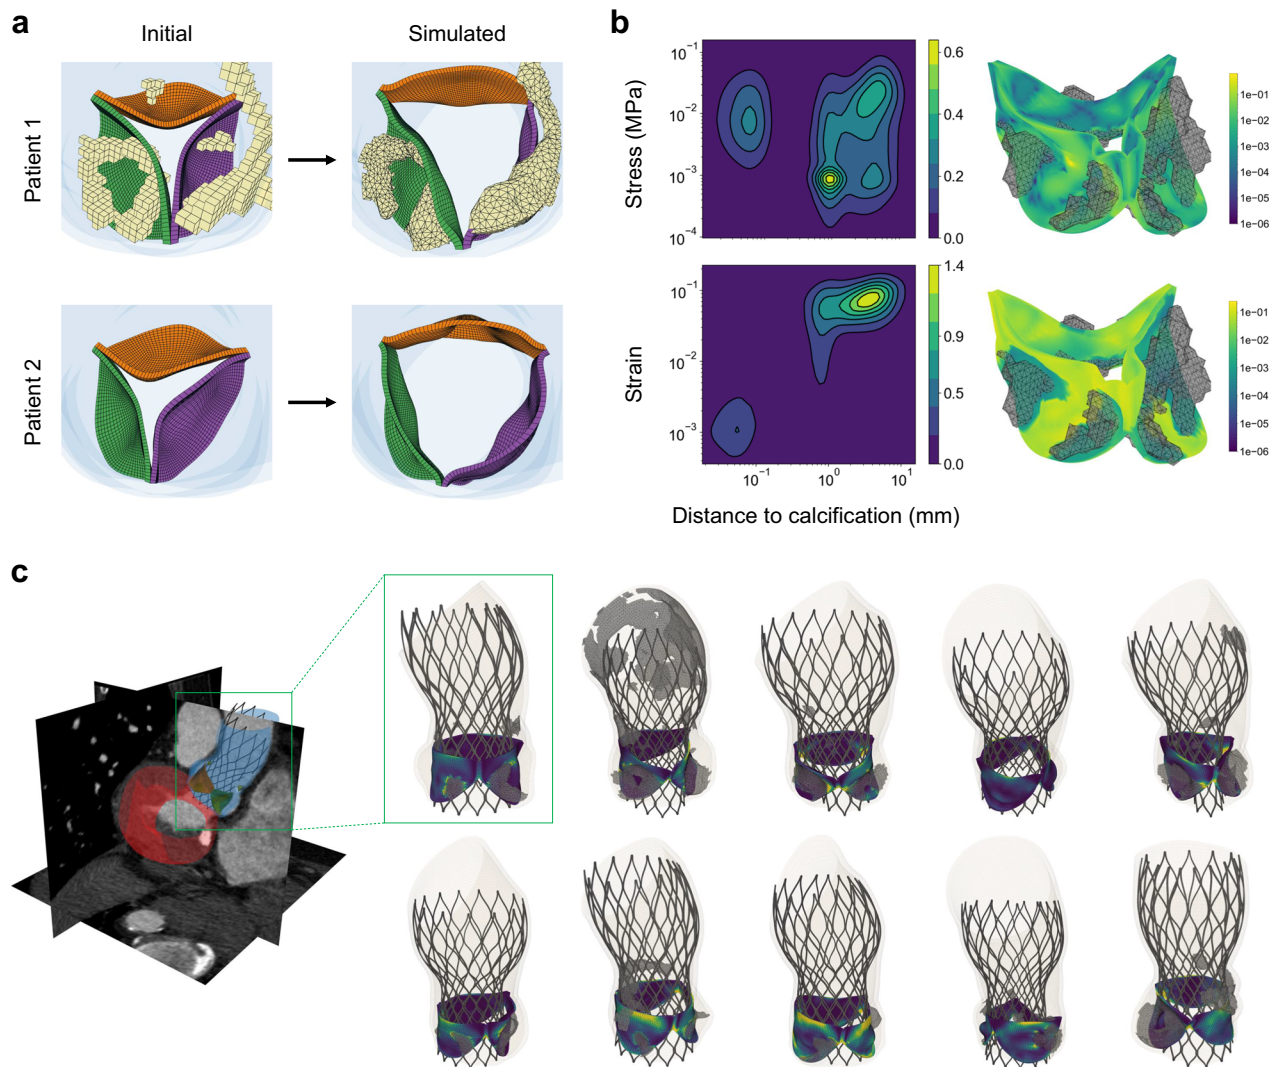


Fig. 4 | Simulation outputs using the final reconstructed meshes. a Valve opening simulations demonstrate the effects of calcification on the final leaflet positions. Yellow is the ground-truth calcification, left is the input valve geometry predicted by DeepCarve, and right is the deformed geometry after finite element analysis. Movement is clearly restricted near calcified regions. **b** Stress (top) and strain (bottom) analyses from valve opening simulations. Left is the Gaussian KDE plot of

stress/strain vs. distance to calcification from the aggregate of 35 test-set patient simulations. Right is one test-set patient with stress/strain overlaid with the valve leaflets, plus the ground-truth calcification (gray) for reference. **c** TAVR stent deployment simulation results. Left: image and simulated geometry overlay. Right: maximum principal stress magnitudes plotted on the aortic valve leaflets for 10 different test-set patients.

From both simulations, we extracted two main pieces of information. First is the effect of the calcification on the leaflet movement. The leaflet movement is clearly restricted near the calcified regions (Fig. 4a). This can be attributed to the high stiffness of the calcification and the aortic wall, which incur both intrinsic resistance from the leaflet elements and external reactive forces from the physical connections to the surrounding structures. The second set of information was the stress/strain values. The stress/strain values are significantly influenced by the location of the calcification. The simulations were performed using completely auto-generated meshes from

DeepCarve and C-MAC, and the stress/strain results correspond well to the location of the ground-truth calcification.

For valve opening, we further performed Gaussian kernel density estimation (KDE)^{51,52} on the aggregated data points of all test-set simulations to estimate the joint distribution of the leaflet nodes' stress/strain values vs. their distances to the nearest calcification (Fig. 4b). The KDE plots (1) confirm our qualitative observations that nodes closer to calcification have low strain, and (2) reveal new insights that the leaflet stress may exhibit a complex pattern at 1–10 mm away from the calcification. More detailed

analyses of the simulation results may lead to important insights on the growth mechanism of calcification as well as the intricate leaflet behavior associated with the positioning of the calcification within the aortic root.

For stent deployment, the key finding was the unmatched robustness of C-MAC's outputs in the complex multi-body simulation. Using the calcification meshing algorithm from DeepCarve, 4 out of the 10 simulations failed due to calcification element distortions. Using C-MAC, all 10 simulations successfully ran to completion. The simulation results in Fig. 4c concur with previous observations that the calcified regions are associated with higher levels of stress. This further confirms the high robustness of our meshing algorithm and the benefit it provides for complex downstream tasks.

Discussion

Our extensive evaluations suggest that C-MAC is a robust fully automated solution for incorporating aortic calcification mesh onto an existing tissue mesh. (1) Our deep learning segmentation model provides a good initial calcification segmentation. (2) Our segmentation post-processing algorithm enhances the calcification-to-tissue anatomical consistency while slightly compromising the segmentation volume accuracy. (3) Our meshing algorithm accurately converts the input voxelgrid segmentation into surface meshes with great element quality along both contact and non-contact surfaces, while preserving the original tissue mesh topology. Furthermore, the non-deep-learning components of our algorithm can be generally applied to any problem that involves attaching voxelgrid segmentation to a clean manifold surface mesh. We demonstrate this versatility with an additional experiment of attaching calcification to the mitral valve and LV (Supplementary Fig. 3).

However, our methods and analyses also possess several limitations. C-MAC has a relatively long processing time of ~1 min compared to ~1 s of deep learning-based methods. Our post-processing algorithm sacrifices volume overlap to achieve anatomical consistency. We were unable to validate our simulation accuracy with clinical data. We were also only able to validate our simulations with one commercial finite element solver (Abaqus) on a limited set of physical problems involving the aortic valve. We will elaborate on these limitations and potential future solutions for the rest of this section.

Unlike the original DMTet paper⁴², we did not train a separate neural network to expedite the DMTet optimization process. Instead, we opted for a per-image optimization approach. The main reason for this choice was the long processing time of TetGen for the background mesh generation. Unlike⁴² where the background mesh was a fixed lattice of regularly spaced tetrahedra, our meshing workflow requires a background mesh that contains the patient-specific tissue geometry in the inner volume. Since the element density and configuration can change drastically based on the leaflet conformation, we could not reliably construct the background mesh without repeating TetGen for each scan separately. This poses a significant restriction on our training scheme since repeating TetGen for every augmented training sample would significantly increase the training time (from ~1 s to ~15 s per sample) or storage (~20 MB per background mesh for hundreds of thousands of samples).

Furthermore, even if we were able to successfully train a deep learning model for the DMTet optimization process, it would only reduce the total run-time from ~1 min to ~40 s because the other major time-consuming processes have nothing to do with DMTet optimization. Therefore, we decided to simply use the proposed per-image optimization approach. This approach has an additional minor benefit that we can slightly randomize the output with different initialization for DMTet optimization if we are unsatisfied with our output.

Our current approach for segmentation post-processing has a trade-off between calcification segmentation accuracy and anatomical consistency because we fix DeepCarve's cardiovascular structures while extending the calcification segments to remove the spatial gaps. We aimed to achieve the best possible solution for this approach using an adaptive closing kernel, but the inherent reduction in volume overlap is likely unavoidable. An

interesting future direction would be to jointly optimize the tissue mesh and the calcification segments to deform the tissue mesh towards the predicted calcification, instead of simply extending the calcification. This approach is technically more challenging because it requires simultaneous optimization of multiple metrics, such as the calcification accuracy, tissue surface accuracy, tissue mesh element quality, and also requires a robust formulation for the selective attraction between only the contact surfaces of the tissue and calcification meshes.

Another important consideration for C-MAC is the usability of the final meshes. Since we optimize for good element quality using DMTet optimization, the key remaining consideration is the manifoldness of the output. One of the main strengths of marching tetrahedra is its ability to generate topologically consistent outputs due to its unambiguous tessellation of the tetrahedral lattice. The same principle applies to the raw non-optimized DMTet output with our custom background mesh and calcification segmentation, which means the output should be manifold surfaces without any modifications. The main issue with the non-optimized DMTet output is that TetGen can occasionally fail due to the highly irregular triangular elements from the raw node-sampled SDF.

We proposed DMTet optimization to maintain the original manifoldness of DMTet while improving the element quality. The element quality improvements are straightforwardly obtained using \mathcal{L}_{DMTet} (Eq. (8)), which optimizes for various desired triangle properties. For the manifoldness, we first argue that `ModifyEdgeCases` (Eq. (6)) does not alter the manifoldness of the original DMTet output. From Eq. (7): (1) For any tetrahedron with an SDF sign-change across the fake node, the output surface will always be a tissue surface element. DeepCarve's tissue meshes all have manifold surfaces by design, so the DMTet output at contact surfaces will also be manifold. (2) For any tetrahedron with an SDF sign-change across a boundary node, the output surface will always only *expand* its borders along the boundary surface, and the amount of expansion will at worst cause coincident faces. This means no self-intersections will be introduced by our modification, so the DMTet output will be manifold.

For DMTetOpt's δ optimization (Algorithm 4), ΔSDF is constrained to only change each nodal SDF's magnitude while keeping its original sign, which prevents any topological changes. The only remaining potential concern is ΔV_{bg} . Given that \mathcal{L}_{DMTet} is only calculated for tetrahedra with nodal SDF sign-changes, the ΔV_{bg} optimization will also only be applied to those elements' nodes. In this case, a potential way for ΔV_{bg} to cause non-manifoldness is to cause self-intersections by heavily displacing the background mesh nodes near the segmentation isosurface. This is strongly discouraged by many of our design choices, such as (1) \mathcal{L}_{DMTet} 's surface smoothness constraints and ΔV_{bg} regularization and (2) our strategy of updating the nodal SDF at the new node coordinate for each DMTetOpt iteration. Empirically, DMTetOpt consistently generates high-quality manifold meshes, leading to a TetGen success rate of 100%.

Clinical validation of solid mechanics simulations is extremely challenging because post-TAVR CTs are not routinely collected and the ground-truth stress/strain values for in vivo anatomical structures are nearly impossible to obtain. The solid mechanics simulations in this work instead illustrate two other important points: (1) our generated meshes can handle complex solid interactions such as TAVR stent deployments and (2) the high-quality mesh output with mesh correspondence allows for large-batch simulations, even with the addition of calcification. Further clinical validation of our simulations will be performed with flow-based simulations in future studies using the more routinely collected Doppler echodata⁵³.

In this work, we propose an automated end-to-end image-to-mesh solution for incorporating calcification onto an existing tissue geometry. Our solution includes (1) a segmentation algorithm that encourages anatomical consistency between the predicted calcification and tissue mesh and (2) a meshing algorithm that establishes complete contact surface matching between the two geometries while maintaining the original tissue mesh topology. Our technique allows for large-batch simulations due to the preservation of inter-patient mesh correspondence of the tissue geometry. We demonstrated the viability of such downstream tasks using solid mechanics

simulations. Our method may help accelerate the development and usage of complex physics-driven simulations for cardiovascular applications.

Methods

Data acquisition and preprocessing

We used the same dataset as DeepCarve with some minor modifications. The original dataset consisted of 80 CTA scans and the corresponding labels for the cardiovascular structures and calcification. 14 of the 80 CT scans were from 14 different patients in the training set of the MM-WHS public dataset⁵⁴. The remaining 66 scans were from 55 IRB-approved TAVR patients at Hartford hospital. Approval of all ethical and experimental procedures and protocols was granted by the Institutional Review Board (IRB) under Assurance No. FWA00021932, and IRB-Panel_A under IRB No. HHC-2018-0042.

We used more than one time point for some of the Hartford patients. Of the 80 total scans, we used 35/10/35 scans for training/validation/testing. We ensured that the testing set had no overlapping patients with the training/validation sets. All evaluations were performed with the 35 test-set cases. All patients had tricuspid aortic valves and varying levels of calcification. Some scans had no aortic calcification. Across all patients in the dataset, the total calcification volume was $851.44 \pm 966.27 \text{ mm}^3$, and calcification was most frequently found closest to the aorta, non-coronary cusp, right coronary cusp, left coronary cusp, and LV, in that order. The histograms of calcification volume, each segmentation voxel's closest structure, and distances of each voxel to the closest structure are provided in Supplementary Fig. 1.

The calcification segmentation was obtained using adaptive thresholding and manual post-processing. The original dataset included calcification that is within some distance away from any of the pointclouds. For this work, we further processed it to only include segments around the aorta. This was mainly to omit calcification along the mitral valve, which would require accurate mitral valve leaflet geometry for us to properly apply our meshing techniques.

Both Gbd and Gch require extra training labels extracted from the original calcification segmentation. Gbd requires voxelgrid SDF, which we obtained using isosurface extraction and point-to-mesh distance from each voxel position. Gch requires calcification pointclouds, which we obtained using isosurface extraction and surface point sampling.

For preprocessing, we rescaled all images and labels to an isotropic spatial resolution of 1.25 mm^3 and cropped them to 128^3 voxels. The default crop center was the center of the bounding box of ground-truth cardiovascular tissue pointclouds. For translation augmentation, the crop center was offset by a random sample from a 3D Gaussian ($\text{stdev} = \text{crop_width}/3$), and the offset amount was capped to prevent any labeled structures being outside the cropped region. The images were not pre-aligned with any additional transforms.

Image intensities were first clipped to [lower bound, upper bound], and further min-max normalized to [0, 1]. The lower and upper bounds were fixed for DeepCarve at $[-158, 864]$ Hounsfield Units (HU) and chosen between three different ranges for calcification segmentation, as specified in the results tables (Table 1).

Deep learning model training details

Dice (a.k.a. DSC)^{55–57} and cross entropy (CE)⁴⁰ are commonly used for training segmentation models. Dice has several advantages over CE, such as being more robust to unbalanced targets and being able to directly optimize a common evaluation metric, but it is ineffective at handling empty targets. A common workaround is to instead optimize the combined DiceCE loss⁵⁸, which can be defined for binary segmentation tasks as

$$\mathcal{L}_{\text{DiceCE}}(y, \hat{y}) = 1 - 2 \frac{\sum_{i=1}^N y_i \hat{y}_i}{\sum_{i=1}^N y_i + \sum_{i=1}^N \hat{y}_i + \epsilon} + \lambda \left[\frac{1}{N} \left(\sum_{i=1}^N y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i) \right) \right] \quad (2)$$

where y_i is the target and \hat{y}_i is the predicted value at each i^{th} voxel, λ is a weighting hyperparameter, and ϵ is an error term that prevents division by 0. To prevent numerical issues with $\log(0)$, the log output is set to have a minimum bound at -100^{59} .

Generalized Dice Loss (GDL) is a variation of the Dice loss that handles multi-class segmentation with a per-class weighting term that mitigates biases towards larger segments⁵⁶. For a binary segmentation task, we can treat the background as its own class, in which case

$$\mathcal{L}_{\text{GDL}}(y, \hat{y}) = 1 - 2 \frac{\sum_{c=1}^2 w_c \sum_i y_{ci} \hat{y}_{ci}}{\sum_{c=1}^2 w_c \sum_i (y_{ci} + \hat{y}_{ci})} \quad (3)$$

where $w_c = \frac{1}{(\sum_i y_{ci} + \epsilon_c)}$. We set $\epsilon_0 = -100$ and $\epsilon_1 = 100$ via hyperparameter tuning, where $c = 0$ is background and $c = 1$ is calcification. The ϵ_c helps prevent division by 0 and helps maintain a smooth Dice-like behavior for both empty and sparse targets. Empirically, GDL performed better than DiceCE for our task, so it was chosen as the main region-based loss for the rest of our experiments.

All variants of distribution-based losses (i.e. CE) and region-based losses (i.e. Dice) penalize each voxel prediction independently, and fail to capture the magnitude of error from the original segmentation boundary. The boundary loss was proposed to combat this phenomenon⁶⁰:

$$\mathcal{L}_{\text{Gbd}}(y, \hat{y}) = \mathcal{L}_{\text{GDL}}(y, \hat{y}) + \lambda \sum_i -\text{SDF}(y)_i \hat{y}_i \quad (4)$$

where $\text{SDF}(y)$ is the signed distance function (SDF) calculated using the ground-truth segmentation y . By this definition, SDF is assumed to be positive on the inside and negative on the outside of the $y = 0.5$ isosurface. The weighting hyperparameter λ is adjusted based on training epochs, similar to⁶⁰. In our case, $\lambda = 0$ at epoch 0 and was linearly increased to $\lambda = 1000$ from epochs 1000–2000. We clipped the SDF from -3 to 3 for numerical stability.

Finally, instead of relying on the integral approximation approach as did⁶⁰, we evaluated the effectiveness of directly optimizing for the symmetric chamfer distance⁶¹ between the predicted and target segmentations:

$$\mathcal{L}_{\text{Gch}}(y, \hat{y}) = \mathcal{L}_{\text{GDL}}(y, \hat{y}) + \lambda \left(\frac{1}{|A|} \sum_{a \in A} \min_{b \in B} \|a - b\|_2^2 + \frac{1}{|B|} \sum_{b \in B} \min_{a \in A} \|b - a\|_2^2 \right) \quad (5)$$

where A and B are pointclouds sampled on the segmentations' extracted surfaces. We used DM Tet to extract the surface from the predicted segmentation, which makes the operation differentiable with respect to the segmentation model parameters⁴². $\lambda = 0$ at epoch 0 and linearly increased to $\lambda = 0.1$ from epochs 1000–2000.

All models were trained for 4000 epochs with the Adam optimizer, at a learning rate of $1e-4$. For data augmentation, we used random translation and b-spline deformation.

We also compared our models' performance against nnU-net, a popular self-configuring DL segmentation method. For fair comparisons, we trained the nnU-net with the same training/validation/test splits as our own models, which effectively discarded nnU-net's default ensembling operations.

For the self-configured nnU-net, we set the channel name to "CT", performed the default "nnUNetv2_plan_and_preprocess" with the "verify_dataset_integrity" option enabled, and trained and predicted with the "3d_fullres" configuration. For the nnU-net with I minmax $[-200, 1500]$, we kept everything consistent other than changing the channel name to "noNorm" and using pre-normalized CT scans as the initial nnUNet_raw data.

Segmentation post-processing details

Algorithm 2. Post-processing of ca2 segmentation

```

1: function PostProcessCa2Seg( $y_0, M_{tissue}$ )
2:    $\{y_1, y_2, y_3\} \leftarrow \text{GroupSegByLeaflets}(y_0, M_{tissue})$ 
3:   while not satisfied do
4:     for  $y_i$  in  $\{y_i\}$  do
5:        $\hat{y}_{tissue} \leftarrow \text{FilterTissueSeg}(y_i, M_{tissue})$ 
6:        $\hat{y}_i \leftarrow \text{AdaptiveClose}(y_i, \hat{y}_{tissue})$ 
7:        $y_i \leftarrow y_i + \hat{y}_i$ 
8:      $y_{ca2} \leftarrow \text{SubtractAndFilterSeg}(\{y_i\}, M_{tissue})$ 
9:   return  $y_{ca2}$ 

```

S_i is assigned in-place, so S_i is updated with each iteration.

(GroupSegByLeaflets) From the initial calcification segmentation, we first subtracted away the regions overlapping with the cardiovascular tissue mesh to split any calcification that may extend over multiple leaflets. To perform voxelgrid subtraction, we converted the tissue mesh to voxelgrid segments using the image stencil operation⁶² at 3x resolution of the initial segmentation, which ensured that the thin leaflet geometries were accurately represented in the voxelgrid format. Then, we subtracted the tissue segments from the initial calcification segments, and then assigned each island of the split calcification to the closest aortic valve leaflet, measured by the mean symmetric chamfer distance.

(FilterTissueSeg) Using the grouped calcification segments and the predicted tissue mesh, we extracted tissue segments that are within some distance away from the calcification. To achieve this, we first converted the tissue mesh to tissue segments, and only kept the segments intersecting with dilated calcification segments.

The kernel for morphological operations was $7 \times 7 \times 7$ voxels and adaptive to the tissue mesh. First, at every voxelgrid cell, we defined an anisotropic kernel that is shaped to be ellipsoidal with the major principal axis pointing in the direction of the closest tissue mesh node’s surface normal. Then, we linearly combined it with an identity kernel, where the combination weights were determined by the voxelgrid cell’s distance away from the tissue mesh. For FilterTissueSeg, we performed the calcification dilation twice with this kernel at 0.8 times the kernel weights.

(AdaptiveClose) We performed one iteration of adaptive morphological closing on the grouped calcification segments and the corresponding filtered tissue segments. We used the same adaptive kernel described in FilterTissueSeg with the original kernel weights.

(SubtractAndFilterSeg) Due to the limited resolution of the voxelgrid representation, directly meshing from the adaptively closed calcification segments results in many undesired protrusions in the final mesh. To avoid this, we filtered the combined segments by subtracting away the tissue segments and performing volume-based island removals. For the subtraction, we tripled the segmentation resolution, similar to GroupSegByLeaflets. The final segmentation for downstream meshing (y_{ca2}) therefore has a 3x voxelgrid resolution compared to the original image.

Background mesh generation details

Algorithm 3. Background mesh generation

```

function GenBGMesh( $M_{tissue}$ )
   $S_{aorta} \leftarrow \text{ExtractAortaSurf}(M_{tissue})$ 
   $S_{offset} \leftarrow \text{OffsetSurf}(M_{tissue})$ 
   $M_{prelim} \leftarrow \text{TetGenAndHollow}(S_{aorta}, S_{offset})$ 
   $M_{bg} \leftarrow \text{CreateFakeElems}(M_{prelim})$ 
  return  $M_{bg}$ 

```

(ExtractAortaSurf) Since we were focusing on aortic calcification, the original tissue surface consisted of the surface elements of the aorta and the aortic valve leaflets, which we can extract using standard meshing libraries⁶².

(OffsetSurf) For the outer bound of the background mesh, we generated an offset surface that includes all areas within 10 voxel spacings away from the aortic surface. The extracted aortic surface and the offset surface were combined by addition, which was easily doable because the

offset surface and the extracted tissue surface are clean manifold surfaces with no intersections by construction.

(TetGenAndHollow) The merged surfaces were processed by TetGen for constrained tetrahedral meshing, and then the elements inside the aortic surface were removed.

(CreateFakeElems) Lastly, we created “fake” tetrahedral elements by adding a “fake” node to all background boundary surfaces, i.e. the aortic surface and the offset surface elements.

TetGenAndHollow and CreateFakeElems are crucial for accurate meshing. The rationale will be explained further in the following subsection.

DMTet optimization details

Algorithm 4. DMTet optimization

```

function DMTetOpt( $y_{ca2}, M_{bg}$ )
  ( $V_{bg}, E_{bg}$ )  $\leftarrow M_{bg}$ 
   $\delta \sim \mathcal{U}_n(-1e-3, 1e-3)$ 
  for  $n_{opt}$  do
    ( $\Delta V_{bg}, \Delta SDF$ )  $\leftarrow \delta$ 
     $\tilde{V}_{bg} \leftarrow V_{bg} + \Delta V_{bg}$ 
     $\tilde{M}_{bg} \leftarrow (V_{bg}, E_{bg})$ 
     $SDF \leftarrow \text{LinearMapAndInterp}(y_{ca2}, \tilde{V}_{bg})$ 
     $SDF \leftarrow \text{ModifyEdgeCases}(SDF) \triangleright \text{Eq. (6)}$ 
     $SDF \leftarrow SDF + \Delta SDF$ 
     $S_{DMTet} \leftarrow \text{DMTet}(\tilde{SDF}, \tilde{M}_{bg}) \triangleright \text{Eq. (1)}$ 
    ( $V, E$ )  $\leftarrow S_{DMTet}$ 
     $\delta \leftarrow \text{Adam}(\delta, \mathcal{L}_{DMTet}(V, E, \Delta V_{bg})) \triangleright \text{Eq. (8)}$ 
   $S_{ca2} \leftarrow \text{CleanMesh}(S_{DMTet})$ 
  return  $S_{ca2}$ 

```

(LinearMapAndInterp) We first converted the predicted segmentation to a simplified voxelgrid SDF via linear mapping $SDF(\hat{y}) = 2\hat{y} - 1$, which means $\hat{y}_i \in \{0, 1\} \rightarrow SDF(\hat{y})_i \in \{-1, 1\}$. Then, we trilinearly interpolated the SDF at each background mesh node.

To demonstrate the need for further processing, let us consider applying marching tetrahedra with the default $SDF_i \in [-1, 1]$ and the preliminary background mesh. Due to the smooth transition of the nodal SDF magnitudes, none of the resulting surface mesh would match exactly with the original background mesh nodes. This is counterproductive, as our desired output is a surface mesh with coincident nodes and edges along the contact surfaces.

(ModifyEdgeCases) To enforce complete contact surface matching, we make two modifications to the nodal SDF after the initial interpolation:

$$SDF_i = \begin{cases} 0 & \text{if background boundary node} \\ -1e12 & \text{if fake node} \\ SDF_i & \text{else} \end{cases} \quad (6)$$

Then, from Eq. (1), this results in the following node assignments:

$$v'_{ab} \approx \begin{cases} v_a & \text{if background boundary node} \\ v_b & \text{if fake node} \\ v'_{ab} & \text{else} \end{cases} \quad (7)$$

Recall that (1) all elements inside the tissue volume were removed via TetGenAndHollow and (2) the fake node was added to the background boundary surface to generate fake tetrahedral elements via CreateFakeElems. Together with Eq. (7), this means that any SDF sign change involving a boundary node will always result in that node being included as a part of the DMTet output. This is exactly our intended output with complete contact surface matching with the original boundary elements.

Applying Eq. (6) is enough for establishing complete contact surface matching. However, another known problem of marching tetrahedra is the irregular elements of the output surface⁶³. Although some related solutions exist, our task also requires that the contact surfaces be preserved during mesh processing. DMTet optimization is the first step of our two-part solution. Constrained surface remeshing is the optional second step.

Similar to the original DMTet¹², we optimized (1) the deformation of the background mesh and (2) the offset from the initial interpolated nodal SDF. Unlike the original paper, we do not train a deep learning model for the prediction. Instead, we perform the optimization for each inference target independently. The rationale for this choice is described in the Discussion section. During the optimization process, we minimize the output’s deviation from the original segmentation by penalizing the background mesh deformation and prohibiting sign changes in the nodal SDF.

The overall idea is to improve the mesh quality while minimizing its deviation from the original segmentation. The loss is a combination of Laplacian smoothing, edge length, edge angle, and deformation penalty (Eq. (8)). The first three losses help improve the overall mesh and individual element qualities, while the deformation penalty helps minimize surface shrinking. Eq. (6) is applied after each optimization step, so only the non-contact nodes’ SDF values are optimized during this process.

$$\begin{aligned} \mathcal{L}_{DMTet}(V, E, \Delta V_{bg}) = & \lambda_0 \sum_{v_i \in V} \left\| \frac{1}{|\mathcal{N}(v_i)|} \sum_{v_j \in \mathcal{N}(v_i)} v_i - v_j \right\|_2 \\ & + \lambda_1 \left[\sum_{v_i \in V} \sum_{v_j \in \mathcal{N}(v_i)} \left(\|v_i - v_j\|_2 - \epsilon \right)^2 \right]^{\frac{1}{2}} \\ & + \lambda_2 \left[\sum_{a \in A(V, E)} (a - \alpha)^2 * \sigma(a) \right]^{\frac{1}{2}} \\ & + \lambda_3 \sum_{v_i \in V_{bg}} \|\Delta v_i\|_2 \end{aligned} \quad (8)$$

(V,E): vertices and edges of the DMTet output. \mathcal{N} : neighboring nodes. A: angles between all edges. ϵ : desired edge length, set to 0.1 for our task. α : desired edge angle, set to 60 degrees. λ_i : hyperparameters, each set to [10, 1, 1, 0.5]. $\sigma(a) = (1 - \text{sigmoid}(a - 10)) + \text{sigmoid}(150)$: filter to minimize the penalty for “good enough” angles, i.e. between 10 and 150 degrees.

We used the Adam optimizer with a learning rate of $1e-2$ and achieved convergence at around $n_{opt} = 100$.

(CleanMesh) For the final DMTet output, we cleaned the mesh by removing overlapping nodes (distance $< 1e-2$) and elements with repeated nodes (i.e. points and edges).

Constrained surface remeshing details

Although the DMTet output is already a suitable manifold surface for the final tetrahedralization, it can be further refined via surface remeshing (Algorithm 5). The main remeshing algorithm we used is ACVD, a Voronoi diagram-based vertex clustering method⁴⁶.

Algorithm 5. Constrained remeshing by vertex-clustering

```
function ConstrainedRemesh( $S_{ca2}$ )
  for  $n_{remesh}$  do
     $S_{contact}$   $S_{free} \leftarrow$  SeparateContactSurf( $S_{ca2}$ )
     $S_{free2} \leftarrow$  ConstrainedClustering( $S_{free}$ )
     $S_{ca2} \leftarrow$  Merge( $S_{contact}$   $S_{free2}$ )
  return  $S_{ca2}$ 
```

To preserve the contact surface elements, we made two modifications to the remeshing steps.

(SeparateContactSurf) First, we split the initial surface into contact and non-contact elements by checking the number of nodes merged with the original tissue mesh’s nodes. If all three nodes of a triangular element are coincident with tissue mesh nodes, then that element is a contact element.

(ConstrainedClustering) We performed the initial ACVD clustering step on the non-contact surface, which assigns a cluster index to each node. The number of initial clusters was set to 80% of the number of nodes. The nodes belonging to both the contact and non-contact elements were preserved by assigning new unique cluster indices. Note that the percentage of vertices actually being clustered progressively decreases with each remeshing step because of our cluster re-assigning scheme. Finally, we performed the standard ACVD triangulation to obtain the remeshed non-contact surface.

(Merge) We merged the contact surfaces and the remeshed non-contact surfaces to obtain the final output surface mesh. This step was easy to perform because we preserved the coincident nodes between the contact and non-contact surfaces using constrained clustering.

The constrained remeshing was repeated as many times as necessary to achieve the desired element size and density. We often stopped at $n_{remesh} = 15$ to prevent oversimplification. The remeshing step is optional, as we can obtain good quality tetrahedral calcification meshes directly from the output of DMTetOpt. However, we found that the remeshing step generally further improves the element quality for downstream applications.

Final tetrahedralization

We obtained the final tetrahedral calcification mesh by applying TetGen on the surface mesh from the previous subsections. This step was straightforward because we designed the surface meshes to be clean and manifold. For node-stitching, we first assigned each connected component of calcification to a leaflet based on the lowest mean symmetric chamfer distance, and nodes from the assigned leaflet or the aortic wall were merged with calcification nodes if the distances were less than $1e-3$. Connected components of calcification that had no merged nodes with tissue surfaces were removed.

Statistical analyses

The scipy package⁵² was used to perform related-sample t-tests with two-sided alternative hypothesis. Metrics from repeated meshing trials were first averaged across each patient.

Data availability

The public MMWHS dataset used in this study can be found at⁵⁴. The de-identified patient data from Hartford Hospital may be made available upon reasonable request to the corresponding author, subject to approval from the Institutional Review Board.

Code availability

Our implementation is available at <https://github.com/danpak94/Deep-Cardiac-Volumetric-Mesh>.

Received: 23 February 2024; Accepted: 26 July 2024;

Published online: 15 August 2024

References

- Greenland, P., LaBree, L., Azen, S. P., Doherty, T. M. & Detrano, R. C. Coronary artery calcium score combined with framingham score for risk prediction in asymptomatic individuals. *JAMA* **291**, 210–215 (2004).
- Chen, J. et al. Coronary artery calcification and risk of cardiovascular disease and death among patients with chronic kidney disease. *JAMA Cardiol.* **2**, 635–643 (2017).
- Witteman, J. M., Kok, F., Van Saase, J. C. & Valkenburg, H. Aortic calcification as a predictor of cardiovascular mortality. *Lancet* **328**, 1120–1122 (1986).
- Nicoll, R. & Henein, M. Y. The predictive value of arterial and valvular calcification for mortality and cardiovascular events. *IJC Heart Vessels* **3**, 1–5 (2014).
- Sangiorgi, G. et al. Arterial calcification and not lumen stenosis is highly correlated with atherosclerotic plaque burden in humans: a histologic study of 723 coronary artery segments using

- noncalcifying methodology. *J. Am. Coll. Cardiol.* **31**, 126–133 (1998).
6. Durham, A. L., Speer, M. Y., Scatena, M., Giachelli, C. M. & Shanahan, C. M. Role of smooth muscle cells in vascular calcification: implications in atherosclerosis and arterial stiffness. *Cardiovasc. Res.* **114**, 590–600 (2018).
 7. Mohler, E. R. Mechanisms of aortic valve calcification. *Am. J. Cardiol.* **94**, 1396–1402 (2004).
 8. Pawade, T., Sheth, T., Guzzetti, E., Dweck, M. R. & Clavel, M.-A. Why and how to measure aortic valve calcification in patients with aortic stenosis. *JACC Cardiovasc. Imaging* **12**, 1835–1848 (2019).
 9. Marquis-Gravel, G., Redfors, B., Leon, M. B. & Genereux, P. Medical treatment of aortic stenosis. *Circulation* **134**, 1766–1784 (2016).
 10. Benfari, G. et al. Concomitant mitral regurgitation and aortic stenosis: one step further to low-flow preserved ejection fraction aortic stenosis. *Eur. Heart J. Cardiovasc. Imaging* **19**, 569–573 (2018).
 11. Slimani, A. et al. Relative contribution of afterload and interstitial fibrosis to myocardial function in severe aortic stenosis. *Cardiovasc. Imaging* **13**, 589–600 (2020).
 12. Ge, L. & Sotiropoulos, F. Direction and magnitude of blood flow shear stresses on the leaflets of aortic valves: is there a link with valve calcification? *J. Biomech. Eng.* **132**, 014505 (2010).
 13. Halevi, R. et al. Fluid–structure interaction modeling of calcific aortic valve disease using patient-specific three-dimensional calcification scans. *Med. Biol. Eng. Comput.* **54**, 1683–1694 (2016).
 14. Berdajs, D., Mosbahi, S., Ferrari, E., Charbonnier, D. & von Segesser, L. K. Aortic valve pathology as a predictive factor for acute aortic dissection. *Ann. Thorac. Surg.* **104**, 1340–1348 (2017).
 15. Weinberg, E. J., Mack, P. J., Schoen, F. J., Garcia-Cardeña, G. & Kaazempur Mofrad, M. R. Hemodynamic environments from opposing sides of human aortic valve leaflets evoke distinct endothelial phenotypes in vitro. *Cardiovasc. Eng.* **10**, 5–11 (2010).
 16. Arzani, A. & Mofrad, M. R. A strain-based finite element model for calcification progression in aortic valves. *J. Biomech.* **65**, 216–220 (2017).
 17. Qin, T. et al. The role of stress concentration in calcified bicuspid aortic valve. *J. R. Soc. Interface* **17**, 20190893 (2020).
 18. Milhorini Pio, S., Bax, J. & Delgado, V. How valvular calcification can affect the outcomes of transcatheter aortic valve implantation. *Expert Rev. Med. Devices* **17**, 773–784 (2020).
 19. Pollari, F. et al. Aortic valve calcification as a risk factor for major complications and reduced survival after transcatheter replacement. *J. Cardiovasc. Computed Tomogr.* **14**, 307–313 (2020).
 20. Wang, Q., Kodali, S., Primiano, C. & Sun, W. Simulations of transcatheter aortic valve implantation: implications for aortic root rupture. *Biomech. Model. Mechanobiol.* **14**, 29–38 (2015).
 21. Sturla, F. et al. Impact of different aortic valve calcification patterns on the outcome of transcatheter aortic valve implantation: a finite element study. *J. Biomech.* **49**, 2520–2530 (2016).
 22. Kurugol, S. et al. Automated quantitative 3d analysis of aorta size, morphology, and mural calcification distributions. *Med. Phys.* **42**, 5467–5478 (2015).
 23. Mahabadi, A. A. et al. Association of aortic valve calcification to the presence, extent, and composition of coronary artery plaque burden: from the rule out myocardial infarction using computer assisted tomography (romicat) trial. *Am. Heart J.* **158**, 562–568 (2009).
 24. Alqhtani, A. M. et al. Quantifying aortic valve calcification using coronary computed tomography angiography. *J. Cardiovasc. Computed Tomogr.* **11**, 99–104 (2017).
 25. Bettinger, N. et al. Practical determination of aortic valve calcium volume score on contrast-enhanced computed tomography prior to transcatheter aortic valve replacement and impact on paravalvular regurgitation: elucidating optimal threshold cutoffs. *J. Cardiovasc. Computed Tomogr.* **11**, 302–308 (2017).
 26. Vlastra, W. et al. Aortic valve calcification volumes and chronic brain infarctions in patients undergoing transcatheter aortic valve implantation. *Int. J. Cardiovasc. Imaging* **35**, 2123–2133 (2019).
 27. Grbic, S. et al. Image-based computational models for tavi planning: from ct images to implant deployment. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013: 16th International Conference, Nagoya, Japan, September 22–26, 2013, Proceedings, Part II 16*, 395–402 (Springer, 2013).
 28. Harbaoui, B. et al. Aorta calcification burden: towards an integrative predictor of cardiac outcome after transcatheter aortic valve implantation. *Atherosclerosis* **246**, 161–168 (2016).
 29. Graffy, P. M., Liu, J., O'Connor, S., Summers, R. M. & Pickhardt, P. J. Automated segmentation and quantification of aortic calcification at abdominal ct: application of a deep learning-based algorithm to a longitudinal screening cohort. *Abdom. Radiol.* **44**, 2921–2928 (2019).
 30. Morganti, S. et al. Simulation of transcatheter aortic valve implantation through patient-specific finite element analysis: two clinical cases. *J. Biomech.* **47**, 2547–2555 (2014).
 31. Loureiro-Ga, M. et al. A biomechanical model of the pathological aortic valve: simulation of aortic stenosis. *Comput. Methods Biomech. Biomed. Eng.* **23**, 303–311 (2020).
 32. Russ, C. et al. Simulation of transcatheter aortic valve implantation under consideration of leaflet calcification. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 711–714 (IEEE, 2013).
 33. Bianchi, M. et al. Patient-specific simulation of transcatheter aortic valve replacement: impact of deployment options on paravalvular leakage. *Biomech. Model. Mechanobiol.* **18**, 435–451 (2019).
 34. Kong, F., Wilson, N. & Shadden, S. A deep-learning approach for direct whole-heart mesh reconstruction. *Med. Image Anal.* **74**, 102222 (2021).
 35. Kong, F. & Shadden, S. C. Learning whole heart mesh generation from patient images for computational simulations. In *IEEE Transactions on Medical Imaging* (IEEE, 2022).
 36. Pak, D. H. et al. Distortion energy for deep learning-based volumetric finite element mesh generation for aortic valves. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24*, 485–494 (Springer, 2021).
 37. Pak, D. H. et al. Weakly supervised deep learning for aortic valve finite element mesh generation from 3d ct images. In *Information Processing in Medical Imaging: 27th International Conference, IPMI 2021, Virtual Event, June 28–June 30, 2021, Proceedings 27*, 637–648 (Springer, 2021).
 38. Pak, D. H. et al. Patient-specific heart geometry modeling for solid biomechanics using deep learning. In *IEEE Transactions on Medical Imaging* (IEEE, 2023).
 39. Fu, X.-M., Liu, Y. & Guo, B. Computing locally injective mappings by advanced mips. *ACM Trans. Graph. (TOG)* **34**, 1–12 (2015).
 40. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, 234–241 (Springer, 2015).
 41. Isensee, F., Jaeger, P. F., Kohl, S. A., Petersen, J. & Maier-Hein, K. H. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat. methods* **18**, 203–211 (2021).
 42. Shen, T., Gao, J., Yin, K., Liu, M.-Y. & Fidler, S. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Adv. Neural Inf. Process. Syst.* **34**, 6087–6101 (2021).
 43. Hang, S. Tetgen, a delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw.* **41**, 11 (2015).
 44. Payne, B. A. & Toga, A. W. Surface mapping brain function on 3d models. *IEEE Comput. Graph. Appl.* **10**, 33–41 (1990).

45. Chan, S. L. & Purisima, E. O. A new tetrahedral tessellation scheme for isosurface generation. *Comput. Graph.* **22**, 83–90 (1998).
46. Valette, S., Chassery, J. M. & Prost, R. Generic remeshing of 3d triangular meshes with metric-dependent discrete voronoi diagrams. *IEEE Trans. Vis. Comput. Graph.* **14**, 369–381 (2008).
47. Martin, C. & Sun, W. Comparison of transcatheter aortic valve and surgical bioprosthetic valve durability: a fatigue simulation study. *J. Biomech.* **48**, 3026–3034 (2015).
48. Mao, W., Wang, Q., Kodali, S. & Sun, W. Numerical parametric study of paravalvular leak following a transcatheter aortic valve deployment into a patient-specific aortic root. *J. Biomech. Eng.* **140**, 101007 (2018).
49. Caballero, A., Mao, W., McKay, R. & Sun, W. The impact of self-expandable transcatheter aortic valve replacement on concomitant functional mitral regurgitation: a comprehensive engineering analysis. *Struct. Heart* **4**, 179–191 (2020).
50. Medtronic LLC. Medtronic corevalve system instructions for use. https://www.accessdata.fda.gov/cdrh_docs/pdf13/P130021S033C.pdf (2017). Accessed: 2023-09-04.
51. Scott, D. W. *Multivariate density estimation: theory, practice, and visualization* (John Wiley & Sons, 2015).
52. Virtanen, P. et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **17**, 261–272 (2020).
53. Ozturk, C. et al. AI-powered multimodal modeling of personalized hemodynamics in aortic stenosis. arXiv preprint arXiv:2407.00535 (2024).
54. Zhuang, X. & Shen, J. Multi-scale patch and multi-modality atlases for whole heart segmentation of mri. *Med. image Anal.* **31**, 77–87 (2016).
55. Drozdal, M., Vorontsov, E., Chartrand, G., Kadoury, S. & Pal, C. The importance of skip connections in biomedical image segmentation. In *International Workshop on Deep Learning in Medical Image Analysis, International Workshop on Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*, 179–187 (Springer, 2016).
56. Sudre, C. H., Li, W., Vercauteren, T., Ourselin, S. & Jorge Cardoso, M. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3*, 240–248 (Springer, 2017).
57. Pak, D. H., Caballero, A., Sun, W. & Duncan, J. S. Efficient aortic valve multilabel segmentation using a spatial transformer network. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, 1738–1742 (IEEE, 2020).
58. Ma, J. et al. Loss odyssey in medical image segmentation. *Med. Image Anal.* **71**, 102035 (2021).
59. Paszke, A. et al. Pytorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **32**, (2019).
60. Kervadec, H. et al. Boundary loss for highly unbalanced segmentation. In *International conference on medical imaging with deep learning*, 285–296 (PMLR, 2019).
61. Wang, N. et al. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European conference on computer vision (ECCV)*, 52–67 (ECCV, 2018).
62. Schroeder, W., Martin, K. M. & Lorensen, W. E. *The visualization toolkit an object-oriented approach to 3D graphics* (Prentice-Hall, Inc., 1998).
63. Treece, G. M., Prager, R. W. & Gee, A. H. Regularised marching tetrahedra: improved iso-surface extraction. *Comput. Graph.* **23**, 583–598 (1999).

Acknowledgements

This work was supported by the National Heart, Lung, and Blood Institute (NHLBI) of the National Institute of Health (NIH), grants F31HL162505, T32HL098069, and R01HL121226. We thank Yuhang Du from Texas Tech University for assisting with running some of the simulations.

Author contributions

D.H.P. designed the algorithm and performed all experiments other than the finite element simulations. M.L. performed the finite element simulations. D.H.P. and M.L. wrote the manuscript. T.K. provided guidance for algorithm assessment. D.H.P., M.L., R.G., and J.S.D. reviewed the results. C.O. and E.R. evaluated the outputs. D.H.P. and R.M. curated the dataset.

Competing interests

D.H.P. is the inventor of a related provisional patent application No. 63/611,903. E.T.R. is a member of the board of directors at Affluent Medical and also serves on the board of advisors for Pumpinheart and Helios Cardio. E.T.R. offers consulting services for Holistick Medical and is a co-founder of Spheric Bio and Fada Medical. The remaining authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41746-024-01202-9>.

Correspondence and requests for materials should be addressed to Daniel H. Pak.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024