Eurographics Workshop on Visual Computing for Biology and Medicine (2026)
R. Krueger, E. Mörth and K. Furmanová (Editors)

Conference Paper

Vessel-straightening-based Registration for Aortic Flow Comparison

Aaron Schroeder[1], Anja Horvat[2], Kai Ostendorf[1], Kathrin Bäumler[3], Domenico Mastrodicasa[4], Nina Kerschbaumer[3], Dominik Fleischmann[3], Bernhard Preim[1], and Gabriel Mistelbauer[3]

[1]Faculty of Computer Science, Otto von Guericke University Magdeburg, Germany
[2]Faculty of Mechanical Engineering and Naval Architecture, University of Zagreb, Croatia
[3]Department of Radiology, Stanford University School of Medicine, USA
[4]Department of Radiology, University of Washington School of Medicine, USA

Abstract

Comparing vascular flow fields across patients, time points, or simulation and imaging modalities requires registration methods that preserve both anatomical correspondence and vector orientation. Generic image-registration methods can align segmented vessel masks, but their deformation fields do not directly provide a physically meaningful transformation for attached blood-flow vectors. We present a model-based registration approach for aortic flow data that uses segmentations, centerlines, local vessel frames, radius scaling, and branch-label-based twist alignment to establish correspondence between vascular domains. Velocity vectors are relocated with the registered anatomy and reoriented using only the rotational component of the local frame transformation, preserving vector magnitudes while adapting directions to the target vessel course. In a synthetic benchmark with 30 aortic fixed–moving pairs, the proposed model-based transformation reduced the mean normalized vector root mean square error (RMSE) from 0.58 for an elastix baseline to 0.27 and reduced the mean angular error from 26.02° to 0.55°. Because these generated geometries are strictly tubular, the model-based transform also achieved high mask overlap, with a mean Dice coefficient of 0.97 compared with 0.98 for elastix. An additional quantitative evaluation on real-patient anatomies with synthetic centerline-aligned flow showed that elastix refinement is important for anatomical overlap in more complex shapes, increasing mean Dice from 0.82 for the model alone to 0.95 while keeping vector errors below the elastix baseline. These results indicate that explicit vessel-model correspondence can improve flow-vector consistency while remaining compatible with conventional anatomical registration refinement.

CCS Concepts

• ***Human-centered computing*** → ***Scientific visualization;*** • ***Computing methodologies*** → *Shape analysis;* • ***Applied computing*** → *Imaging;*

1. Introduction

Blood flow in large vessels is shaped by patient-specific anatomy. In the aorta, curvature, branching, local diameter changes, and pathological remodeling interact with the flow field and influence clinically relevant quantities such as flow direction, pressure, and wall-related hemodynamics. Aortic dissection is a particularly challenging example: the intimal flap separates true and false lumina, entry and re-entry tears connect both channels, and the vessel can remodel substantially over time [SKW21, YMN22]. Both 4D flow magnetic resonance imaging (MRI) and computational fluid dynamics (CFD) have therefore been used to study complex aortic and dissection hemodynamics [CWG*12, TS10, BVS*20].

Many clinically and scientifically relevant questions are *comparative*. Examples include follow-up analysis, pre- and post-treatment comparison, cohort comparison, and comparison of measured 4D flow MRI with simulated CFD. Comparative visualization methods can help users inspect such differences [VP04, GAW*11, BEGP18], but the comparison depends on accurate spatial correspondence. Before local velocities, cross-sectional profiles, or streamlines can be compared, the vascular datasets must be registered.

For vascular flow data, registration is not only an anatomical image-alignment problem. The flow field is an attached vector field: each voxel carries a velocity vector whose direction and magnitude are part of the data. If the vessel is bent, straightened, stretched, or otherwise deformed during registration, simply moving the vector components as image channels can produce directions that are geometrically inconsistent with the registered vessel. This is especially visible in streamline visualizations, where small localorientation errors accumulate into implausible trajectories. Thus, vascular flow registration must align the vessel domain and transform velocity vectors so they remain meaningful in the registered anatomy.

Generic rigid, affine, and non-rigid registration methods are powerful tools for image alignment, but their objective functions are usually formulated around image similarity or local contour matching. Such methods are not necessarily organized around the vessel centerline.

We propose a *vessel-straightening-based* registration method for vascular flow fields. The method takes segmentations of a fixed and a moving vessel tree and derives centerline-based local frames. The moving flow field is transformed into the fixed geometry by mapping through a straightened vessel coordinate system. During this process, the spatial transformation relocates each vector in the vessel lumen, and the rotational component of the local frame transformation is applied to the vector itself. This rotation-only vector transformation intentionally preserves velocity magnitude while adapting vector direction to the registered vessel course.

We evaluate the approach on aortic data, including aortic dissection anatomies as a geometrically complex test case. Our evaluation uses synthetic centerline-aligned flow rather than the pathology-specific true/false-lumen hemodynamics that make dissection clinically challenging. The current implementation uses manually assigned aortic branch labels to align the twist of the local frames. We test whether a vessel-straightening model provides a useful deformation for aortic flow-field registration and improves vector transformation compared with a generic registration baseline.

Motivating Applications

The proposed registration framework is motivated by several application scenarios. In longitudinal follow-up, registration can support visual inspection of how the anatomy and hemodynamics of one patient's aorta change throughout disease progression. For treatment assessment, pre- and post-intervention datasets can be compared to study how the intervention changed vessel course and hemodynamics. At the cohort level, registration can support comparison of common flow features and patient-specific deviations. Finally, 4D flow MRI–CFD comparison requires measured and simulated vector fields to be aligned, using the same pipeline after resampling tetrahedral CFD data to a volume.

These scenarios motivate the quantitative evaluation, which uses synthetic aortic flow data with known structure in simplified synthetic aortas as well as real patient anatomy.

The contributions of our work are:

- a vessel-straightening-based registration method for aortic flow-field data;
- a centerline- and radius-guided deformation model that follows the tubular organization of the aorta rather than relying only on local image similarity; and
- a quantitative synthetic benchmark for controlled evaluation of flow-vector registration.

2. Related Work

2.1. Medical Image Registration

Medical image registration is the estimation of a spatial transformation that aligns the moving image (the one being transformed) with the fixed image (the target). Rigid and affine models handle global pose differences, while non-rigid models such as volumetric B-splines (three-dimensional tensor-product deformation fields, distinct from the 1D curve fits used elsewhere in this paper for centerlines) represent local anatomical variation. Toolkits such as elastix provide mature implementations of intensity-based registration, multi-resolution optimization, and configurable transformation models [KSM*10]. Sotiras et al. provide a broader survey of deformable medical image registration, situating B-spline and diffeomorphic approaches within the wider state of the art [SDP13].

However, generic registration methods are not formulated around the tubular structure of vessels. They optimize image similarity or contour alignment, but they do not necessarily preserve a meaningful relationship between the centerline, cross-sectional frames, and the attached blood-flow vectors. For scalar images, a deformation mainly defines where intensities are resampled. For flow data, the deformation also raises the question of how velocity vectors should be reoriented after their positions have changed. This distinction motivates vessel-aware registration models for flow-field comparison.

2.2. Vascular Models and Tubular Reformation

Vascular anatomy is characterized by an approximate centerline, local cross-sections, radius variation, and branch structure. These structures are central to how physicians and visualization systems inspect vessels. Constructing this representation from a segmentation is itself a distinct problem: Selle et al. formalize the extraction of a vessel-tree graph, including skeletonization, branch labeling, and cross-sectional radius profiles, for hepatic vasculature [SPSP02]. We rely on the same class of skeleton-to-graph construction (Sec. 3.2) to build the fixed and moving vessel-tree representations from which our registration model is derived. Curved planar reformation, for example, maps curved vessels into views that are easier to inspect along the vessel course [KFW*02]. Straightening methods extend this idea by transforming tubular structures into parameterized domains that support comparison and visualization. Angelelli and Hauser used straightening for side-by-side visualization of tubular flow data [AH11], and Zhou et al. proposed a general straightening approach for 3D tubular objects [ZZWY25]. Unlike these approaches, which use straightening to support direct visualization of a tubular object, we use the straightened domain as an intermediate correspondence model that reorients an attached vector field between two distinct vessel geometries.

The local coordinate system used during straightening is also important. Frames that twist unnecessarily along the vessel can introduce artificial rotation in cross-sectional views and in transformed vector fields. Rotation-minimizing or minimum-twist frames provide a way to define stable local orientations along space curves [WJZL08, FM18]. Thus, our method uses the aortic centerline and local frames to construct a deformation that follows the vessel course. Together with these frames, the straightened domain forms the registration model, described in Sec. 3, that maps a moving flow field into a fixed aortic geometry.

2.3. Comparative Visualization of Medical Flow

Comparative visualization provides strategies for analyzing differences between spatial or temporal datasets, including juxtaposition, superposition, explicit difference visualization, and linked interaction [VP04, GAW*11]. In medical flow visualization, these strategies are used to compare simulations, measurements, treatment outcomes, or patient-specific cases. Van Pelt et al. presented comparative blood-flow visualization for cerebral aneurysm treatment assessment [vPGL*14], Behrendt et al. proposed a framework for visual comparison of 4D PC-MRI aortic blood-flow data [BEGP18], and Meuschke et al. introduced COMFIS for comparative visualization of simulated medical flow data [MVE*22]. Interactive probing and illustrative visualization have also been used to support exploration of 4D MRI blood-flow data [vPOBB*11, BMGS13].

These works demonstrate the value of visual comparison, but they also underline the need for reliable spatial correspondence. If two flow fields remain in different anatomical coordinate systems, apparent differences may reflect misregistration rather than hemodynamic variation, and this ambiguity cannot be fully resolved after the fact. This is particularly relevant for aortic data, where curvature, branching, and pathology can produce large geometric differences. Our work addresses the registration step that precedes such comparisons, with a focus on transforming the moving vector field in a geometrically consistent way.

2.4. Vector-Field Transformation

Spatially transforming a scalar image and spatially transforming directional data are different operations. When a coordinate warp is applied to data with orientation information, the position of each sample and the orientation attached to that sample must both be considered. This issue is well known in diffusion tensor imaging, where tensor-valued images require explicit reorientation during spatial transformation [APBG01]. A related line of work formulates diffeomorphic registration directly for vector-valued image data: Cao et al. jointly optimize a dense diffeomorphic deformation and a corresponding vector-reorientation rule by minimizing a vector-field similarity functional over the full image domain, without assuming any particular anatomical structure [CMWY05]. Our setting differs in that the transformed vector field is attached to a tubular anatomical structure with a well-defined longitudinal and cross-sectional organization. Rather than solving a general-purpose diffeomorphic optimization problem, we derive both the spatial correspondence and the vector-reorientation rule directly, in closed form, from the vessel centerline, local frames, and radius model (Sec. 3). This is computationally cheaper and yields an explicit, interpretable transformation, at the cost of being restricted to anatomy that admits this tubular parameterization. Although blood-flow velocity vectors are not diffusion tensors, the underlying concern is similar: a deformation can change the local coordinate frame in which directional information should be interpreted.

For vascular flow-field registration, this means that treating the three velocity components as independent scalar channels is inappropriate. The vector must be relocated according to the spatial deformation and reoriented according to the local transformation. Our method applies the rotational component of the vessel-frame transformation to each 3D velocity vector. This preserves velocity magnitude while adapting direction to the registered vessel course. The resulting approach occupies the gap between generic image registration, which optimizes image similarity, and comparative flow visualization, which assumes that the compared fields already have meaningful spatial correspondence.

3. Vessel-straightening-based Flow-Field Registration

Our method registers a moving vascular flow field to a fixed vascular anatomy by using the tubular organization of the vessel as an explicit deformation model. Local coordinate frames along the centerline define a straightened coordinate system in which corresponding locations along the moving and fixed aortas can be related. The resulting spatial transformation is then applied to the moving velocity volume, and each velocity vector is reoriented with the rotational component of the corresponding local frame transformation.

3.1. Problem Definition

Let the fixed dataset be represented by a vascular domain $\Omega_f \subset \mathbb{R}^3$ and, where available, a fixed velocity field $\mathbf{v}_f : \Omega_f \to \mathbb{R}^3$. The moving dataset is represented by a vascular domain $\Omega_m \subset \mathbb{R}^3$ and a moving velocity field $\mathbf{v}_m : \Omega_m \to \mathbb{R}^3$. The domains are derived from binary segmentations of the aorta and its relevant branches. The velocity fields may originate from 4D flow MRI or from volumetrically resampled CFD data. In both cases, the velocity samples, segmentation, and vessel model of each dataset must share one physical coordinate system before inter-dataset registration. However, for our evaluation we only use synthetic flow data, which provides a known ground truth for quantitative assessment of the registration quality. For measured MRI data, the vascular segmentations are derived directly from the MRI volume. For CFD data, the binary lumen mask is obtained from the labeled volumetric mesh used for the simulation and is resampled to a voxel grid before the registration pipeline is applied.

The goal is to transform the moving velocity field into the fixed geometry. We denote the spatial transformation from moving physical coordinates to fixed physical coordinates by $T_{m\to f}$. For resampling, we use the transformation in pull form: each fixed-grid target position $\mathbf{x}_f$ is mapped to a moving position $\mathbf{x}_m$, where the moving field is interpolated. The output is a registered velocity field $\hat{\mathbf{v}}_m : \Omega_f \to \mathbb{R}^3$ that is defined in the fixed coordinate system and can therefore be compared with $\mathbf{v}_f$.

In longitudinal follow-up and treatment assessment, moving and fixed domains correspond to different examinations of the same patient. In cohort analysis, they correspond to different patient. In MRI–CFD comparison, one velocity field is measured and the other simulated. The registration model is independent of the velocity source because it is derived from the segmented vascular geometry.

3.2. Vascular Model

For each dataset, the vessel mask is skeletonized using a 3D thinning approach [LKC94], and the resulting skeleton is converted into a directed graph representing the vessel tree [Mis13]. The centerline and vessel tree are computed automatically from the segmentation mask by these algorithms, without additional regularization or manual correction. The vessel tree consists of the main aortic centerline and available branch centerlines. The main aortic centerline is represented as an ordered sequence of points $\mathbf{c}(s)$ parameterized by arc length or by a normalized longitudinal coordinate s. At each centerline sample, the tangent direction $\mathbf{t}(s)$ is estimated from the local centerline geometry.

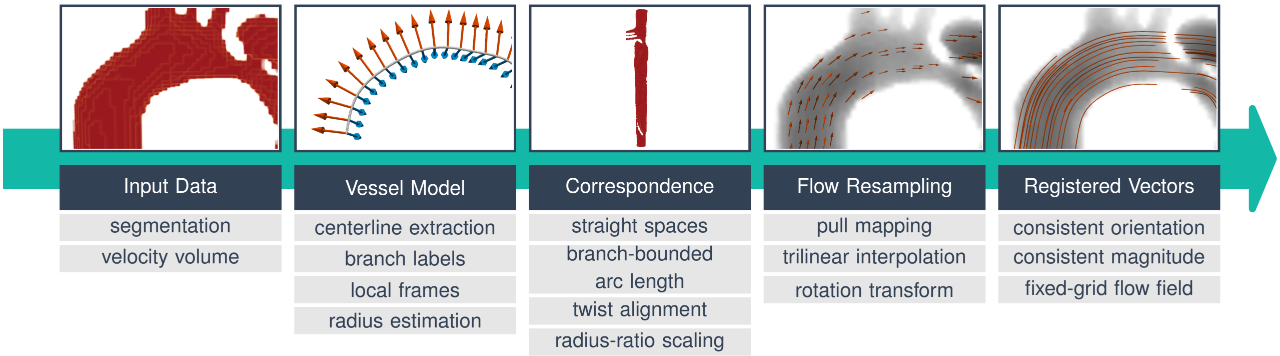

Figure 1: *Overview of the proposed aortic flow-field registration pipeline. Segmentations are used to extract centerlines, branch labels, local frames, and vessel radii. Fixed and moving anatomies are mapped into straightened vessel coordinates, where branch-bounded arc length and local radius scaling define the spatial mapping. The moving velocity field is sampled in pull form on the fixed grid, and each sampled vector is reoriented using the rotation between corresponding local vessel frames while preserving its magnitude.*

The local coordinate system of the aorta is described by a frame $(\mathbf{t}(s), \mathbf{n}(s), \mathbf{b}(s))$, where $\mathbf{t}(s)$ is the longitudinal direction and $\mathbf{n}(s)$ and $\mathbf{b}(s)$ span the cross-sectional plane. We use minimum-twist frames to avoid artificial rotation of the cross-sectional axes along curved vessel segments, ensuring anatomically consistent inter-dataset alignment.

The segmentation is also used to estimate a local vessel radius along the main aortic centerline. For each centerline sample, a cross-section orthogonal to the tangent is considered, and the local radius is estimated from the inscribed vessel region in that plane. Specifically, the radius is defined by the maximum inscribing circle that fits inside the vessel mask in the corresponding cross-sectional slice. The radius is used to account for diameter differences between the moving and fixed aortas during straightening. This produces a deformation that follows the longitudinal course of the vessel and its local cross-sectional scale, which can vary substantially between datasets, especially in aortic dissection.

Branches are used to align the twist of the local aortic frames. The current implementation uses manually assigned branch labels for the aortic branches. Because the evaluation in this work is limited to thoracic anatomy, the labeled branches are restricted to the brachiocephalic trunk, the left common carotid artery, and the left subclavian artery. Branches used for correspondence are determined by identifying labels that are present in both the fixed and moving vessel trees. The number of labels available for a given pair therefore depends on how many of these three arch branches fall within the imaged field of view for both datasets and can be identified as individual branches. These common labels provide anatomical orientation cues, by identifying corresponding branch directions in the fixed and moving datasets. At each centerline location with a common branch label, the cross-sectional frame is constrained so that its primary axis points toward that branch, which keeps the corresponding branch ostia aligned between the fixed and moving datasets after registration. Between such constrained locations, and outward from the constrained location nearest to each centerline endpoint, the frame orientation is propagated as a minimum-twist frame. No separate convention is used to fix the orientation at the two centerline endpoints themselves. The branches themselves are not explicitly registered as separate deformable structures.

For aortic dissection cases, the current model treats the segmented aorta as a single vascular domain and extracts one main aortic centerline: the true and false lumen segmentations are merged with a morphological closing operation before centerline extraction. True and false lumina are not modeled as separate centerline systems. This design keeps the deformation model robust, in the sense that it avoids the added topological complexity of tracking two lumina connected by entry and re-entry tears, and focused on the overall vessel course, but it also limits the ability to represent lumen-specific correspondences in dissected anatomy.

3.3. Straightened Registration Model

The registration is constructed by mapping both datasets through a straightened vessel coordinate system. The implementation iterates over cross-sectional slices of the curved vessel. For each slice, voxel positions inside the curved vessel mask are expressed relative to the local centerline frame and the corresponding transform entries are written to their positions in a straight-space volume. Thus, the straight-space representation stores a dense mapping between straightened coordinates and the original curved physical coordinates. Because neighboring cross-sectional slices may overlap inside the vessel mask, the same curved-space position can occur more than once in the straight-space volume. When several slices write to the same straight-space location, the entry produced by the most recently processed slice overwrites earlier entries. This last-write resolutions effect is largest where consecutive cross-sections overlap the most, i.e., in tightly curved vessel segments. Conversely, a curved-space location that no cross-sectional slice reaches during this traversal receives no straight-space entry. Such locations are not assigned an arbitrary correspondence but are excluded from resampling by the valid-vessel-region masking described below.

Conceptually, a physical point $\mathbf{x}$ near the aortic centerline is represented with respect to a local centerline sample $\mathbf{c}(s)$:

$$\mathbf{x} = \mathbf{c}(s) + u\mathbf{n}(s) + w\mathbf{b}(s), \tag{1}$$

where s is the longitudinal coordinate and (u,w) are cross-sectional coordinates. This maps the curved aorta into a straightened coordinate representation (s,u,w).

Fixed and moving straightened coordinates are related by corresponding longitudinal positions along the main aortic centerline. For a target point in the fixed vessel, the fixed straightening transform yields its coordinates (s_f,u_f,w_f). The corresponding moving longitudinal coordinate s_m is obtained from the centerline correspondence. The correspondence is based on normalized arc length between common branch locations, rather than on a single global normalized arc-length parameter over the entire aorta. Common branch labels therefore define corresponding vessel intervals, and positions within each interval are matched by their relative arc length. The same branch information is used to align the orientation of the cross-sectional frames at branch locations.

Cross-sectional coordinates are adjusted using the local radius ratio between the fixed and moving vessel models. Let $r_f(s_f)$ and $r_m(s_m)$ denote the local radius estimates. The normalized cross-sectional position of the fixed point is mapped to the moving vessel by scaling the radial coordinates according to the moving radius:

$$u_m = u_f \frac{r_m(s_m)}{r_f(s_f)}, \qquad w_m = w_f \frac{r_m(s_m)}{r_f(s_f)}. \tag{2}$$

The moving physical position is then reconstructed with the moving centerline frame:

$$\mathbf{x}_m = \mathbf{c}_m(s_m) + u_m\mathbf{n}_m(s_m) + w_m\mathbf{b}_m(s_m). \tag{3}$$

This provides the inverse mapping required for pull-based resampling from the moving flow volume to the fixed grid.

The dense transformation is evaluated only for positions in the fixed vascular region and in a narrow support region around it. For positions outside the valid vessel mask, the registered velocity is set to zero or marked invalid. This avoids interpreting interpolated vectors outside the vascular domain as meaningful flow data. After the vessel tree has been extracted, the mask is used to define valid voxels and to estimate local radii. It does not otherwise change the interpolation rule near the vessel boundary.

3.4. Transformation of the Blood-Flow Vector Field

As motivated in Sec. 2.4, the spatial transformation defines where the moving velocity field is sampled but not how the sampled vector should be oriented. We therefore transform each vector using the local relationship between the moving and fixed vessel frames. For a fixed target position $\mathbf{x}_f$, the registration determines the corresponding fixed frame $F_f(s_f) = [\mathbf{t}_f(s_f), \mathbf{n}_f(s_f), \mathbf{b}_f(s_f)]$ and moving frame $F_m(s_m) = [\mathbf{t}_m(s_m), \mathbf{n}_m(s_m), \mathbf{b}_m(s_m)]$. The rotation that maps vectors from moving to fixed local frame is

$$R_{m\to f} = F_f(s_f) F_m(s_m)^\top. \tag{4}$$

The registered velocity vector is then computed as

$$\hat{\mathbf{v}}_m(\mathbf{x}_f) = R_{m\to f}\, \mathbf{v}_m(\mathbf{x}_m), \tag{5}$$

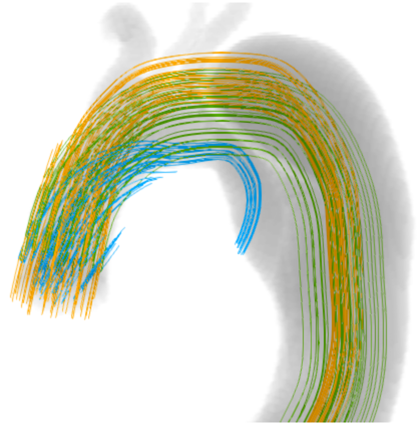

Figure 2: *Streamline demonstration for a real-anatomy case with synthetic flow. Streamlines generated from the registered flow fields of elastix (blue), model (green), and model+elastix (yellow) are shown, with the vascular mask in the background for reference. The elastix example deviates from the vessel course, while the model-based transformations produce streamlines that follow the vessel geometry.*

where $\mathbf{v}_m(\mathbf{x}_m)$ is obtained by trilinear interpolation of the moving velocity field.

This rotation-only transformation preserves the velocity magnitude of the sampled moving vector. It changes the vector direction according to the deformation of the vessel coordinate frame, but it does not scale the vector according to local stretching or compression of the spatial mapping, thus avoiding artificial changes in magnitude that could be introduced by unstable local scale factors.

A more general vector transformation can be constructed from the Jacobian of the dense deformation. The Jacobian can be decomposed into a rotation and a symmetric stretch component. Applying only the rotation gives the transformation used in this work, while applying the full Jacobian additionally scales and shears the vector according to local expansion or compression of the spatial mapping. However, in strongly curved or narrow vessel regions, Jacobian-based scaling can produce extreme magnitude changes that are difficult to interpret as blood-flow velocities. We informally observed this behavior when testing the Jacobian-based transformation on a subset of cases, which led us to exclude it from the current method. A systematic quantitative ablation of this design decision is left to future work.

3.5. Implementation Details

The method is implemented as a segmentation-driven preprocessing and resampling pipeline. Input segmentations are used to compute the vessel tree, the main aortic centerline, local minimum-twist frames, branch-label orientation constraints, and radius estimates. The fixed grid is traversed in physical coordinates, each valid target position is mapped into the moving physical coordinate system through the straightened representation, and the moving velocity field is sampled with trilinear interpolation. Vectors outside the valid vessel mask are set to zero.

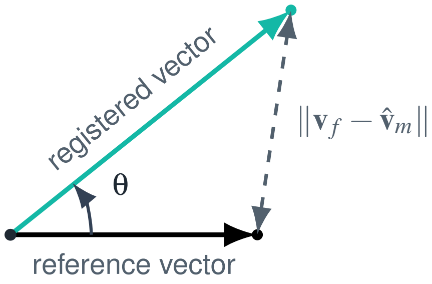

Figure 3: *Illustration of the flow-vector error metrics. Two vectors differ in length and direction. The arc indicates the angular error, the dashed connector between endpoints illustrates the vector difference whose squared magnitude contributes to the vector RMSE.*

4. Results

The evaluation is designed to test two linked questions. First, we evaluate whether the proposed vessel-straightening registration brings the moving vascular domain into accurate correspondence with the fixed domain using the Dice coefficient. Second, we evaluate whether the transformed moving velocity field remains geometrically consistent with the registered vessel course using flow-vector similarity and direction-error measures.

4.1. Data and Experimental Setup

We evaluate the method in two controlled quantitative settings. First, we use a synthetic benchmark with generated vessel geometries and centerline-aligned flow. This supports repeated testing across a number of generated cases because the expected flow direction after registration is known. The synthetic benchmark comprises 30 generated fixed–moving pairs, and the real-anatomy evaluation comprises four real aortic dissection fixed–moving pairs with synthetic centerline-aligned flow.

In the synthetic benchmark, paired fixed and moving vascular domains are generated by perturbing a base aortic centerline and vessel radius while maintaining tubular structure. The synthetic velocity field is generated such that flow is primarily aligned with the local centerline direction. Across each synthetic vessel cross-section, the radial velocity profile follows the modified power-law representation of fully developed pipe flow given by Salama [Sal21, Eq. 13]. This makes the benchmark suitable for testing whether registration preserves the expected relationship between the vessel course and the attached velocity vectors.

Synthetic centerlines are generated by perturbing the control points of a smoothed base spline with independent Gaussian offsets ($\sigma = 6\,\text{mm}$); no clipping is applied to these offsets. The local vessel radius is perturbed by a smoothed multiplicative Gaussian noise factor (pre-smoothing $\sigma = 0.4$, smoothed with a window of 15 samples), clamped to $[0.80, 1.20]$ of the base radius. Generated cases are not rejected or filtered by any additional plausibility criterion. Random number generators are seeded non-deterministically for each run, so exact perturbation instances are not reproducible across runs, although the generation procedure and its parameters are.

Second, we use real aortic dissection anatomies and generate synthetic flow in the same centerline-aligned manner. These cases are challenging because dissections have uniquely complex morphology and large patient-specific variation. This preserves the same quantitative tests while replacing the generated vessel geometry with realistic anatomy. The streamline rendering in Figure 2 demonstrates the resulting vector fields using streamlines.

The synthetic and real-anatomy test data, registration outputs, and evaluation results supporting this paper are publicly available [SHO*26]. The real aortic dissection anatomies used for evaluation are derived from the patient-specific fluid–structure interaction models of Baeumler et al. [BRPS*25] and the enlarged entry- and exit-tear models of Zimmermann et al. [ZBL*25].

4.2. Baseline Method

We compare the proposed vessel-straightening method to an elastix baseline applied to the vascular segmentations. The baseline uses affine registration followed by B-spline refinement, with 3D multi-resolution registration, recursive image pyramids, adaptive stochastic gradient descent, and advanced normalized correlation. Both stages use six resolution levels, 2000 optimizer iterations per level, random-coordinate sampling with 4000 spatial samples, and new samples at every iteration. The B-spline stage uses a final grid spacing of 8 voxels. B-spline interpolation order 1 is used during registration, and order 3 is used for final resampling. The resulting deformation resamples the moving flow field into the fixed domain. For vector transformation, we compute the deformation Jacobian and apply only its rotational component to the vectors, matching the rotation-only reorientation used by the proposed method.

For the model+elastix variant, elastix is used as a refinement of the spatial correspondence after the proposed vessel-model registration. The refinement improves anatomical matching by relocating already transformed vectors. It does not add vector rotation, shear, or scaling from elastix. Thus, the vector orientations in the model+elastix result remain those produced by the vessel-frame-based transformation. Concretely, this refinement adds a second resampling step on top of the vessel-model output: for every fixed-grid position, the already vessel-model-transformed vector is transformed by the displacement field, whose magnitude is bounded by the 8-voxel grid spacing used for the B-spline stage. This relocation is applied everywhere in the registered domain, but its magnitude is small relative to the vessel-model deformation. Its main visible effect is the Dice improvement reported in Sec. 4.4.

4.3. Evaluation Metrics

All quantitative metrics are computed in the valid vessel region where the fixed and registered moving fields are both defined. For anatomical alignment, we report overlap of the vessel masks using the Dice coefficient,

$$D = \frac{2|\Omega_f \cap \hat{\Omega}_m|}{|\Omega_f| + |\hat{\Omega}_m|},$$

where $\hat{\Omega}_m$ is the registered moving mask.

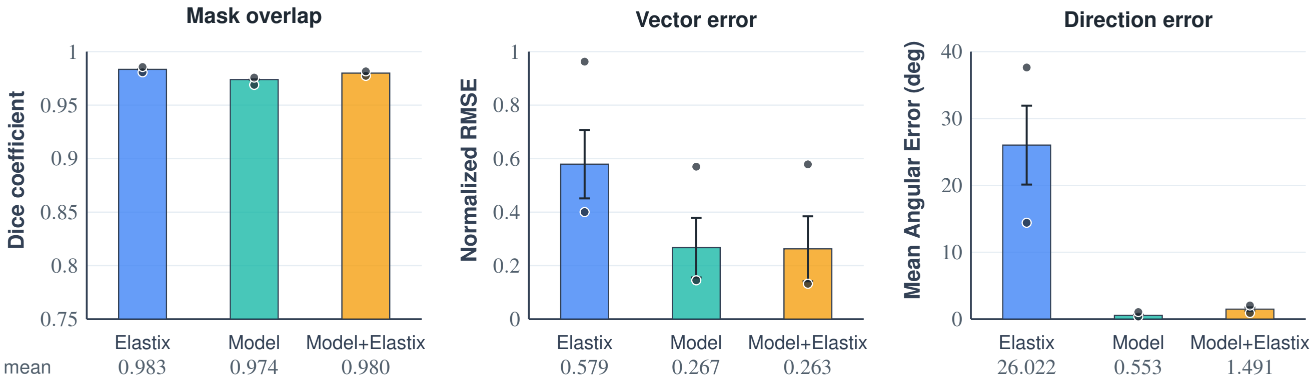

Figure 4: *Synthetic geometry results. Bars show mean values, whiskers show one standard deviation, and points show min/max cases. Model-based registration strongly reduces vector RMSE and angular error while preserving high mask overlap.*

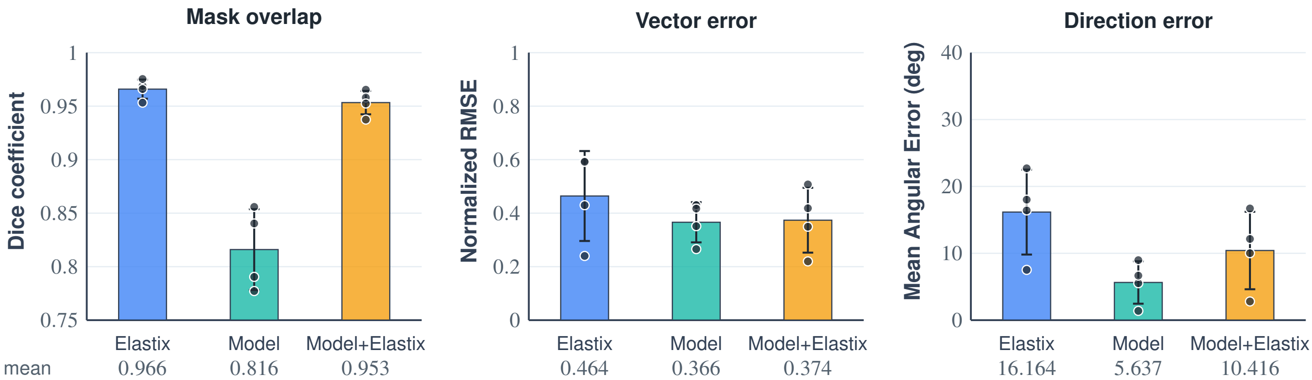

Figure 5: *Aortic dissection anatomy with synthetic flow. Bars show mean, points show individual cases, and whiskers show one standard deviation. Elastix refinement substantially improves overlap while retaining the model-based vector transformation.*

For flow-field comparison, we use metrics that evaluate both vector magnitude and vector direction, as illustrated in Figure 3. Evaluating the absolute vector error, we compute the normalized vector RMSE,

$$\mathrm{NRMSE}_{\mathrm{vec}} = \frac{\sqrt{\frac{1}{N}\sum_i \|\mathbf{v}_f(i) - \hat{\mathbf{v}}_m(i)\|^2}}{s_v},$$

where s_v is the mean velocity magnitude of the fixed reference field for the corresponding benchmark case. This metric is the root-mean-square vector difference after registration, normalized by the mean reference velocity; lower values indicate closer agreement. Isolating the directional error, the angular error is computed as

$$\theta_i = \arccos \frac{\mathbf{v}_f(i) \cdot \hat{\mathbf{v}}_m(i)}{\|\mathbf{v}_f(i)\| \, \|\hat{\mathbf{v}}_m(i)\|}.$$

For the synthetic benchmark, we additionally compute the centerline-tangency error, defined as the angle between the registered moving velocity vector and the local tangent direction of the fixed centerline.

4.4. Anatomical Registration Accuracy

Figure 4 summarizes the aggregate synthetic-benchmark results for 30 fixed–moving pairs with generated geometry and synthetic flow. The Dice scores show high mask overlap for both the elastix baseline and the proposed vessel-model registration, with mean Dice coefficients of 0.983 and 0.974, respectively. This small difference is expected because the generated geometries are strictly tubular with smoothly varying diameter, which matches the assumptions of the vessel-model transform. Adding the local elastix refinement increases the mean Dice coefficient only slightly to 0.980. Thus, for the synthetic geometries, the benefit of the refinement is negligible, and the main distinction between methods is their handling of the attached velocity vectors.

4.5. Flow-Vector Preservation

The flow-vector results show a clearer separation between the generic image-registration baseline and the vessel-model-based transformation. Although elastix produces high mask overlap, its transformed flow vectors remain substantially less consistent with the fixed synthetic flow field. The elastix baseline has a mean normalized vector RMSE of 0.579 and a mean angular error

Proceedings published by Eurographics - The European Association for Computer Graphics.

of 26.022°. In contrast, the model-based vector transformation reduces the mean normalized vector RMSE to 0.267 and the mean angular error to 0.553°. The model+elastix variant gives a similar mean normalized vector RMSE of 0.263, while its mean angular error increases to 1.491° because the elastix refinement relocates vectors without applying additional vector rotation.

Figure 5 repeats the same quantitative evaluation on real-patient aortic dissection anatomies with synthetic centerline-aligned flow. Here the unique morphology and variability of the real anatomy make the elastix refinement important for mask overlap: the model+elastix variant increases mean Dice from 0.816 for the model alone to 0.953, close to the elastix baseline at 0.966. The model-based transform reduces mean normalized vector RMSE from 0.464 to 0.366 and mean angular error from 16.164° to 5.637°. The model+elastix variant keeps the vector errors below the elastix baseline, with a mean normalized vector RMSE of 0.374 and mean angular error of 10.416°.

The angular-error result is particularly important because the synthetic flow field is constructed to follow the centerline direction. A low angular error therefore indicates that the transformed moving vectors are not only spatially relocated into the fixed vessel but also reoriented consistently with the fixed vessel course. Figure 6 illustrates the spatial distribution of this error for one comparison, showing where local vector directions deviate after registration. The elastix result contains larger high-error regions, whereas the model-based results keep most of the vessel at low angular error. A localized region of higher angular error is visible at the tightest bend of the aortic arch in all three results. At this location, the vessel frame orientation changes over a short distance, so small vector-placement errors are amplified by the angular-error metric. This local error is larger in the model+elastix result than in the model-only result because the elastix refinement relocates the already vessel-frame-transformed vectors (Sec. 4.2) without applying any additional rotation. At the tightest bend, this relocation moves vectors to sample positions where the correctly oriented flow field changes most rapidly, increasing local angular error even though the same refinement improves overall mask overlap.

Figure 6: *Demonstration of angular error in the vascular volume for elastix, model-based registration, and model+elastix from left to right. Red regions indicate larger local angular error.*

5. Discussion and Conclusion

The results show that vessel-aware registration can improve the transformation of vascular flow fields beyond what is achieved by optimizing mask overlap alone. This was observed both for generated synthetic geometries with synthetic flow and for real-patient anatomies with synthetic centerline-aligned flow. In both settings, the vessel-model transformation better preserved the expected flow direction after registration than the elastix baseline. This supports the central design choice of the method: spatial correspondence is derived from the vessel centerline, local frames, and radius estimates, while velocity vectors are reoriented with the rotational component of the local frame transformation. For the strictly tubular synthetic geometries, the model-based transform also achieved high mask overlap, so the added elastix refinement had little anatomical benefit. For real-patient anatomies, however, the refinement was important for recovering mask overlap while retaining the vector orientations produced by the vessel-frame-based transformation.

The main implication, discussed in Section 3.4, is that the deformation used for flow-field registration must define how vector directions are reinterpreted in the registered anatomy, not only relocate scalar-valued vector components. Our approach provides an explicit model for this step, which is particularly relevant for comparisons based on streamlines, local flow direction, or cross-sectional velocity profiles.

5.1. Limitations

Several limitations remain. The current evaluation uses synthetic flow fields for both the synthetic and real-patient anatomy cases, so the reported errors should be interpreted as controlled evidence of vector-transformation behavior rather than clinical validation of measured flow. Relatedly, the streamline renderings used to illustrate registered vector fields depict a single, centerline-aligned flow instant. Streamlines computed from measured pulsatile aortic flow would need to be interpreted as instantaneous snapshots rather than particle trajectories, since aortic flow is highly unsteady. The method also depends on the quality of the input segmentation, centerline, branch labels, and radius estimates. Errors in these intermediate structures can directly affect both the spatial mapping and the vector reorientation. The local radius estimate, defined by the maximum inscribed circle in each cross-section, can itself be sensitive to segmentation noise and is not uniquely defined at bifurcations. Per-case diameter and accumulated cross-sectional axis-angle profiles for the real-anatomy registration pairs are provided as supplementary material, illustrating this sensitivity beyond the aggregate metrics reported in Sec. 4. The straight-space domain mapping is also not strictly bijective: overlapping cross-sections in tightly curved regions are resolved by the last-write heuristic described in Sec. 3.3 rather than an explicit rule, and this ambiguity is not separately quantified in the reported metrics. Near the vessel wall, trilinear interpolation of the moving velocity field can blend valid interior samples with zero-valued background voxels outside the segmented lumen, which can locally bias the registered near-wall velocity magnitude. This boundary effect is not separately corrected or quantified.

The current model uses a single main centerline and does not separately model true and false lumina in aortic dissection. Branches are used to compute the vessel tree and to align frame twist, but they are not explicitly registered as deformable structures. Cross-sectional shape differences are represented only through local radius scaling, which cannot capture all wall-shape differences in complex anatomy. This shortcoming can be partly mitigated by the elastix refinement. However, in regions where the local bend radius of the vessel approaches the local lumen radius, this same refinement can locally increase angular error even as it improves overall mask overlap (Fig. 6), because relocation without re-rotation is most consequential exactly where vector orientation changes fastest with position. Finally, the rotation-only vector transformation intentionally preserves velocity magnitude. This is useful for avoiding unstable magnitude changes, but it does not enforce physical conservation under local stretching or compression, so incompressibility of the flow field is not explicitly enforced by the transformation.

5.2. Future Work

Future work should extend the evaluation to measured flow in a larger set of real aortic cases and report robustness with respect to segmentation quality, centerline extraction, branch labeling, and anatomical variability. Detailed longitudinal and MRI–CFD studies remain important application targets, but they require application-specific validation. Methodologically, the model could be extended with explicit branch registration, separate true- and false-lumen correspondences for dissection cases, and richer cross-sectional shape matching. Further work should also investigate regularized Jacobian-based transformations to determine when they are physically meaningful or introduce artifacts. An alternative to registering the velocity field itself would be to compute derived hemodynamic quantities directly in the untransformed moving domain, where the original velocity field still strictly satisfies its governing conservation laws, and to map only these derived results onto the fixed grid. This would avoid vector reorientation entirely, at the cost of restricting downstream analysis to a fixed set of precomputed quantities rather than a fully queryable registered vector field. Future work should compare this alternative workflow with the vessel-frame-based transformation proposed here, to clarify the trade-offs between physical fidelity in the moving domain and flexible inspection of the flow field in the fixed domain.

In summary, this work demonstrates that an explicit vessel model can make aortic flow-field registration more geometrically meaningful. By separating anatomical correspondence from vector reorientation, the approach preserves flow direction more reliably than a generic image-registration baseline and can be combined with local anatomical refinement when realistic vessel shapes require closer mask alignment.

Acknowledgements

We used OpenAI Codex to assist with language polishing, grammar corrections, LaTeX formatting, and figure-code management. All code generated by Codex was reviewed, edited, and approved by the authors.

References

[AH11] ANGELELLI P., HAUSER H.: Straightening tubular flow for side-by-side visualization. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (Dec. 2011), 2063–2070. doi:10.1109/tvcg.2011.235. 2

[APBG01] ALEXANDER D. C., PIERPAOLI C., BASSER P. J., GEE J. C.: Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Transactions on Medical Imaging 20*, 11 (2001), 1131–1139. doi:10.1109/42.963816. 3

[BEGP18] BEHRENDT B., EBEL S., GUTBERLET M., PREIM B.: A framework for visual comparison of 4D PC-MRI aortic blood flow data. *Eurographics Workshop on Visual Computing for Biology and Medicine* (2018). doi:10.2312/VCBM.20181236. 1, 3

[BMGS13] BORN S., MARKL M., GUTBERLET M., SCHEUERMANN G.: Illustrative visualization of cardiac and aortic blood flow from 4D MRI data. In *Proc. of IEEE Pacific Visualization Symposium (PacificVis)* (Feb. 2013), IEEE, pp. 129–136. doi:10.1109/PacificVis.2013.6596137. 3

[BRPS*25] BÄUMLER K., ROLF-PISSARCZYK M., SCHUSSNIG R., FRIES T.-P., MISTELBAUER G., PFALLER M. R., MARSDEN A. L., FLEISCHMANN D., HOLZAPFEL G. A.: Assessment of aortic dissection remodeling with patient-specific fluid–structure interaction models. *IEEE Transactions on Biomedical Engineering 72*, 3 (2025), 953–964. doi:10.1109/TBME.2024.3480362. 6

[BVS*20] BÄUMLER K., VEDULA V., SAILER A., SEO J., CHIU P., MISTELBAUER G., CHAN F., FISCHBEIN M., MARSDEN A., FLEISCHMANN D.: Fluid–structure interaction simulations of patient-specific aortic dissection. *Biomechanics and Modeling in Mechanobiology 19*, 5 (Oct. 2020), 1607–28. doi:10.1007/s10237-020-01294-8. 1

[CMWY05] CAO Y., MILLER M. I., WINSLOW R. L., YOUNES L.: Large deformation diffeomorphic metric mapping of vector fields. *IEEE Transactions on Medical Imaging 24*, 9 (2005), 1216–1230. doi:10.1109/TMI.2005.853923. 3

[CWG*12] CLOUGH R. E., WALTHAM M., GIESE D., TAYLOR P. R., SCHAEFFTER T.: A new imaging method for assessment of aortic dissection using four-dimensional phase contrast magnetic resonance imaging. *Journal of Vascular Surgery 55*, 4 (2012), 914–23. doi:10.1016/j.jvs.2011.11.005. 1

[FM18] FAROUKI R. T., MOON H. P.: Rational frames of minimal twist along space curves under specified boundary conditions. *Advances in Computational Mathematics 44*, 5 (Mar. 2018), 1627–1650. doi:10.1007/s10444-018-9599-3. 2

[GAW*11] GLEICHER M., ALBERS D., WALKER R., JUSUFI I., HANSEN C. D., ROBERTS J. C.: Visual comparison for information visualization. *Information Visualization 10*, 4 (Sept. 2011), 289–309. doi:10.1177/1473871611416549. 1, 2

[KFW*02] KANITSAR A., FLEISCHMANN D., WEGENKITTL R., FELKEL P., GROLLER E.: CPR - curved planar reformation. In *IEEE Visualization* (2002), IEEE, pp. 37–44. doi:10.1109/visual.2002.1183754. 2

[KSM*10] KLEIN S., STARING M., MURPHY K., VIERGEVER M. A., PLUIM J. P. W.: elastix: A toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging 29*, 1 (2010), 196–205. doi:10.1109/TMI.2009.2035616. 2

[LKC94] LEE T.-C., KASHYAP R. L., CHU C.-N.: Building skeleton models via 3-D medial surface/axis thinning algorithms. *CVGIP: Graphical Models and Image Processing 56*, 6 (1994), 462–478. doi:10.1006/cgip.1994.1042. 3

[Mis13] MISTELBAUER G.: *Smart interactive vessel visualization in radiology*. PhD thesis, Faculty of Computer Science of the Technical University of Vienna, 2013. doi:10.34726/hss.2013.22462. 3

[MVE*22] MEUSCHKE M., VOSS S., EULZER P., JANIGA G., ARENS C., WICKENHÖFER R., PREIM B., LAWONN K.: COMFIS - comparative visualization of simulated medical flow data. In *Eurographics*

Workshop on Visual Computing for Biology and Medicine (2022). doi:10.2312/VCBM.20221185. 3

[Sal21] SALAMA A.: Velocity profile representation for fully developed turbulent flows in pipes: A modified power law. *Fluids 6*, 10 (2021), 369. doi:10.3390/fluids6100369. 6

[SDP13] SOTIRAS A., DAVATZIKOS C., PARAGIOS N.: Deformable medical image registration: A survey. *IEEE Transactions on Medical Imaging 32*, 7 (2013), 1153–1190. doi:10.1109/TMI.2013.2265603. 2

[SHO*26] SCHROEDER A., HORVAT A., OSTENDORF K., BÄUMLER K., MASTRODICASA D., FLEISCHMANN D., PREIM B., MISTELBAUER G.: Vessel-straightening-based registration for aortic flow comparison – test data. GitHub repository, 2026. URL: https://github.com/aaschr/vcbm_2026_registration_for_flow_comparison. 6

[SKW21] SHERK W. M., KHAJA M. S., WILLIAMS D. M.: Anatomy, pathology, and classification of aortic dissection. *Techniques in Vascular and Interventional Radiology 24*, 2 (June 2021), 100746. doi:10.1016/j.tvir.2021.100746. 1

[SPSP02] SELLE D., PREIM B., SCHENK A., PEITGEN H.-O.: Analysis of vasculature for liver surgical planning. *IEEE Transactions on Medical Imaging 21*, 11 (2002), 1344–1357. doi:10.1109/TMI.2002.801166. 2

[TS10] TAYLOR C., STEINMAN D.: Image-based modeling of blood flow and vessel wall dynamics: Applications, methods and future directions. *Annals of Biomedical Engineering 38*, 3 (Mar. 2010), 1188–203. doi:10.1007/s10439-010-9901-0. 1

[VP04] VERMA V., PANG A.: Comparative flow visualization. *IEEE Transactions on Visualization and Computer Graphics 10*, 6 (Nov. 2004), 609–624. doi:10.1109/tvcg.2004.39. 1, 2

[vPGL*14] VAN PELT R., GASTEIGER R., LAWONN K., MEUSCHKE M., PREIM B.: Comparative blood flow visualization for cerebral aneurysm treatment assessment. *Computer Graphics Forum 33*, 3 (June 2014), 131–140. doi:10.1111/cgf.12369. 3

[vPOBB*11] VAN PELT R., OLIVAN BESCOS J., BREEUWER M., CLOUGH R. E., GROLLER M. E., TER HAAR ROMENIJ B., VILANOVA A.: Interactive virtual probing of 4D MRI blood-flow. *IEEE Transactions on Visualization and Computer Graphics 17*, 12 (Dec. 2011), 2153–2162. doi:10.1109/TVCG.2011.215. 3

[WJZL08] WANG W., JÜTTLER B., ZHENG D., LIU Y.: Computation of rotation minimizing frames. *ACM Transactions on Graphics 27*, 1 (Mar. 2008), 1–18. doi:10.1145/1330511.1330513. 2

[YMN22] YUAN X., MITSIS A., NIENABER C. A.: Current understanding of aortic dissection. *Life 12*, 10 (Oct. 2022), 1606. doi:10.3390/life12101606. 1

[ZBL*25] ZIMMERMANN J., BÄUMLER K., LOECHER M., CORK T., MARSDEN A. L., ENNIS D. B., FLEISCHMANN D.: TBADFlow: FSI simulations and MRI data of type B aortic dissection models. Dataset, 2025. doi:10.25740/tz375fg1985. 6

[ZZWY25] ZHOU Y., ZHOU J., WANG Z., YANG G.: SATO: Straighten any 3D tubular object. *IEEE Transactions on Medical Imaging 44*, 9 (Sept. 2025), 3882–3891. doi:10.1109/tmi.2025.3571896. 2