

Image Based Rendering for Motion Compensation in Angiographic Roadmapping

Christian Unger

Martin Groher

Nassir Navab

Computer Aided Medical Procedures (CAMP), Technische Universität München, Germany

christian.unger@breakbe.at, groher@in.tum.de, navab@in.tum.de

Abstract

2D angiographic roadmapping is used frequently during image guided interventions to superimpose vessel structures onto currently acquired fluoroscopic images. While the fluoroscopic images, acquired with 12-15 frames per second, show patient bone anatomy as well as the current location of the inserted catheter, the roadmap delineates vessels to provide path information and to avoid accidental vessel wall punctures during catheter advancement.

This technique successfully reduces the injection of contrast agent, which is hazardous for the patient; however, it suffers from inaccuracy due to inevitable patient movement, which yields a misalignment of the roadmap laid over the current fluoroscopic frame.

We propose a method for rigid patient motion compensation via the trifocal tensor and Image Based Rendering (IBR). The method uses two contrasted and slightly shifted views and the current fluoroscopic frame.

Different to the existing solutions, we perform the motion compensation inherently in 3D, increasing reliability and accuracy of the resulting vascular rendering. Moreover, with the IBR technique, we avoid an explicit reconstruction, thus achieving reasonable results even for very small patient movements, which are common in interventional scenarios.

1. Introduction

The variety and application of minimally-invasive, endovascular treatments is increasing rapidly in hospitals. Common to all these interventions is the navigation of a catheter (or guidewire) through the human vascular system to a region of interest. This navigation is conducted under X-ray fluoroscopy visualizing injected catheter and bone anatomy. Since the catheter is to be guided through the vessel system, which is not visible under X-ray imaging, contrast agent is administered through the catheter to produce X-ray angiograms that visualize vascular structures.




Figure 1. Roadmapping on leg anatomy. Top-left: a contrasted view with dark disks showing markers attached to the patient skin; top-right: a non-contrast view (thin dark line is the catheter); bottom-left: a roadmap without compensation; bottom-right: a motion-compensated roadmap computed using our approach

Large amount of injected contrast can be toxic for the patient, thus only few contrasted images are acquired during an endovascular procedure. In order to avoid blind catheter navigation, *roadmapping* techniques have been proposed to fuse previously acquired angiograms (roadmaps) with current fluoroscopic image frames [19, 13].

In clinical practice, the de-facto standard for roadmapping is a mere overlay of 2D roadmap and current fluoroscopic frame. One major problem of this method occurs once the patient moves, which is very likely to happen during these time-consuming procedures, see Fig. 1. The movement causes the roadmap to be misaligned with the catheter visualized in the fluoroscopic frame, which disturbs navigation and may lead to accidental perforation of

the vessel wall - a serious injury to the patient. In neurological procedures or peripheral vessel studies this motion can be assumed to be rigid, i.e. is not accompanied by organ/vascular deformation as e.g. in abdominal catheterizations.

We propose a novel method for rigid motion compensation of angiographic roadmaps, which is based on view synthesis using the trifocal tensor [1]. We acquire two roadmap images, which are slightly shifted against each other. This shift can be caused either by a designated patient movement or by changing the viewpoint of the C-arm. A movement by the patient is preferred as it resembles the natural motion which is to be compensated for. The roadmap images are processed to yield graph representations of the vascular tree structure, composed of centerline and bifurcation points. Prior to the actual procedure, we attach at least 6 radio-opaque markers to the patient skin. The use of a fixed number of markers guarantees a known level of accuracy and reproducibility, which is required for all clinical solutions. We automatically detect and match these markers in all images, and compute the trifocal tensor to relate the two roadmap images with the current fluoroscopic (non-contrasted) frame, see Fig. 1. We extract the fundamental matrix between the two roadmap views from the tensor, and use it together with the graph properties of vasculature to determine dense correspondences on vessel centerlines. In this step we use the property of near projective invariance of 3D tubular structures [14] easing the difficult task of finding correspondences in transparent fluoroscopic X-ray images. From the centerline correspondences and the diameter information of the vascular structures, we render a synthetic vessel to the fluoroscopic frame using trilinear point transfer. This rendered vessel structure is inherently compensated for any 3D rigid transformation, see Fig. 1.

Since our method avoids an explicit 3D reconstruction of vessel structures, it successfully renders views for very small baselines and can thus compensate for small patient motion - a property, which is of particular importance in this specific scenario. Moreover, the compensation is carried out in 3D, different to existing angiographic motion compensation techniques that have been proposed in the context of coronary artery stabilization [6], or retrospective enhancement of digitally subtracted angiograms [15]. In some endovascular interventions, preoperative 3D data (e.g. Computed Tomography Angiography) is available and 2D roadmap motion compensation can be achieved by 2D-3D registration [9]. In most peripheral vessel studies or treatments, however, these high-resolution data is not acquired before the intervention, which is thus carried out using cost-effective intraoperative images only. Modern C-arms provide an interventional tomographic cone-beam reconstruction [23], which is, together with 2D-3D registration, used for roadmapping in some clinics [8]. Such routines are

accompanied with a dramatic increase of injected contrast agent. Our method offers an accurate 3D motion compensation without additional 3D data acquisition, or heavy increase in contrast administration. It provides a general solution to the problem of patient motion compensation, which can be used for stationary, as well as for low-cost mobile C-arm devices.

With regard to angiographic images, calibrated 3D-view geometry has been used for the 3D reconstruction of vessel networks as proposed in [22, 2]. To the best of our knowledge, utilizing the relationship of uncalibrated images for motion compensation of vessel structures has not been addressed so far.

The remainder of this paper is structured as follows. Section 2 introduces the different steps of the motion compensation algorithm, detailing 2D marker extraction, detection, and trifocal tensor computation, the determination of dense vessel centerline correspondences, and the vascular rendering method. Section 3 explains the experimental setup including the creation of ground truth data, the evaluation of rendering errors and its dependence on the number of markers and different sources of measurement noise in the rendering pipeline. These tests are followed by studies on real data, where a focus was laid on the derivation of meaningful quantitative measures that capture the quality of the rendering result. We conclude with a discussion in Section 4.

2. Method

Image based rendering as described in [1] uses the trifocal tensor to compute the position of points in the third view \mathcal{I}'' from a correspondence across the first and second view, \mathcal{I} and \mathcal{I}' . We compute the tensor from corresponding marker positions, which are extracted and matched automatically. Since we want to render vascular structures, we are only interested in corresponding points on vessels. With the centerline and diameter of the vessel, we are able to partition the vessel into small patches, whose back-projections are approximately cylindrical segments. The corresponding segments are subject to image based rendering. Since the projection of the vessel's 3D centerline is *nearly projectively invariant* [14] to the centerline of the 2D projected vessel we are justified to approximate the projected 3D centerline with the 2D centerline. This allows us to find corresponding points despite the transparent nature of X-ray angiograms. In order to automatically retrieve a synthesized vessel tree in view \mathcal{I}'' from views \mathcal{I} and \mathcal{I}' , we perform the following steps:

1. Detection and matching of markers in \mathcal{I} , \mathcal{I}' , and \mathcal{I}'' and computation of the trifocal tensor.
2. Extraction of two centerline graphs \mathcal{G} and \mathcal{G}' from the segmented vessels of the roadmap images \mathcal{I} , \mathcal{I}' . For

Table 1. Notation

\mathbf{X}	<i>bold upper-case letter</i> : point in 3D
\mathbf{x}	<i>bold lower-case letter</i> : point in 2D
x	<i>roman lower-case letter</i> : scalar
F_{21}	fundamental matrix between view \mathcal{I} and \mathcal{I}'
$d(\mathbf{x}, \mathbf{y})$	euclidean distance between \mathbf{x} and \mathbf{y}
$'$	<i>single prime</i> : entities of view \mathcal{I}'
$''$	<i>double prime</i> : entities of view \mathcal{I}''
$\mathbf{x} \leftrightarrow \mathbf{x}'$	correspondency: \mathbf{x} corresponds to \mathbf{x}'
$\tilde{\mathbf{x}}$	true point (without measurement noise)
$\hat{\mathbf{x}}$	point estimated with a computational method

our experiments, we use a semi-automatic segmentation technique as proposed in [3], however, any other (fully automatic) procedure might be employed, see for instance [18, 4]. The creation of graphs from vascular segmentations will be described shortly in 2.2.

3. Graph matching between \mathcal{G} and \mathcal{G}' and the computation of correspondences of all points lying on the centerlines of the vessels.
4. Image Based Rendering on small portions of the vasculature by dividing the vessel into small quadrilateral patches and warping each patch individually.

In the following, we will use the notation summarized in Table 1. All points will be represented by homogeneous coordinates.

2.1. Computation of the Trifocal Tensor

Before the procedure we manually attach X-ray-absorbing concentric balls to the skin of the patient in the area of interest. The projections of these balls can be observed as small dark disks distributed throughout the image, see Fig. 1, 2. We need to extract these markers in all three images \mathcal{I} , \mathcal{I}' and \mathcal{I}'' and match them in order to calculate the trifocal tensor from point-point-point correspondences [10, 20].

The extraction of markers is done using gradient-based Hough transform for circles. In the first step a Canny edge detection operator is applied to the input image and in the second step the Hough transform for circles is run on the filtered image. The projections of the markers are approximately of the same size and therefore the result of the Hough transform can be improved by using a restricted range of diameters.

Having computed the marker positions in the three views, we match corresponding ones with the RANSAC-algorithm [7, 11].

Finally the trifocal tensor is computed by a maximum likelihood estimation [20]. From an initial guess computed by the algebraic minimization algorithm [10] the following cost function is minimized iteratively using a Levenberg-Marquardt optimization over the entries of projection ma-

trices P' , P'' and the 3D reconstructions M_i of the marker correspondences:

$$\sum_{i=1}^{n_M} d(\mathbf{m}_i, \mathbf{P}M_i)^2 + d(\mathbf{m}'_i, \mathbf{P}'M_i)^2 + d(\mathbf{m}''_i, \mathbf{P}''M_i)^2 \quad (1)$$

2.2. Generating Centerline Correspondences

We extract the centerlines of the segmented vessels in views \mathcal{I} and \mathcal{I}' . The extracted centerlines are then converted to graphs: $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V', E')$, where $V \subset \mathbb{P}^2$ is the set of vertices, represented by their spatial coordinates, and $E \subset V \times V$ the set of edges. The individual shape of the centerline between adjacent vertices is also stored in each edge. Finally, we compute sparse and dense correspondences between the two graphs.

2.2.1 Graph Creation

An approximate skeleton from a given segmentation is extracted by a topological thinning algorithm as proposed in [16]. For creating a graph from this skeleton, we employ a wave propagation algorithm [24]. The graph \mathcal{G} is created with vertices at branching points of the vessel tree and edges in between. Points on the centerline between two adjacent vertices are stored as edge labels (to record the geometric progression of the edge) and are ordered with respect to a starting and end branching point: $E \ni e \mapsto \{\mathbf{c}_1, \dots, \mathbf{c}_{N(e)}\}$. Diameter information is computed using a euclidean distance map of the segmentation, and is stored in \mathcal{G} . In order to remove branches that are created during the topological thinning but are not part of the centerline, edges are removed that are shorter than the vessel diameter at the respective bifurcation position. Fig. 2 shows the extracted graph superimposed onto an angiographic image.




Figure 2. The graph data: vertices (red), edges (green) computed from the segmentation (blue) and markers (orange)

2.2.2 Graph Matching

The previous step is applied to views \mathcal{I} and \mathcal{I}' and produces two graphs $\mathcal{G} = (V, E)$ and $\mathcal{G}' = (V', E')$. Now,

a topological graph matching is done only on the bifurcation points of the two graphs in order to find a sparse set of correspondences between \mathcal{I} and \mathcal{I}' . Thereafter, centerline points are matched on corresponding edges (which have a corresponding start and end bifurcation) to create a dense set of centerline correspondences.

Finding Sparse Correspondences The core of our topological graph matching routine is based on the method of TSAI and FU [21]. This method is particularly suitable since it permits the integration of *geometrical* constraints derived from our three-view scenario. This approach employs a search algorithm to find the mapping with least cost. More precisely, it searches for a *subgraph isomorphism* between \mathcal{G} and \mathcal{G}' (and is therefore robust against outliers). The cost of a graph mapping μ is directly computed from it: the subgraph \mathcal{G}_μ of \mathcal{G} induced by the mapping μ can be transformed into a subgraph of \mathcal{G}' by performing several graph edit operations on \mathcal{G}_μ (such as vertex substitution/deletion or edge insertion/substitution/deletion). As every graph edit operation has a specific, pre-defined cost, the cost of a graph mapping is the sum of all induced edit operations. In our case, the edge insertion and deletion costs are set to a much larger value than the other costs, because we are mainly interested in preserving the topological information. In the worst case the method has a computational complexity of $\mathcal{O}(n^2m^n)$ (where $m = |V'|$ and $n = |V|$).

In its most general form every (not mapped) vertex in V' is considered as a possible match for every $\mathbf{v} \in V$. We can confine the set of possible matches, since we expect corresponding vertices to fulfill the epipolar constraint¹; consequently, only those vertices $\mathbf{v}' \in V'$ are allowed that are near to the epipolar line of \mathbf{v} (below a *distance threshold*). Geometrically speaking, for every $\mathbf{v} \in V$ the set of restricted possible matches is found along the epipolar line in view \mathcal{I}' . This improvement avoids the computation of geometrically invalid solutions and therefore it does not violate the optimality of the procedure (among the space of valid solutions, the mapping with least cost is expanded).

Another particularity is motivated by the entire creation process of the two graphs \mathcal{G} and \mathcal{G}' . As already mentioned, the contrasted vessels of two different views are first segmented and then converted into a graph representation. The propagation of the contrast agent usually differs between the two views and therefore there may be vessel branches that have a better or worse contrast in one view. These differences mainly affect the position of graph endpoints (vertices of degree one), but also bifurcation points may not be detected if whole branches are missing. Because of that, corresponding endpoints do not fulfill the epipolar constraint in general and corresponding bifurcations may not be present. We treat these issues by adding additional vertices into both

graphs: $\forall \mathbf{v} \in V$ and $\forall e' \in E'$ we add all points of intersection between the epipolar line $F_{21}\mathbf{v}$ and e' as new vertices into V' and split up e' respectively. The same procedure applies analogously to vertices $\mathbf{v}' \in V'$ and edges $e \in E$.

When the graph matching has finished, the method in general returns more than one solution that are optimal in a *topological* sense. From these solutions we choose the one that has a minimal *geometric* cost:

$$\sum_{(\mathbf{v}, \mathbf{v}') \in \mu} |\text{distanceFromEpipolarLine}(\mathbf{v}, \mathbf{v}')| \quad (2)$$

Dense Centerline Correspondences Given the mapping of the bifurcations, we want to find correspondences between centerline points on edges $e \in E$, $e' \in E'$. The dense correspondences are only generated on the centerline points of the vessel graph, neither on the whole image nor the segmented part of the image. Ideally this means to take every point on a centerline in view \mathcal{I} and to find the corresponding point on the corresponding centerline in view \mathcal{I}' . Due to noise or false correspondences the corresponding point may not always be found.

The basic idea is to simultaneously traverse centerline points on edges, whose start and end bifurcations have a correspondence in the second view. In the i -th step during this traversal, the current position on the centerline in view \mathcal{I} is \mathbf{c}_i . Together with the fundamental matrix F_{21} the corresponding position $\bar{\mathbf{c}}'_i$ on the centerline in view \mathcal{I}' must ideally lie on the epipolar line:

$$\bar{\mathbf{c}}'_i{}^T F_{21} \mathbf{c}_i = 0 \quad (3)$$

Since the centerlines are represented by a discrete set of image positions, we linearly interpolate between consecutive points via a function $f' : \mathbb{R} \rightarrow \mathbb{R}^3$:

$$t \mapsto f'(t) := \mathbf{c}'_{\lfloor t \rfloor} + (t - \lfloor t \rfloor)(\mathbf{c}'_{\lfloor t \rfloor + 1} - \mathbf{c}'_{\lfloor t \rfloor}) \quad (4)$$

with $\lfloor x \rfloor = \max\{n \mid n \in \mathbb{N}, n \leq x\}$. In the i -th step we are interested in $t_i \in \mathbb{R}$ fulfilling:

$$f'(t_i)^T F_{21} \mathbf{c}_i = 0 \quad (5)$$

From a geometric point of view (5) means that two consecutive positions of the centerline (e.g. \mathbf{c}'_k and \mathbf{c}'_{k+1}) have to lie in different half-planes separated by the epipolar line. Algebraically, with $\mathbf{l} = F_{21}\mathbf{c}_i$ the condition to be fulfilled reads as follows (with $k = \lfloor t_i \rfloor$):

$$f'(k)^T \mathbf{l} \cdot f'(k+1)^T \mathbf{l} \leq 0 \quad (6)$$

At this point, we have to be careful with the selection of such pairs, since condition (6) may be fulfilled arbitrarily often along an edge, depending on the shape of the vessel (imagine a curved vessel intersected multiple times by the

¹ $\mathbf{v}'^T F_{21} \mathbf{v} = 0$, iff $\mathbf{v} \leftrightarrow \mathbf{v}'$

epipolar line). Furthermore, it might also be the case that there is no pair for which equation (6) holds (if the epipolar line is erroneous, or the vessel has not been segmented properly, and thus no intersection occurs).

We effectively reduce these ambiguities by ensuring an ordering on the centerline points:

$$(\mathbf{c}_i \leftrightarrow f'(t_i)) \wedge (\mathbf{c}_j \leftrightarrow f'(t_j)) \wedge (i < j) \implies (t_i < t_j) \quad (7)$$

The described steps yield corresponding points on the centerlines of \mathcal{G} and \mathcal{G}' : $\mathbf{c}_i \leftrightarrow (\tilde{\mathbf{c}}'_i = f'(t_i))$.

2.3. Image Based Rendering

From the dense centerline correspondences we can compute the location of the centerline in view \mathcal{I}'' by using trifocal point transfer. The main idea for the rendering is to compute small quadrilateral areas of the projections of small parts of the vessels. Since a vessel can be approximated by small cylindrical segments whose projections will be of quadrilateral shape, the images of the vessels may be partitioned into quadrilateral patches.

In order to perform the rendering we identify a quadrilateral patch (a projection of a cylindrical segment) in view \mathcal{I} for which we can make use of the centerline and diameter information. Then, we find the four corresponding points in view \mathcal{I}' using trifocal point transfer. These correspondences fully determine the location of the quadrilateral in the third view \mathcal{I}'' . We employ a backward warping technique to render all patches to \mathcal{I}'' .

In order to compute a corner point \mathbf{p} of a quadrilateral in view \mathcal{I} , we exploit the diameter D_i and the direction \mathbf{d}_i (with \mathbf{d}_i^\perp geometrically perpendicular to \mathbf{d}_i and of unit length) of the centerline in point \mathbf{c}_i : $\mathbf{p}_{i,\{1,2\}} = \mathbf{c}_i \pm \frac{D_i}{2} \mathbf{d}_i^\perp$. The direction \mathbf{d}_i of the centerline in \mathbf{c}_i can be approximated with the difference of the preceding and successive centerline points, $\mathbf{d}_i = \mathbf{c}_{i+1} - \mathbf{c}_{i-1}$.

3. Experiments and Results

In our experiments we conduct two tests. First, we analyze the error as propagated to a rendered point in \mathcal{I}'' . This error depends on the number of markers used for tensor computation as well as on the error introduced in marker detection and centerline point creation and matching. This test is applied to synthetically created data sets with a known ground truth.

The second test applies the algorithm to real data sets where a reference rendering is given by administration of contrast agent in the third view \mathcal{I}'' . Here, meaningful measures for evaluating the quality of the IBR will be derived to show the algorithm's performance for large and small patient motion. In these tests, we used an epipolar line distance threshold of 15 pixels for the graph matching.

3.1. Monte-Carlo Tests

For testing the accuracy of our algorithm as well as its sensitivity to noise, we create synthetic data by first acquiring contrasted X-ray images from three views with a C-arm during a peripheral vessel study. We segment the vessel structures using the semi-automatic Graph Cut algorithm [3]. For retrieving synthetic data, we then compute the trifocal tensor with our method from marker correspondences $\mathbf{m}_i \leftrightarrow \mathbf{m}'_i \leftrightarrow \mathbf{m}''_i$, $1 \leq i \leq n$. Thereafter, we reconstruct the 3D marker positions \mathbf{M}_i , which are reprojected to get exact 2D marker correspondences $(\bar{\mathbf{m}}_i, \bar{\mathbf{m}}'_i, \bar{\mathbf{m}}''_i) = (\mathbf{P}\mathbf{M}_i, \mathbf{P}'\mathbf{M}_i, \mathbf{P}''\mathbf{M}_i)$. We also perform this reconstruction and reprojection for centerline correspondences to retrieve correct vessel points together with corresponding information. Note that this approach for synthetic data creation obeys the typical baseline setting of a real angiographic procedure.

3.1.1 Simulating Measurement Noise

From our scenario, we extracted four important parameters affecting the error introduced during the rendering process. First, the accuracy of the IBR with regard to the number of markers used to determine the trifocal tensor is evaluated. Second, different sources of measurement noise are introduced. Error was added to the location of the markers in all three views following an isotropic Gaussian distribution with zero mean and variance σ_M^2 . Moreover, 1D Gaussian noise was added to the centerline points of views \mathcal{I} and \mathcal{I}' with zero mean and a variance of σ_1^2 , and σ_2^2 , respectively. This error was applied perpendicular to the centerline direction since the algorithm regards the centerlines of both views as continuous curvilinear structures, where point selection and correspondence detection is not biased along the curve.

During the simulation we perform the following steps for 1000 times and average the propagated error:

1. Add noise to the marker positions: $(\bar{\mathbf{m}}_i, \bar{\mathbf{m}}'_i, \bar{\mathbf{m}}''_i) \mapsto (\mathbf{m}_i, \mathbf{m}'_i, \mathbf{m}''_i)$, $1 \leq i \leq n$ and compute the trifocal tensor \mathcal{T} from the modified markers.
2. Add noise to a random centerline position in view \mathcal{I} : $\bar{\mathbf{c}}_k \mapsto \mathbf{c}_k$.
3. Add noise to all centerline points in view \mathcal{I}' .
4. Compute the point of intersection of the modified centerline with $F_{21}\mathbf{c}_k$. This determines $\tilde{\mathbf{c}}'_k$.
5. Transfer the point to the third view: $\tilde{\mathbf{c}}''_k = \text{transfer}_{\mathcal{T}}(\mathbf{c}_k, \tilde{\mathbf{c}}'_k)$ and compute the euclidean distance of real and estimated centerline point: $d(\bar{\mathbf{c}}''_k, \tilde{\mathbf{c}}''_k)$.

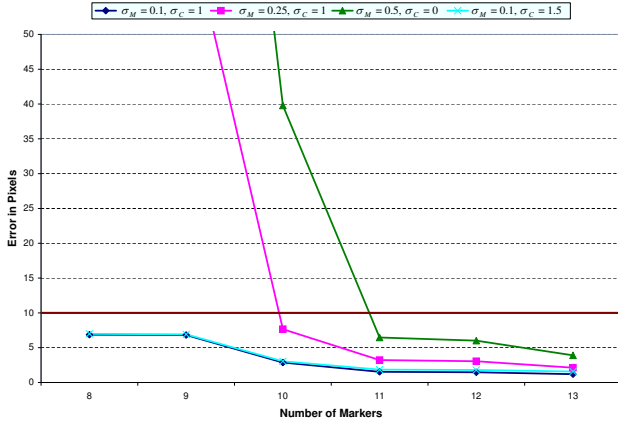


Figure 3. The influence of the number of markers on the trifocal transfer error.

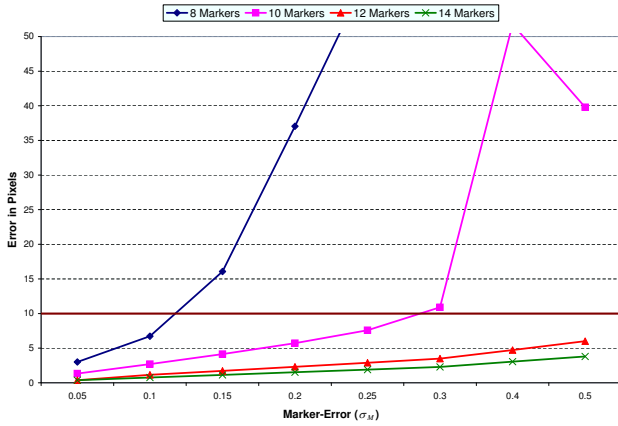


Figure 4. The influence of σ_M on the trifocal transfer error.

3.1.2 Results

The influence of the number of markers, the marker error defined by σ_M , and the centerline error defined by $\sigma_C = \sigma_1 = \sigma_2$ is shown in Fig. 3 and Fig. 4. Here, the transfer was regarded successful with a maximal error of 10 pixels², which is highlighted by a dark red line in the diagrams. As expected, the error observed in the third view drops when using more markers yielding an accurate solution for 8 markers with a low level of noise, and for 10 markers with a high level of noise. In clinical situations the number of markers that can be attached to the patient is limited, however, a conservative choice of 11 markers is still acceptable according to our clinical partners. Regarding noise, we observe from different test runs that the marker error σ_M has a great impact on the solution, while the centerline error σ_C does not influence the accuracy of the transfer much. The marker detection, however, can be made rather robust when using Hough transform for circles, making it possible to achieve subpixel accuracy [12].

²in a 1024×1024 image, pixel size 0.2mm × 0.2mm

Table 2. Similarity Coefficients for Patient Data Sets

	DSC	S_e	S_p	$\mu_{CLD} (\sigma_{CLD})$
Fig. 5 top	0.77	0.86	0.997	1.5 (1.9)
Fig. 5 bottom	0.79	0.75	0.997	1.9 (3.3)

3.2. Tests on Real Data

For testing the performance of the algorithm on clinical data, we acquired two data sets of a peripheral vessel study different in the extent of patient motion. In one case, the patient’s leg was nudged by the interventionalist resulting in a marker displacement between 10 and 25 pixels. In the other case, the patient was asked to hold still (displacement below 2 pixels). By successfully compensating both large and small patient motion, we demonstrate our method’s high “capture range”, proving the possibility for a seamless integration into clinical workflow. In this evaluation, we always acquire three *contrasted* views, thus creating a reference vessel structure in view \mathcal{I}'' for comparison.

The comparison is conducted by the evaluation of two quality measures, pixel coverage, and centerline distance. While the former evaluates the whole rendering process, the latter provides information on the quality of point transfer using the trifocal tensor, leaving out possible errors induced by the approximation of the centerline using quadrilaterals.

Pixel-Coverage During the IBR, we modify pixels in the third view \mathcal{I}'' . However, we touch only those pixels in \mathcal{I}'' that belong to the vasculature, i.e. that belong to the segmented part of the images \mathcal{I} and \mathcal{I}' . Therefore, the collection of all pixels modified by the image based rendering $\tilde{\mathcal{S}}''$, is an approximation of the reference segmentation \mathcal{S}'' of \mathcal{I}'' . For every pixel p there are four possible cases (the \in -relation between a pixel and a segmentation denotes that the pixel is part of the segmentation):

T_P : $p \in \tilde{\mathcal{S}}''$ and $p \in \mathcal{S}''$: true positives (green)

F_P : $p \in \tilde{\mathcal{S}}''$ and $p \notin \mathcal{S}''$: false positives (red)

F_N : $p \notin \tilde{\mathcal{S}}''$ and $p \in \mathcal{S}''$: false negatives (blue)

T_N : $p \notin \tilde{\mathcal{S}}''$ and $p \notin \mathcal{S}''$: true negatives

From these quantities we compute three similarity coefficients, the Dice Similarity Coefficient (DSC), Sensitivity (S_e), and Specificity (S_p). These coefficients are commonly used in the context of evaluating image segmentation results [5], and are particularly suitable for the comparison of reference segmentation and computed rendering:

$$DSC = \frac{2T_P}{2T_P + F_P + F_N} \quad (8)$$

$$S_e = \frac{T_P}{T_P + F_P} \quad (9)$$

$$S_p = \frac{T_N}{T_N + F_P} \quad (10)$$

Nearest-Neighbour Centerline-Distance For every centerline correspondence $\mathbf{c}_i \leftrightarrow \tilde{\mathbf{c}}'_i$ the location of the centerline $\tilde{\mathbf{c}}'_i$ in view \mathcal{I}'' is computed by trifocal point transfer. We are interested in the distance to the (unknown) true location of the centerline $\bar{\mathbf{c}}'_i$, which can be approximated by computing the closest symmetric centerline position:

$$\operatorname{argmin}_k d(\tilde{\mathbf{c}}'_i, \mathbf{c}''_k) + \operatorname{argmin}_i d(\bar{\mathbf{c}}'_i, \mathbf{c}''_k) \quad (11)$$

$\mathbf{c}''_k, 1 \leq k \leq N''$ are the points of the centerline of the reference segmentation in view \mathcal{I}'' . Consequently every transferred centerline position implies a distance. For the comparison the mean (μ_{CLD}) and standard deviation (σ_{CLD}) of all summed distances are computed.

Besides the analysis of these quality measures for our proposed method, we also compare our results to explicit vessel reconstruction and reprojection techniques. Here, we want to demonstrate the suitability of the IBR method for small baseline changes that occur when the patient is moving. In particular, we test our method against a 3D reconstruction technique using the optimal triangulation as described in [11], page 318, followed by a projection into view \mathcal{I}'' . The fundamental matrix is estimated from the two views \mathcal{I} and \mathcal{I}' using an ML approach.

3.2.1 Results

Fig. 5 illustrates the pixel coverage of the two renderings for large and small patient motion.





Figure 5. Resulting qualitative pixel coverage (above: large motion, 11 markers, below: small motion, 12 markers). Blue areas are *missed* regions of the ground truth. Green areas are *covered* regions of the ground truth. Red areas are *wrongly drawn* regions.

Some general remarks on the pixel coverage: the red regions should be as small as possible, since they represent *wrongly drawn* parts and may be a risk for the patient (there are no vessels in reality, but the compensated view shows some). The blue regions correspond to *missing* parts (there are vessels in reality, but they have not been rendered) and

are not a risk for the patient, but are inconvenient for the surgeon. Finally, the green regions are those parts of the image where vessels have been rendered correctly. The resulting pixel coverage for Fig. 5 can be taken from Tab. 2. The comparison with explicit reconstructions of the same scenes confirms our expectation that three views are crucial for compensating patient motion in angiographic interventions. In fact, the reconstruction with fundamental matrices computed from two views fails completely at transferring the vasculature. However, a projection of the vessel tree, which is reconstructed via fundamental matrices extracted from the trifocal tensor yields results similar to our inherent compensation method. Fig. 6 illustrates the considerable difference between the fundamental matrices and underlines the necessity for three views in order to robustly estimate the motion.



(a) Trifocal Tensor

(b) Fundamental Matrices

Figure 6. Epipolar lines using the trifocal tensor (left) and triangulation (right). The epipolar lines correspond to (manually selected) marker positions in view \mathcal{I} . In this case, rotational motion happened along a vertical axis – therefore the result of the trifocal tensor is much more meaningful.

4. Concluding Remarks

In this paper, we present a novel method for the compensation of rigid patient motion in fluoroscopic images using two previously acquired contrasted images. With radio-opaque markers attached to the patient, we recover the geometric relationship between the two previously acquired contrasted roadmap views and each incoming fluoroscopic frame. An intelligent use of vessel topology and geometry allows us to compute correspondences of the vasculature and to render it at the right location into the non-contrasted fluoroscopic frames. Thereby, we not only circumvent the difficult task of finding reliable correspondences in transparent X-ray images, but also present a solution, which like any clinically acceptable interventional solution results in guaranteed, reproducible outcome. Different to existing techniques for motion compensation, we inherently solve for a 3D transformation without using previously acquired 3D information, or increasing contrast agent injection con-

siderably. We keep our method as general as possible, in order to make it easy to use for any commercial C-arm system without need for performing a tedious device calibration.

Key issues for a successful motion compensation are the number and configuration of available markers. The previous section showed that 11 markers yield a stable rendering result in all cases. Our clinical partners confirm that the presence of these small circular structures in the fluoroscopic images (see Fig. 1) does not affect their interventional workflow. In fact, the markers will not influence their ability in localizing the catheter as well as the superimposed roadmap. As the markers are attached by the operating physician, the obvious degenerate configurations are easy to avoid. And the chance of having 11 points on complex degenerate configurations [17, 11] is negligible.

The rough segmentation of vessel structures within a contrasted angiographic image is a prerequisite of our algorithm. Different methods have been proposed to accomplish such vessel segmentation on fluoroscopic images without user interaction (see e.g. [18, 4]). Since our method provides a robust search for corresponding points on vessel structures including outlier detection and removal, even an algorithm yielding a suboptimal segmentation can be employed. Currently, we are acquiring the ethical permission to evaluate this solution on a controlled set of patients.

Acknowledgements. The authors want to thank Johannes Rieger and Marcus Treitl from Klinikum Innenstadt, Ludwig-Maximilian Universität München for supporting us with clinical data and their medical expertise.

References

- [1] S. Avidan and A. Shashua. Novel view synthesis by cascading trilinear tensors. *IEEE Trans. on Vis. and Comp. Graphics*, 4(4):293–306, 1998.
- [2] C. Blondel, R. Vaillant, F. Devernay, G. Malandain, and N. Ayache. Automatic trinocular 3D reconstruction of coronary artery centerlines from rotational X-ray angiography. In *CARS*, 2002.
- [3] Y. Boykov and G. Funka-Lea. Graph cuts and efficient N-D image segmentation. *IJCV*, 70(2):109–131, 2006.
- [4] A. Can, H. Shen, J. Turner, H. Tanenbaum, and B. Raysam. Rapid automated tracing and feature extraction from retinal fundus images using direct exploratory algorithms. *IEEE Trans. on Information Technology in Biomedicine*, 3(2):125–138, 1999.
- [5] R. Cardenes, M. Bach, Y. Chi, I. Marras, R. de Luis, M. Anderson, P. Cashman, and M. Bultelle. Multimodal evaluation for medical image segmentation. In *Computer Analysis of Images and Patterns*, pages 229–236, 2007.
- [6] M.-P. Dubuisson-Jolly, C.-C. Liang, and A. Gupta. Optimal polyline tracking for artery motion compensation in coronary angiography. In *ICCV*, 1998.
- [7] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. of the ACM*, 24(6):381–395, 1981.
- [8] S. Gorges, E. Kerrien, M.-O. Berger, Y. Troussset, J. Pescatore, R. Anxionnat, and L. Picard. Model of a vascular C-Arm for 3D augmented fluoroscopy in interventional radiology. In *MICCAI*, pages 214–222, 2005.
- [9] M. Groher, F. Bender, R. Hoffmann, and N. Navab. Segmentation-driven 2D-3D registration for abdominal catheter interventions. In *MICCAI*, 2007.
- [10] R. Hartley. Minimizing algebraic error in geometric estimation problems. In *IEEE CVPR*, pages 469–476, 1998.
- [11] R. Hartley and A. Zisserman. *Multiple View Geometry in computer vision*. Cambridge University Press, 2nd edition, 2003.
- [12] Q. Ji and Y. Xie. Randomised hough transform with error propagation for line and circle detection. *Pattern Analysis and Applications*, 6:55–64, 2003.
- [13] M. Kukuk and S. Napel. Rotational roadmapping: A new image-based navigation technique for the interventional room. In *MICCAI*, pages 636–643, 2007.
- [14] A. Liu, E. Bullitt, and S. Pizer. 3D/2D registration via skeletal near projective invariance in tubular objects. In *MICCAI*, pages 952–963, 1998.
- [15] E. Meijering. *Image Enhancement in Digital X-ray Angiography*. PhD thesis, Utrecht University, 2000.
- [16] K. Palágyi, E. Sorantin, E. Balogh, A. Kuba, C. Halmi, B. Erdohelyi, and K. Hausegger. A sequential 3D thinning algorithm and its medical applications. In *IPMI*, pages 409–415, 2001.
- [17] A. Shashua and S. Maybank. Degenerate n point configurations of three views: Do critical surfaces exist? Tr 96-19, Hebrew University, Computer Science, 1996.
- [18] C. Steger. An unbiased detector of curvilinear structures. *IEEE Trans. PAMI*, 20(2):113–125, 1998.
- [19] J. Tobis, W. Johnston, S. Montelli, E. Henderson, W. Roeck, B. Bauer, O. Nalcioglu, and W. Henry. Digital coronary roadmapping as an aid for performing coronary angioplasty. *Am. J. of Cardiology*, 56:237–241, 1985.
- [20] P. H. S. Torr and A. Zisserman. Robust parameterization and computation of the trifocal tensor. *Image and Vision Computing*, 15:591–605, 1997.
- [21] W. H. Tsai and K. S. Fu. Subgraph error-correcting isomorphisms for syntactic pattern recognition. *IEEE Trans. Syst. Man Cybern.*, 13:48–61, 1983.
- [22] C. Venaille, D. Mischler, R. Collorec, J. Catros, and J. Coatrieux. Automatic 3D-reconstruction of vascular networks from three projections: A simulated annealing approach. In *IEEE Engineering in Medicine and Biology*, 1989.
- [23] K. Wiesent, K. Barth, N. Navab, P. Durlak, T. Brunner, O. Schuetz, and W. Seissler. Enhanced 3-D-reconstruction algorithm for C-arm systems suitable for interventional procedures. *IEEE Trans. Med. Imag.*, 19(5):391–403, 2000.
- [24] C. Zahlten, H. Jürgens, and H. O. Peitgen. Reconstruction of branching blood vessels from CT-data. In *Eurographics Workshop of Visualization in Scientific Computing*, pages 161–168, 1994.