The Journal of Open Source Software

multimodars: A Rust-powered toolkit for multi-modality cardiac image fusion and registration

Anselm W. Stark [1,2¶], **Marc Ilic** [1,2], **Ali Mokhtari** [1,2], **Pooya Mohammadi Kazaj** [1,2], **Christoph Gräni** [1], **and Isaac Shiri** [1]

1 Department of Cardiology, Inselspital, Bern University Hospital, University of Bern, Switzerland **2** Graduate School for Cellular and Biomedical Sciences, University of Bern, Bern, Switzerland ¶ Corresponding author

DOI: 10.21105/joss.10131

Software

- Review
- Repository
- Archive

Editor: Erick Martins Ratamero

Reviewers:

- @jonaspleyer
- @richardkoehler
- @crnh

Submitted: 28 August 2025
Published: 07 May 2026

Summary

Coronary artery anomalies (CAAs) and coronary artery disease (CAD) require precise morphological and functional assessment for diagnosis and treatment planning. Cardiac computed tomography angiography (CCTA) provides a comprehensive 3D coronary anatomy but lacks the sub-millimeter resolution and dynamic tissue detail available from intravascular imaging, such as intravascular ultrasound (IVUS) and optical coherence tomography (OCT).

The `multimodars` package is a general-purpose toolkit that registers high-resolution intravascular pullbacks to CCTA-derived centerlines, producing locally enhanced fusion 3D vessel representations. Developed initially to quantify dynamic lumen changes in CAAs, the toolkit produces high-fidelity models suitable for visualization, geometric analysis, and patient-specific modeling. It implements four alignment paradigms (full, double-pair, single-pair, single) to compare pullbacks acquired under different hemodynamic states (e.g., rest vs. pharmacologic stress) or at different clinical timepoints (e.g., pre- vs. post-stenting). `Multimodars` targets reproducible multimodal fusion for both specialized CAA research and general CAD applications, producing quantitative lumen metrics (minimum lumen area, stenosis fraction, elliptic ratio, per-frame deformation) and optional interpolated deformation animations alongside the 3D models.

Statement of Need

Building reliable 3D coronary models requires combining complementary imaging modalities. Intravascular imaging offers exceptional local resolution but lacks whole-vessel context and 3D orientation. CCTA provides the global 3D geometry but suffers from limited spatial resolution and artifacts like blooming. `Multimodars` fills a critical gap for researchers in **cardiac imaging**, **interventional cardiology**, and **biomedical engineering** who require high-fidelity lumen models for:

- Automated quantification of vessel deformation under stress.
- Patient-specific Computational Fluid Dynamics (CFD) and Fluid-Structure Interaction (FSI) simulations, which require high-fidelity lumen geometry unavailable from either modality alone (Fan et al., 2014; Kilic et al., 2020).
- Longitudinal and interventional comparisons (e.g., pre- vs. post-stenting) within a consistent registration framework.
- Digital twin and 3D-printing workflows for biomechanical device testing.

The package accepts CSV and NumPy inputs, including data formats produced by the AIVUS-CAA software (Stark et al., 2025), providing a standardized pipeline from raw image

segmentation to final 3D fusion. The bidirectional NumPy API additionally allows integration with any external segmentation tool.

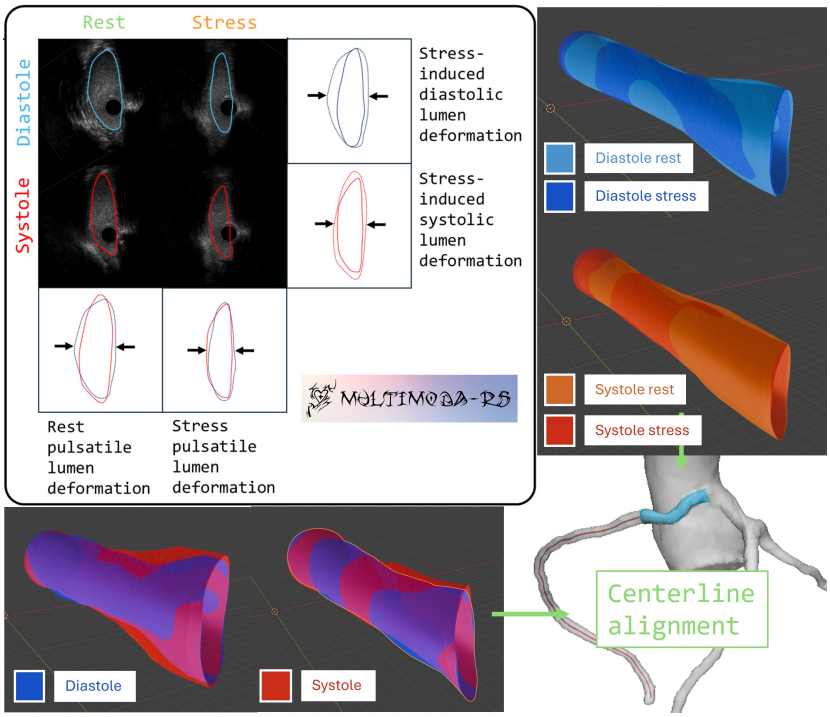

Figure 1: Figure 1: Illustration of `multimodars` processing modes and their clinical use. The 'full' mode returns four geometry pairs to analyze rest and stress hemodynamics (rest pulsatile deformation; stress pulsatile deformation; stress-induced diastolic deformation; stress-induced systolic deformation). The 'double-pair' mode returns pulsatile deformation in rest and stress. 'Single-pair' compares any two states (e.g., pre-/post-stent). 'Single' aligns frames within one pullback. A worked example of all four modes is available in the intravascular notebook included in the repository.

State of the Field

Prior research has established the clinical value of CCTA/intravascular fusion (Bourantas et al., 2013; Giessen et al., 2010; Ilic et al., 2026; Kilic et al., 2020; Wu et al., 2020), but two barriers persist. First, most existing fusion solutions are tied to proprietary vendor hardware or closed-source commercial workstations, limiting academic transparency; the closest open-source precedent, IVUSAngio (Doulaverakis et al., 2013), relies on biplane X-ray angiography rather than CCTA and provides no multi-state support. Second, no existing toolkit is specifically tailored for **multi-state** analysis—comparing rest versus pharmacologic stress or pre- versus post-intervention—while maintaining a consistent alignment across pullbacks. Published IVUS analyses of anomalous aortic origin of coronary arteries (AAOCA) quantify pulsatile deformation only under resting conditions (Formato et al., 2023); extending such analyses to stress conditions requires the cross-state registration logic implemented in `multimodars`. While packages like `trimesh` or `SimpleITK` provide general mesh and registration utilities, they do not offer the domain-specific coronary alignment logic (e.g., cumulative rotation propagation to preserve vessel torsion, ellipticity-weighted inter-pullback harmonization, CCTA centerline morphing) required here.

Software Design

Architectural choices and trade-offs

We chose **Rust** for the core backend to leverage its memory safety and hierarchical data parallelism (via the Rayon crate). This allows the toolkit to handle the significant computational load of multiscale angular searches across hundreds of image frames without the performance limitations found in pure Python implementations.

The data model uses a compact typed structure (PyContourPoint, PyContour, PyFrame, PyGeometry) that maps losslessly to (N,4) NumPy arrays (e.g., frame_id, x, y, z). This choice balances the performance of low-level data structures with the usability of the Python data science ecosystem.

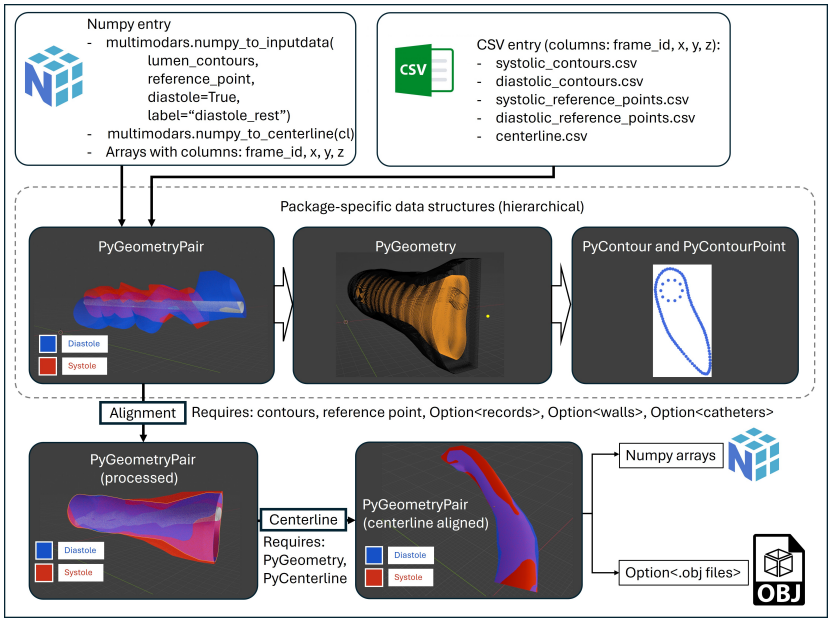

Figure 2: Figure 2: Overview of the intravascular module. Top: Input data in (N,4) format (frame_id, x, y, z) is accepted from CSV files or arbitrary NumPy arrays. Middle: Data are organized into a four-level hierarchy — PyGeometryPair → PyGeometry → PyFrame → PyContour → PyContourPoint. Bottom: Frames are aligned intra-pullback within each PyGeometry and inter-pullback within a PyGeometryPair; the resulting geometry can be registered to a CCTA centerline and exported as Wavefront OBJ files or converted back to NumPy arrays.

Alignment and fusion algorithms

Alignment is a two-stage pipeline producing spatially and rotationally consistent mappings both within pullbacks (intra-pullback) and between pullbacks (inter-pullback). It further registers 3D models with a CCTA-derived centerline, morphs the CCTA mesh to match intravascular dimensions, and stitches the two surfaces together.

- **Intra-pullback:** The proximal frame is the rotational reference. Sequentially, each proximal→distal neighbor is aligned by centroid translation and a rotation search minimizing a point-set distance derived from directed Hausdorff distances. Rotation employs a multiscale angular search (coarse → fine, e.g., 1° → 0.1° → 0.01°) with cumulative rotation propagation to preserve vessel torsion (See Figure 4). Naive brute-force complexity scales as $O(n \times \frac{R}{S} \times m^2)$ (n = frames, m = points per contour, R = angular range, S = step size); the multiscale refinement reduces this to an effective $O(n \times (R + c) \times m^2)$ for small step sizes, making runtime far less sensitive to angular

resolution while preserving alignment accuracy.

- **Inter-pullback:** Inter-pullback alignment harmonizes distal centroids, averages slice spacing to align z-coordinates, and applies a rigid rotation to minimize mean directed distances across corresponding frames; ellipticity-weighted similarity prioritizes non-round stenotic slices.
- **CCTA-Centerline alignment:** The `multimodars` package implements a three-point (aortic-, cranial- and caudal direction) anatomical registration and a manual alignment mode. It additionally utilizes Hausdorff distances to CCTA mesh points for ambiguous anatomies: centerlines are resampled to contour spacing, centroids are translated to matched points, normals are aligned by cross-product computations, and an optional interpolated UV-mapped mesh is produced for visualization and downstream modeling.
- **CCTA mesh fusion:** A CCTA STL surface mesh is labeled by propagating a rolling sphere along coronary centerlines to assign vertices to aorta, RCA, or LCA regions. For CAAs, where the intramural vessel runs within the aortic wall, a ray-casting occlusion step corrects systematic mislabeling: rays cast from the aortic centerline detect intersecting wall faces, and topologically connected coronary vertices are reclassified as aortic. The labeled mesh is then subdivided into proximal, anomalous, and distal sub-regions; each region is morphed radially along the local centerline normal to match the aligned intravascular lumen dimensions; and the replaced CCTA segment is excised to expose a boundary ring. The intravascular mesh is stitched to this ring with a triangulated patch, and the resulting surface is repaired and isotropically remeshed (via optional PyMeshLab).

Aligned pairs expose a `get_summary()` method returning minimum lumen area, maximum stenosis percent, stenosis length, and a per-frame table of lumen area and elliptic ratio for downstream analysis.

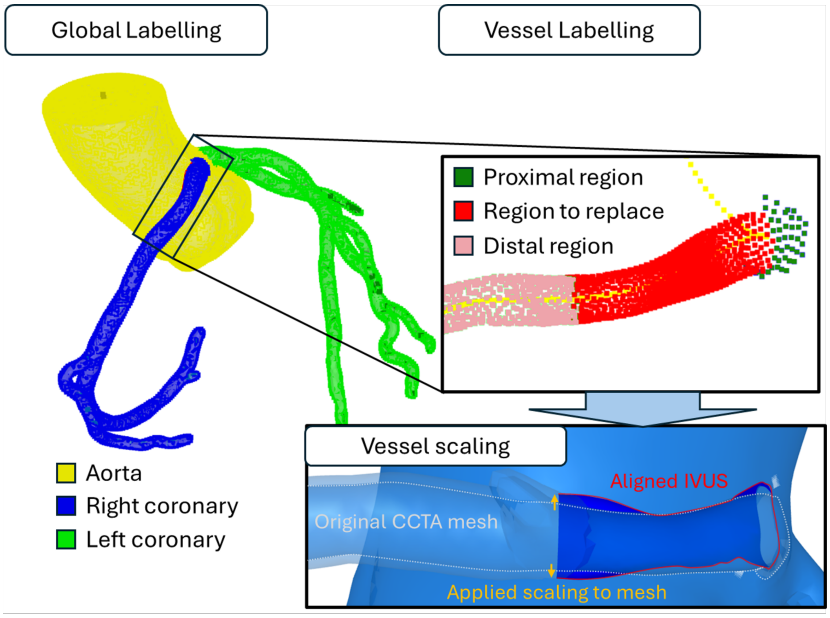

Figure 3: Figure 3: Top left shows the initial labeling of the coronary arteries and aorta based on centerlines. In a second step the regions to be replaced are identified based on the aligned 3D intravascular model. In a last step the CCT proximal and distal borders can be adjusted to best match intravascular borders.

Alignment algorithm performance and parallelization

Rust (Rayon) provides hierarchical data parallelism: candidate rotation angles within each frame comparison are evaluated concurrently across all available cores, with independent pullbacks additionally processed in parallel via scoped threads (Crossbeam). Typical production workflows downsample contours to 200–500 points/frame to balance sub-pixel accuracy and compute time.

A reproducible benchmark suite is included in the repository (`benchmarks/`) and documented in the package reference (ReadTheDocs). Key results on an Intel Xeon Gold 6234 (8 physical cores, 16 logical processors):

- **Algorithm vs. brute-force:** the multiscale search outperforms exhaustive brute-force evaluation by **5.5×** at a 0.1° step size and **10.3×** at 0.05°, with the gap widening rapidly at finer resolutions as brute-force must evaluate every candidate angle.
- **Parallelization:** parallelizing at the angle-evaluation level (rather than the point-rotation level) provides sufficient rayon tasks to utilize cores effectively. Brute-force scales **6.5×** and the optimized search **4.2×** from 2 to 16 cores.
- **Combined gain:** relative to brute-force at 2 cores, the optimized algorithm at 16 cores is **38.5×** faster, with algorithm choice and core count contributing multiplicatively.

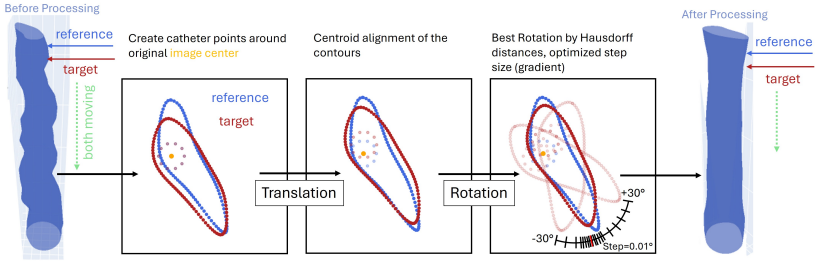

Figure 4: Figure 4: Multiscale intra-pullback alignment workflow (coarse-to-fine angular search and centroid propagation).

Implementation, reproducibility and usage

The package is built with `maturin` and available on PyPI. Output includes Wavefront OBJ files for intravascular meshes, STL section exports for downstream biomechanical simulation, and optional interpolated intermediate meshes for deformation animations. The project includes example notebooks, sample data, and CI tests; documentation and tutorials are hosted on ReadTheDocs (ReadTheDocs).

Research Impact Statement

`Multimodars` was motivated by the need to quantify dynamic lumen deformation in CAAs, where rest/stress and pulsatile comparisons are diagnostically critical. High-resolution fusion enables quantitative assessment of stress-induced deformation and supports patient-specific hemodynamic modeling. We successfully implemented this fusion approach to unveil a distinct compression pattern not visible in IVUS or CCTA alone in a case of an 11-year-old athlete with an anomalous aortic origin of a coronary artery (Stark et al., 2026). Beyond CAA, the `single-pair` mode supports general CAD workflows such as pre- vs. post-stenting lumen comparison, and the `single` mode enables standalone geometry reconstruction for any clinical pullback. We hope to foster a research community that leverages `multimodars` to standardize multimodal coronary fusion and accelerate the development of personalized interventional or computational strategies.

Human Participants Research Policy

The example data included in this repository originates from the NARCO study, a prospective clinical study investigating coronary artery anomalies. All data have been fully anonymized prior to inclusion. The study was conducted in accordance with the Declaration of Helsinki and received approval from the local ethics committee (Kantonale Ethikkommission Bern, KEK 2020-00841). All participants provided written informed consent. The study is registered with ClinicalTrials.gov (NCT04475289).

AI Usage Disclosure

No generative AI was used for architectural design or core algorithms. Claude (Anthropic) and ChatGPT (OpenAI) were used for creating documentation and docstrings, bug fixing, and minor inline code changes. For this manuscript these tools were only used for grammatical changes.

Acknowledgments

None

References

Bourantas, C. V., Papafaklis, M. I., Athanasiou, L., Kalatzis, F. G., Naka, K. K., Siogkas, P. K., Takahashi, S., Saito, S., Fotiadis, D. I., Feldman, C. L., & others. (2013). A new methodology for accurate 3-dimensional coronary artery reconstruction using routine intravascular ultrasound and angiographic data: Implications for widespread assessment of endothelial shear stress in humans. *EuroIntervention*, *9*(5), 582–593. https://doi.org/10.4244/EIJV9I5A94

Doulaverakis, C., Tsampoulatidis, I., Antoniadis, A. P., Chatzizisis, Y. S., Giannopoulos, A. A., Kompatsiaris, Y., & Giannoglou, G. D. (2013). IVUSAngio tool: A publicly available software for fast and accurate 3D reconstruction of coronary arteries. *Computers in Biology and Medicine*, *43*(11), 1793–1803. https://doi.org/10.1016/j.compbiomed.2013.08.013

Fan, R., Tang, D., Yang, C., Zheng, J., Bach, R., Wang, L., Muccigrosso, D., Billiar, K., Zhu, J., Ma, G., Maehara, A., & Mintz, G. S. (2014). Human coronary plaque wall thickness correlated positively with flow shear stress and negatively with plaque wall stress: An IVUS-based fluid-structure interaction multi-patient study. *BioMedical Engineering OnLine*, *13*(1), 32. https://doi.org/10.1186/1475-925X-13-32

Formato, G. M., Agnifili, M. L., Arzuffi, L., Rosato, A., Ceserani, V., Zuniga Olaya, K. G., Secchi, F., Deamici, M., Conti, M., Auricchio, F., Bedogni, F., Frigiola, A., & Lo Rito, M. (2023). Morphological changes of anomalous coronary arteries from the aorta during the cardiac cycle assessed by IVUS in resting conditions. *Circulation: Cardiovascular Interventions*, *16*(7), e012636. https://doi.org/10.1161/CIRCINTERVENTIONS.122.012636

Giessen, A. G. van der, Schaap, M., Gijsen, F. J., Groen, H. C., Walsum, T. van, Mollet, N. R., Dijkstra, J., Vosse, F. N. van de, Niessen, W. J., Feyter, P. J. de, & others. (2010). 3D fusion of intravascular ultrasound and coronary computed tomography for in-vivo wall shear stress analysis: A feasibility study. *The International Journal of Cardiovascular Imaging*, *26*(7), 781–796. https://doi.org/10.1007/s10554-009-9546-y

Ilic, M., Häner, J., Lehmann, J., Stark, A. W., Illi, J., Gräni, C., Bentele, M., Baumann-Zumstein, P., Busch, J., & Haeberlin, A. (2026). A comprehensive workflow for CCTA and OCT data fusion with 3D printing validation: Advancing patient-specific testing

environments for percutaneous coronary intervention devices. *BioMedical Engineering OnLine*, *25*(15). https://doi.org/10.1186/s12938-025-01501-6

Kilic, Y., Safi, H., Bajaj, R., Serruys, P. W., Kitslaar, P., Ramasamy, A., Tufaro, V., Onuma, Y., Mathur, A., Torii, R., & others. (2020). The evolution of data fusion methodologies developed to reconstruct coronary artery geometry from intravascular imaging and coronary angiography data: A comprehensive review. *Frontiers in Cardiovascular Medicine*, *7*, 33. https://doi.org/10.3389/fcvm.2020.00033

Stark, A. W., Kazaj, P. M., Balzer, S., Ilic, M., Bergamin, M., Kakizaki, R., Giannopoulos, A., Haeberlin, A., Räber, L., Shiri, I., & others. (2025). Automated intravascular ultrasound image processing and quantification of coronary artery anomalies: The AIVUS-CAA software. *Computer Methods and Programs in Biomedicine*, *272*, 109065. https://doi.org/10.1016/j.cmpb.2025.109065

Stark, A. W., Mokhtari, A., Kakizaki, R., Bigler, M. R., Siepe, M., Shiri, I., Räber, L., & Gräni, C. (2026). An 11-year-old athlete with dynamic stenosis of coronary artery anomaly with normal FFR-dobutamine but abnormal iFR-dobutamine. *Case Reports*, *31*(12), 107028. https://doi.org/10.1016/j.jaccas.2026.107028

Wu, W., Samant, S., De Zwart, G., Zhao, S., Khan, B., Ahmad, M., Bologna, M., Watanabe, Y., Murasato, Y., Burzotta, F., & others. (2020). 3D reconstruction of coronary artery bifurcations from coronary angiography and optical coherence tomography: Feasibility, validation, and reproducibility. *Scientific Reports*, *10*(1), 18049. https://doi.org/10.1038/s41598-020-74264-w