

Residual Motion Compensation in ECG-Gated Interventional Cardiac Vasculature Reconstruction

C Schwemmer^{1,2}, C Rohkohl³, G Lauritsch³, K Müller^{1,2} and J Hornegger^{1,2}

¹Pattern Recognition Lab, Department of Computer Science, Friedrich-Alexander-Universität Erlangen-Nürnberg, Martensstr. 3, D-91058 Erlangen, Germany

²Erlangen Graduate School in Advanced Optical Technologies (SAOT), Paul-Gordan-Str. 6, D-91052 Erlangen, Germany

³Siemens AG, Healthcare Sector, Siemensstr. 1, D-91301 Forchheim, Germany

E-mail: chris.schwemmer@cs.fau.de

Abstract. 3-D reconstruction of cardiac vasculature from angiographic C-arm CT (rotational angiography) data is a major challenge. Motion artefacts corrupt image quality, reducing usability for diagnosis and guidance. Many state-of-the-art approaches depend on retrospective ECG-gating of projection data for image reconstruction. A trade-off has to be made regarding the size of the ECG-gating window. A large temporal window is desirable to avoid undersampling. However, residual motion will occur in a large window, causing motion artefacts.

We present an algorithm to correct for residual motion. Our approach is based on a deformable 2-D-2-D registration between the forward projection of an initial, ECG-gated reconstruction, and the original projection data. The approach is fully automatic and does not require any complex segmentation of vasculature, or landmarks. The estimated motion is compensated for during the backprojection step of a subsequent reconstruction. We evaluated the method using the publicly available CAVAREV platform and on six human clinical datasets. We found a better visibility of structure, reduced motion artefacts, and increased sharpness of the vessels in the compensated reconstructions compared to the initial reconstructions. At the time of writing, our algorithm outperforms the leading result of the CAVAREV ranking list. For the clinical datasets, we found an average reduction of motion artefacts by $13\pm 6\%$. Vessel sharpness was improved by $25\pm 12\%$ on average.

This is an author-created, un-copyedited version of an article accepted for publication in Physics in Medicine and Biology. IOP Publishing Ltd is not responsible for any errors or omissions in this version of the manuscript or any version derived from it. The Version of Record is available online at <http://dx.doi.org/10.1088/0031-9155/58/11/3717>

1. Introduction

In complex interventional cardiac procedures, three-dimensional information can provide easier assessment and better guidance (Hetterich et al. 2010, Gutleben et al. 2011). In current clinical practice, this 3-D imaging is performed by moving the patient to a CT scanner or acquiring a CT prior to the intervention. Ideally, 3-D imaging should be performed in the interventional suite using C-arm CT. Imaging during the intervention would make up-to-date information available. In the cardiac environment, such an approach is currently limited by the temporal resolution of available C-arm systems. Due to the slow rotation speed, resulting in acquisition times of several seconds, heart motion corrupts a straightforward 3-D reconstruction. Streak artefacts, motion blur, and reduced sharpness and visibility of structure are typical problems associated with such an approach.

If an ECG signal is recorded simultaneously with the acquisition, a retrospective gating of X-ray projection data can be performed: Only images from a specific heart phase contribute to the reconstruction (Desjardins & Kazerooni 2004). However, ECG data does not necessarily correspond to the exact motion state of the heart (Desjardins & Kazerooni 2004). In addition, ECG-gating sparsifies data and therefore reconstruction becomes a strongly ill-posed problem, limiting the quality and usability of the resulting 3-D data. To compensate for this, not a strict gating, but a temporal window with a certain width around the selected motion state is used (Rohkohl et al. 2008). A wider gating window is desirable to get a high signal-to-noise ratio and little undersampling artefacts, but then the residual motion in the gated projection data corrupts image quality. Therefore, gating is a trade-off between undersampling artefacts (narrow gating window) and motion artefacts (wide gating window).

In the literature, several approaches have been proposed to account for residual motion due to non-ideal ECG-gating. A general distinction can be made between those that do a full 3-D motion estimation and compensation (Blondel et al. 2006, Schomberg 2007, Keil et al. 2009, Prümmer et al. 2009, Tang et al. 2012), and those that perform a 2-D estimation (Taguchi et al. 2007, Hansis, Schäfer, Dössel & Grass 2008, Schwemmer et al. 2012). 3-D motion compensation corrects motion in image space, while a 2-D method works in projection space. Due to possible overlay of structures along the X-ray beam direction, a 2-D method is only an approximation. All structures along the beam are affected by the same transformation during motion

compensation. On the other hand, a full 3-D estimation is a strongly ill-posed problem with high computational demands. Motion periodicity assumptions and/or regularisation are therefore often used (Zeng et al. 2005, Hansis et al. 2009). However, the object of interest here, i.e. contrasted vasculature, is very sparse compared to the image volume. In addition, due to the ECG-gating, the residual motion within the gating window can be assumed to be reasonably small. Therefore an approximate 2-D estimation in projection space might be sufficient.

In the literature, a couple of projection-based methods can be found, which are discussed shortly. If landmarks are available in the acquired data (e.g. vascular stents or catheter-based markers), these can be tracked and a landmark-based registration can be performed (Perrenot et al. 2007). However, the availability of such landmarks cannot be assumed for all data. Recently, a model-based learning approach was proposed as another method. A previously learnt model is registered to the actual data (Lebois et al. 2011). This means that an extensive training phase is needed in advance. A method based on registration in projection space was proposed in (Hansis, Schäfer, Dössel & Grass 2008). It requires a segmentation of vasculature centrelines in the acquired projection data, which is difficult (Jandt et al. 2009).

In this paper, a projection-based method for compensation of residual motion in ECG-gated data is presented. We estimate residual motion by a deformable, multi-resolution, 2-D–2-D registration in projection space without requiring complex pre-processing steps like vessel centreline segmentation. The method is fully automatic, no user interaction is required.

In the next two sections, our method, the experimental setup, and the evaluation process are discussed in detail. In Section 4, the results are presented, followed by a discussion in Section 5. Section 6 concludes this paper.

2. Methods

In this section, the individual steps of our algorithm are discussed in detail. The outline is given by Figure 1. In step 1, an initial ECG-gated reconstruction is performed. In step 2, a thresholding operation removes non-vascular structures from the volume. This sparse volume is then forward-projected using the acquisition geometry. In these forward projection images (FwP), the region of interest, i.e. the region containing the contrasted vascular structure, is determined automatically. In step 3, the original projection images are pre-processed using a morphological

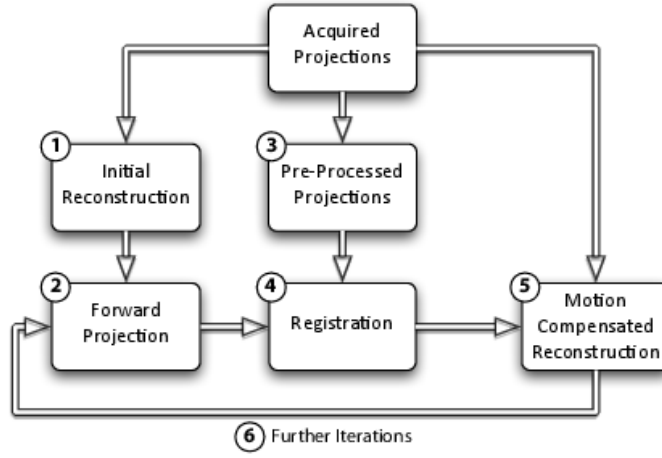


Figure 1: Illustration of the proposed algorithm.

top-hat filter and a thresholding operation to reduce background structures. In step 4, the FwP are registered to the pre-processed original projection images using deformable 2-D-2-D registration in a multi-resolution scheme. Finally, in step 5, the estimated 2-D motion field per projection is compensated for in a subsequent ECG-gated reconstruction. In step 6, the procedure may be repeated for additional refinement.

2.1. Initial Reconstruction

We perform an initial ECG-gated reconstruction by inserting a weighting function λ into a standard FDK-type algorithm (Schäfer et al. 2006, Rohkohl et al. 2008). Let $h_r \in [0, 1]$ be the reference heart phase, at which reconstruction shall be carried out. Heart phase is the relative distance between two consecutive R-peaks in the ECG. The ECG-gated FDK reconstruction $f_{h_r} : \mathbb{R}^3 \mapsto \mathbb{R}$ at a voxel $\mathbf{x} \in \mathbb{R}^3$ is given by

$$f_{h_r}(\mathbf{x}) = \sum_{i=1}^N \lambda(i, h_r) \cdot w(i, \mathbf{x}) \cdot p_F(i, A(i, \mathbf{x})) \quad , \quad (1)$$

where N is the number of projection images, $w : \mathbb{N} \times \mathbb{R}^3 \mapsto \mathbb{R}$ is the FDK distance weight and $p_F(i, \mathbf{u}) : \mathbb{N} \times \mathbb{R}^2 \mapsto \mathbb{R}$ is the filtered, redundancy- and cosine-weighted projection data of the i -th image at pixel position \mathbf{u} . The pixel position for backprojection to the voxel \mathbf{x} is determined by the perspective projection

$A : \mathbb{N} \times \mathbb{R}^3 \mapsto \mathbb{R}^2, (i, \mathbf{x}) \mapsto A(i, \mathbf{x}) = \mathbf{u}$. The perspective projection A can be calculated using pre-calibrated projection matrices (Wiesent et al. 2000). The ECG-gating weighting function λ used here is a cosine-window defined by

$$\lambda(i, h_r) = \begin{cases} \cos^a \left(\frac{d(h(i), h_r)}{\omega} \pi \right) & \text{if } d(h(i), h_r) \leq \frac{\omega}{2} \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where $h(i)$ is the heart phase of the i -th projection image according to the ECG, $\omega \in [0, 1]$ controls the width and $a \geq 0$ controls the shape of the gating window. The distance measure d is defined as $d(h_1, h_2) = \min_{j \in \{0, 1, -1\}} |h_1 - h_2 + j|$.

Since the resulting reconstruction suffers from undersampling artefacts, we also perform a streak reduction. A formal derivation of this post-processing step was presented in (Rohkohl et al. 2008). In practice, streak reduction is integrated into the reconstruction as follows:

$$\hat{f}_{h_r}(\mathbf{x}) = \sum_{i=1+N_{\text{ign}}}^{N-N_{\text{ign}}} \lambda(j_i, h_r) \cdot w(j_i, \mathbf{x}) \cdot p_{\text{F}}(j_i, A(j_i, \mathbf{x})) \quad , \quad (3)$$

with an enforced ordering of $\lambda(j_1, h_r) \cdot w(j_1, \mathbf{x}) \cdot p_{\text{F}}(j_1, A(j_1, \mathbf{x})) \leq \dots \leq \lambda(j_N, h_r) \cdot w(j_N, \mathbf{x}) \cdot p_{\text{F}}(j_N, A(j_N, \mathbf{x}))$ and $j_i \in \{1, \dots, N\}$, $N_{\text{ign}} \in \mathbb{N}_0$. This means the N_{ign} smallest and largest contributions to each voxel are ignored during reconstruction.

2.2. Forward Projection Generation

2.2.1. Thresholding. Only contrasted vascular structure is of interest for the registration algorithm, which means only high-contrast structure needs to be retained. Since a selective contrasting is performed during the acquisition, vascular structure is the predominant high-contrast structure in the image. Therefore, a simple thresholding operation can be employed to remove background structures: Only the $t_r \in [0, 1]$ percentile of the largest voxel values is retained, while the other voxels are set to 0. We denote the thresholded initial reconstruction with $\hat{f}'_{h_r}(\mathbf{x})$.

2.2.2. Forward Projection. After the thresholding of the initial reconstruction, maximum intensity forward projection images $p_{\text{fwp}}(i, \mathbf{u}) : \mathbb{N} \times \mathbb{R}^2 \mapsto \mathbb{R}$ are generated using the original acquisition geometry. Only those projections with a gating weight

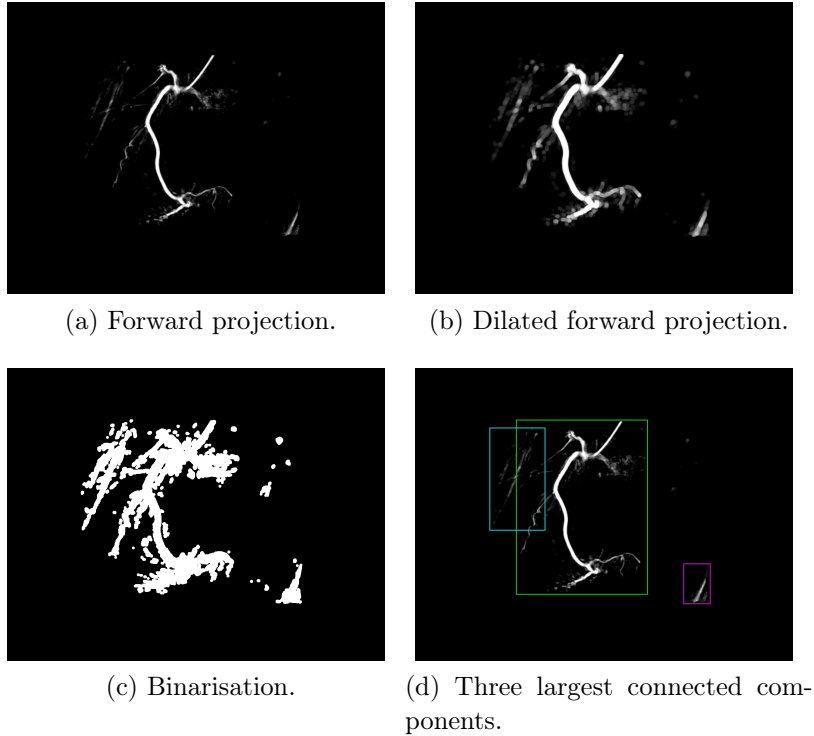


Figure 2: Automatic ROI determination.

$\lambda(i, h_r) > 0$ need to be generated, since others will not be used by the final reconstruction. We use a ray-casting approach as in (Galigekere et al. 2003):

$$p_{\text{fwp}}(i, \mathbf{u}) = \max_{\mathbf{x} \in L_{i, \mathbf{u}}} \hat{f}'_{h_r}(\mathbf{x}) \quad , \quad (4)$$

where $L_{i, \mathbf{u}} = \{\mathbf{x} \in \mathbb{R}^3 \mid A(i, \mathbf{x}) = \mathbf{u}\}$ is the set of voxels representing the virtual X-ray beam from the i -th source position to pixel position \mathbf{u} on the detector. FwP are generated in parallel on graphics hardware (Weinlich et al. 2008).

2.2.3. Region of Interest Detection. The registration process can be sped up and stabilised by defining a region of interest (ROI) for the evaluation of the objective function. This ROI is automatically determined from the FwP by applying the following algorithm to every forward projected image (cf. Figure 2):

- (i) Dilate the image using morphological operations (Dougherty & Lotufo 2003) with a circular structuring element.




Figure 3: Original (left) and pre-processed (right) projection image. Images were cropped to the ROI shown in Figure 6b for better visualisation.

- (ii) Binarise the image by setting every pixel with an intensity > 0 to 1.
- (iii) Detect and label connected components, compute component sizes as the number of pixels in each component.
- (iv) Keep only the largest object and calculate its rectangular bounding box.

The final ROI is the bounding box covering the bounding boxes from all projection images, with an added safety margin to account for incomplete structure in the initial reconstruction. Steps 1–4 can be executed in parallel for all FwP, only the final ROI calculation needs to be done after all individual bounding boxes are found.

2.3. Pre-Processing of Original Projections

Since the contrasted vessels are small objects with a larger scale background, a background reduction of the original projection images $p(i, \mathbf{u})$ improves vessel contrast and increases stability of the registration process. We use a morphological top-hat filter (Hansis, Schäfer, Dössel & Grass 2008) for this step. The filter has a circular structuring element of radius r . Let $\mathcal{N}_r(\mathbf{u})$ be the neighbourhood around pixel position \mathbf{u} defined by the structuring element. Then the top-hat filtered image is given by

$$p_{\text{th}}(i, \mathbf{u}) = p(i, \mathbf{u}) - \max_{\mathbf{u}' \in \mathcal{N}_r(\mathbf{u})} \tilde{p}_{\text{th}}(i, \mathbf{u}') \quad \text{and} \quad (5)$$

$$\tilde{p}_{\text{th}}(i, \mathbf{u}) = \min_{\mathbf{u}'' \in \mathcal{N}_r(\mathbf{u})} p(i, \mathbf{u}'') \quad . \quad (6)$$

This means that morphological opening (erosion followed by dilation) is performed and the result is subtracted from the original image. The effect of this operation is that all image structures larger than the structuring element are removed.

After filtering, a thresholding operation retains only the $t_p \in [0, 1]$ percentile of the largest pixel values. After both steps, most of the non-vascular background is removed (cf. Figure 3) and we denote the pre-processed projection images with $p_{\text{bgr}}(i, \mathbf{u})$.

2.4. Image Registration

During image registration, a mapping between the space of the pre-processed projection images $p_{\text{bgr}}(i, \mathbf{u})$ and the FwP $p_{\text{fwp}}(i, \mathbf{u})$ is established. After registration, $p_{\text{bgr}}(i, \mathbf{u})$ is similar to $p_{\text{fwp}}(i, \mathbf{M}(i, \mathbf{u}))$, where $\mathbf{M}(i, \mathbf{u}) = \mathbf{u}' : \mathbb{N} \times \mathbb{R}^2 \mapsto \mathbb{R}^2$ is the motion vector field for the i -th image. Similarity is defined by the objective function of the registration algorithm. We employ a multi-resolution scheme with an affine motion model on all resolution levels, and an additional uniform cubic B-spline motion model (Unser 1999) on the highest level only:

$$\mathbf{M}(i, \mathbf{u}) = \mathbf{M}_{\text{affine}}(i, \mathbf{u}) + \mathbf{M}_{\text{spline}}(i, \mathbf{u}) \quad . \quad (7)$$

This separation of affine and deformable motion as well as the multi-resolution scheme reduce the susceptibility of the deformable registration to local minima and also increase convergence speed. The affine motion model has 6 degrees of freedom, representing anisotropic scaling, shear, rotation and translation. The B-spline motion model is parameterised by the number of control points c in each dimension, resulting in $2 \cdot c^2$ degrees of freedom. The choice of c influences the relation of smoothness (small c) to flexibility (large c) of the motion model.

The resize factor between the different resolution levels was 2. Downscaling of the images between resolution levels was performed by bilinear interpolation and pre-smoothing with a Gaussian kernel ($\sigma = 1$).

We used normalised cross-correlation (NCC) as the similarity measure, which is a common measure for multi-modality registration problems (Russakoff et al. 2003). Purely intensity-based metrics like sum of squared differences do not work for this problem, since the grey values of $p_{\text{bgr}}(i, \mathbf{u})$ and $p_{\text{fwp}}(i, \mathbf{u})$ differ due to the gated

reconstruction and the maximum intensity forward projection. NCC is insensitive to these grey value differences, while being less computationally intensive than mutual information or similar multi-modality metrics. The NCC is only evaluated within the ROI, which also decreases computation time. NCC for projection image i is defined as

$$\text{NCC}_i(p_{\text{bgr}}, p_{\text{fwp}}, \mathbf{M}) = \frac{1}{n-1} \left\langle \frac{\mathbf{P}_{i,\text{bgr}}}{\|\mathbf{P}_{i,\text{bgr}}\|}, \frac{\mathbf{P}_{i,\text{fwp}}}{\|\mathbf{P}_{i,\text{fwp}}\|} \right\rangle, \quad (8)$$

where n is the number of pixels in the ROI, $\langle \star, \star \rangle$ is the inner product and $\|\star\|$ the L^2 norm. $\mathbf{P}_{i,\text{bgr}}(\mathbf{u}) = p_{\text{bgr}}(i, \mathbf{u}) - \overline{p_{\text{bgr}}}(i)$ and $\mathbf{P}_{i,\text{fwp}}(\mathbf{u}) = p_{\text{fwp}}(i, \mathbf{M}(i, \mathbf{u})) - \overline{p_{\text{fwp}}}(i)$, where $\overline{p_\star}(i)$ is the mean intensity of an image. The transformed images $p_{\text{fwp}}(i, \mathbf{M}(i, \mathbf{u}))$ were only computed at those pixels necessary for the evaluation of the NCC, i.e. within the ROI. Bilinear interpolation was used.

Registration was driven by a gradient descent optimisation method. Since we perform registration on a per-image basis, all projections can be processed in parallel.

2.5. Motion Compensated Reconstruction

After registration, the motion vector field $\mathbf{M}(i, \mathbf{u})$ is known for every projection image. The estimated motion can be compensated for in the backprojection step

$$f_{h_r, \mathbf{M}}(\mathbf{x}) = \sum_{i=1}^N \lambda(i, h_r) \cdot w(i, \mathbf{x}) \cdot p_{\text{F}}(i, \mathbf{M}(i, \mathbf{A}(i, \mathbf{x}))) . \quad (9)$$

Backprojection can be implemented very efficiently on graphics hardware (Hofmann et al. 2011). The same is true for B-spline evaluation (Ruijters et al. 2008). Therefore, the motion compensated backprojection can be carried out completely parallelised on the graphics card.

The same streak reduction as in Equation 3 is also used for the motion compensated reconstruction. For the sake of conciseness, this is not shown in Equation 9, since the streak reduction is not the focus here.

2.6. Further Iterations

The motion compensation algorithm can be used in an iterative manner: The output from step 2.5 can be used as input for step 2.2. Registration accuracy may improve when using input images that contain less artefacts and better contrast than the initial reconstruction.

Table 1: Clinical datasets used for the evaluation.

Dataset	3-D Img. Vol. [mm ³]	Heart Rate [bpm]	#Projs. Used
LCA1	140 ² × 101	77 ± 0.1	53
LCA2	152 ² × 107	58 ± 0.4	53
LCA3	152 ² × 114	52 ± 0.7	54
RCA1	152 ² × 110	68 ± 1.5	53
RCA2	131 ² × 109	71 ± 2.1	53
RCA3	143 ² × 119	54 ± 1.9	54

3. Experimental Setup and Evaluation Methods

3.1. Experimental Setup

3.1.1. CAVAREV. CAVAREV (Rohkohl et al. 2010) is a publicly available platform for the evaluation of cardiac vasculature reconstruction algorithms. We used the cardiac motion-only dataset for this experiment, since we assume a breath-hold protocol. This dataset consists of 133 simulated projection images created from a software phantom that shows a thorax and contrasted left and right coronary arteries. Each projection image has a size of 960×960 pixels and an isotropic pixel size of 0.32 mm. The reconstructed 3-D volumes have an isotropic voxel size of 0.5 mm and a size of 98^3 mm³. The gating parameters were selected as $\omega = 0.4$ and $a = 4$, with $h_r = 0.90$ and $N_{\text{ign}} = 3$.

3.1.2. Human Clinical Datasets. Six human clinical datasets were used for the evaluation (cf. Table 1): In LCA1, LCA2 and LCA3, a left coronary artery was imaged. The patient in dataset LCA2 had a total occlusion in the proximal part of the LAD, which means that no contrast agent reached the LAD beyond this point. In RCA1, RCA2 and RCA3, a right coronary artery was imaged. All datasets were acquired using a five second rotational angiography with selective contrast agent administration (1–2 ml/s) on a Siemens zeego C-arm device (Siemens AG, Healthcare Sector, Forchheim, Germany). The datasets are a representative cross-section, fulfilling the prerequisites: No patient and breathing motion and good contrast in the vessels over the whole scan.

Source-isocentre-distance was ~ 80 cm and source-detector-distance ~ 120 cm. Each dataset consists of 133 projection images with a size of 1240×960 pixels and an isotropic pixel size of 0.308 mm. The reconstructed 3-D volumes have an isotropic voxel size of 0.5 mm. The gating parameters were selected as $\omega = 0.4$ and $a = 4$, with $h_r = 0.75$ and $N_{\text{ign}} = 3$ for all human datasets.

3.1.3. Common Parameters. Thresholding was performed at $t_r = 0.005$ and $t_p = 0.2$. The size of the morphological kernel for top-hat filtering was $r = 3.85$ mm. Dilation for ROI detection was performed using a circular structuring element with a radius of 1.54 mm. The size of the safety margin for the final bounding box was ~ 3 mm. The number of B-spline control points was set to $c = 6$ in each dimension. We employed a multi-resolution registration scheme using 3 levels, with only the affine transform on the lower two levels and the additional deformable B-spline transform on the highest level. The maximum number of optimisation steps was set to 400 for the affine (200 on each level) and 250 for the deformable registration. Optimisation was stopped if the gradient magnitude of the NCC was below $3 \cdot 10^{-4}$. Two iterations of the whole algorithm were performed for each dataset. Registration times were obtained on two Intel® Xeon® E5540 CPUs (2.53 GHz) with 16 threads in total and 16 GB of memory. The graphics hardware used for forward and backprojections was an NVIDIA® Quadro® FX 5800 GPU with 4 GB of graphics memory.

3.2. Evaluation

Qualitative evaluation was carried out visually. Quantitative evaluation of the CAVAREV dataset was done using the metric provided by the platform, which is explained shortly in the following subsection. For a quantitative evaluation of the human datasets, we employed two metrics: Image artefacts in a sub-volume that encloses the reconstructed vessels, and vessel sharpness. Since we perform high-contrast imaging and use a streak-reduced image reconstruction, streak artefacts can be controlled (and suppressed by the display window) reasonably well. Therefore, we did not measure the streak artefact level of the reconstructed images.

3.2.1. CAVAREV Metric. Since our goal is the reconstruction of a single image volume at a specific reference motion state h_r , we used the metric introduced as

Q_{3D} in (Rohkohl et al. 2010). A short overview is given here, full details can be found in (Rohkohl et al. 2010). For each of the 133 time points (corresponding to the 133 projection images) of the CAVAREV dataset, a static binary volume f_i^{morph} representing the coronary vasculature at the corresponding motion state exists. A reconstruction f submitted to CAVAREV is compared to all 133 ground truth volumes in the following way:

$$Q_i(f) = \max_{a \in \{0, \dots, 255\}} \text{dsc} \left(f_i^{\text{morph}}, T(f, a) \right) , \quad (10)$$

where $T(f, a)$ is a threshold function that binarises f at threshold a :

$$T(f, a)(\mathbf{x}) = \begin{cases} 1 & f(\mathbf{x}) \geq a \\ 0 & f(\mathbf{x}) < a \end{cases} \quad (11)$$

The similarity is evaluated using the Dice similarity coefficient

$$\text{dsc}(f_1, f_2) = 2 \frac{\sum_x f_1(x) \cdot f_2(x)}{\sum_x f_1(x) + f_2(x)} , \quad (12)$$

which ranges from 0 (no overlap) to 1 (perfect spatial overlap) since f_1 and f_2 are binary. $Q_i(f)$ is called the motion phase-dependent reconstruction quality. Since CAVAREV accepts reconstructions with 8 bit quantisation, all possible thresholds are tried in this way and the best is used for the quality metric at this time point i . This simulates a manual thresholding operation. The quality measure Q_{3D} is then computed as

$$Q_{3D}(f) = \max_{i \in \{1, \dots, 133\}} Q_i(f) . \quad (13)$$

3.2.2. Image Artefacts. We estimated image artefacts caused by residual motion using a noise estimator defined in (Yang & Tai 2010). This approach excludes image structures from the noise level calculation. We used it to quantify motion blur (smoothed out structure) and ghosting artefacts (smooth artefacts in areas without structure). The original approach is a 2-D measure. For the evaluation of volume data, we performed a slice-wise calculation on HU-scaled reconstructed images. Slicing was performed along the z-axis.

First, a slice is divided into non-overlapping blocks and over- or under-saturated blocks are excluded. Artefacts can drive parts of a quantised image into saturation. Therefore, as in (Yang & Tai 2010), blocks with a mean intensity within 7% of the boundaries of the data range are disregarded.

Then, homogeneous blocks are identified with the help of a Sobel edge detector and an automatic threshold (using the lower 10% percentile, as in (Yang & Tai 2010)). The Sobel edge detector basically computes the first derivative of the block, its 3×3 and 5×5 convolution kernels are given in (Yang & Tai 2010). After the thresholding operation, only blocks without strong edges remain.

Finally, a noise estimate for each remaining block is calculated as (Yang & Tai 2010)

$$\sigma_n = \left(\frac{\pi}{2}\right)^{\frac{1}{2}} \frac{1}{6N_H^*} \sum_{B_H^*} |\mathbf{I}(x, y) \otimes \mathbf{L}_A| \quad , \quad (14)$$

where B_H^* is the set of voxels belonging to all homogeneous, non-saturated blocks, N_H^* is the number of voxels in B_H^* , $\mathbf{I}(x, y)$ the voxel value and \mathbf{L}_A a convolution kernel. Following (Yang & Tai 2010), the kernel is defined to be


$$\mathbf{L}_A = \begin{bmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{bmatrix} \quad . \quad (15)$$

We used the same block size of 5×5 voxels as in the original publication. The region where noise was estimated in was a cuboid containing the reconstructed vascular tree. Its volume was $\sim 1 \cdot 10^6 \text{ mm}^3$ (left coronary arteries) and $\sim 7 \cdot 10^5 \text{ mm}^3$ (right coronary arteries), respectively.

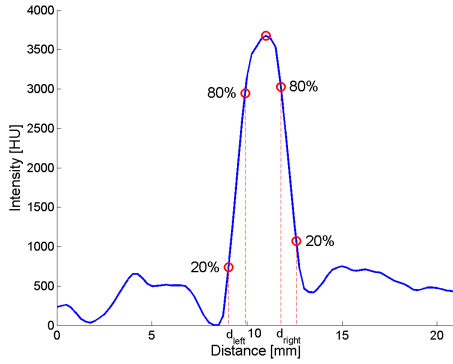
Our implementation of the algorithm is available online at (Schwemmer 2012).

3.2.3. Vessel Sharpness. For the evaluation of the vessel sharpness, we used a measure introduced in (Li et al. 2001) and extended it to a semi-automatic evaluation of whole vessel branches (Figure 4 illustrates the process):


Semi-Automatic Placement of Measurement Lines. First, a semi-automatic centreline segmentation of one vessel branch using a manually placed seed point is performed for each reconstructed volume (Friman et al. 2008) (cf. Figure 4a). A cubic B-spline is fit to the detected centreline for a continuous representation. Regularly spaced evaluation points e_i are placed on this spline with a spacing of 1 mm. At each point, 10 radial line segments perpendicular to the spline are distributed with an angular spacing of 20° . Each line segment has a length of 21 mm. We denote the set of 10 profile lines at e_i with L_i . Figure 4b shows an example for this.



(a) Centreline extracted from dataset RCA3.



(c) Profile plot for one line segment. Red circles denote 20%, 80% and maximum points.



(b) Line segments at 5 evaluation points along the centreline.

Figure 4: Vessel sharpness measurement process.

Profile Measurement. The following algorithm is applied to each profile line l at each evaluation point (cf. Figure 4c) (Li et al. 2001). It is important to note that all intensity measurements are taken from the original, non-windowed volumes.

- (i) Find the maximum intensity value along the line. This separates the profile line into two parts, which are denoted left and right (of the maximum).
- (ii) Find the minimum intensity values of the left and right part.
- (iii) Find the 20% and 80% points between minimum and maximum for both parts individually.
- (iv) Let d_{left} and d_{right} be the distances between the respective 20% and 80% points.




Figure 5: Coronary artery segments used for the sharpness evaluation.

(v) Vessel sharpness for this profile line is then

$$s_l = \frac{2}{d_{\text{left}} + d_{\text{right}}} . \quad (16)$$

For the maximum detection, we constrained the search space to 1/8 of the profile line around its centre.

Vessel Sharpness. We define vessel sharpness at an evaluation point e_i as

$$s_{e_i} = \frac{1}{10} \sum_{L_i} s_l . \quad (17)$$

Using the individual sharpness values s_{e_i} , we characterised the sharpness of a reconstruction in two ways: The overall mean sharpness is the mean value of all s_{e_i} of that reconstruction. In addition, we calculated mean sharpness values for segments along the vessel branch that follow the 15-segment coronary artery classification system (Austen et al. 1975). Left coronary arteries were evaluated at segments 6–8 for the LAD, segments 11, 13 and 14 for the LCX, and segment 5 for the main branch. Right coronary arteries were evaluated at segments 1–4. Figure 5 shows an overview of the segments used.




Figure 6: Automatically determined ROIs drawn over the pre-processed original projection data.

4. Results

4.1. Region of Interest Detection

The bounding box was always calculated successfully for the complete coronary tree. All vessels that were visible in the initial reconstruction and thus in the FwP were included in the bounding boxes. Figure 6 shows two typical examples of the final regions of interest, drawn over the pre-processed original projection data. We observed that in the presence of strong motion, small distal vessel parts can extend beyond the calculated ROI (cf. arrow in Figure 6b). This is due to the incomplete nature of the initial reconstruction. A larger safety margin could possibly include these parts, but a compromise between a tight ROI and including all distal vessel needs to be made. On the other hand, the ROI only defines the area for the calculation of the similarity measure. The motion vector field is defined over the whole image.

The average size of the ROIs over all datasets was $36 \pm 7\%$ of the full projection image size. This resulted in an average reduction of computation time by 41% for left coronary arteries and 24% for right coronary arteries compared to registration without an ROI.

Table 2: Registration time and number of optimisation steps per projection for left (LCA) and right (RCA) coronary arteries. The number of steps shown correspond to two affine registrations and the deformable registration, respectively.

	Time [s]	1 st Iter.		2 nd Iter.		
		Number of Steps		Time [s]	Number of Steps	
		Affine	Deformable		Affine	Deformable
LCA	8 ± 2	104 ± 55	31 ± 27	11 ± 3	111 ± 61	40 ± 35
RCA	29 ± 7	120 ± 67	99 ± 80	28 ± 9	122 ± 69	109 ± 74




Figure 7: Average initial NCC value by heart phase of projection image. The NCC values of each dataset were normalised to $[0, 1]$ before averaging.

4.2. Registration

Table 2 lists the average registration time and number of optimisation steps per projection by type of dataset for the human clinical data. The number of steps until convergence for the affine registration was similar for both types of datasets and in the first and second iteration. The number of steps needed for the deformable registration was 2.5–3 times as much for the RCA datasets as for the LCA datasets, with a very high standard deviation. This increase is also almost identically reflected by the registration times, meaning the deformable registration dominates the total time. An analysis of the time spent in the affine and deformable iterations showed

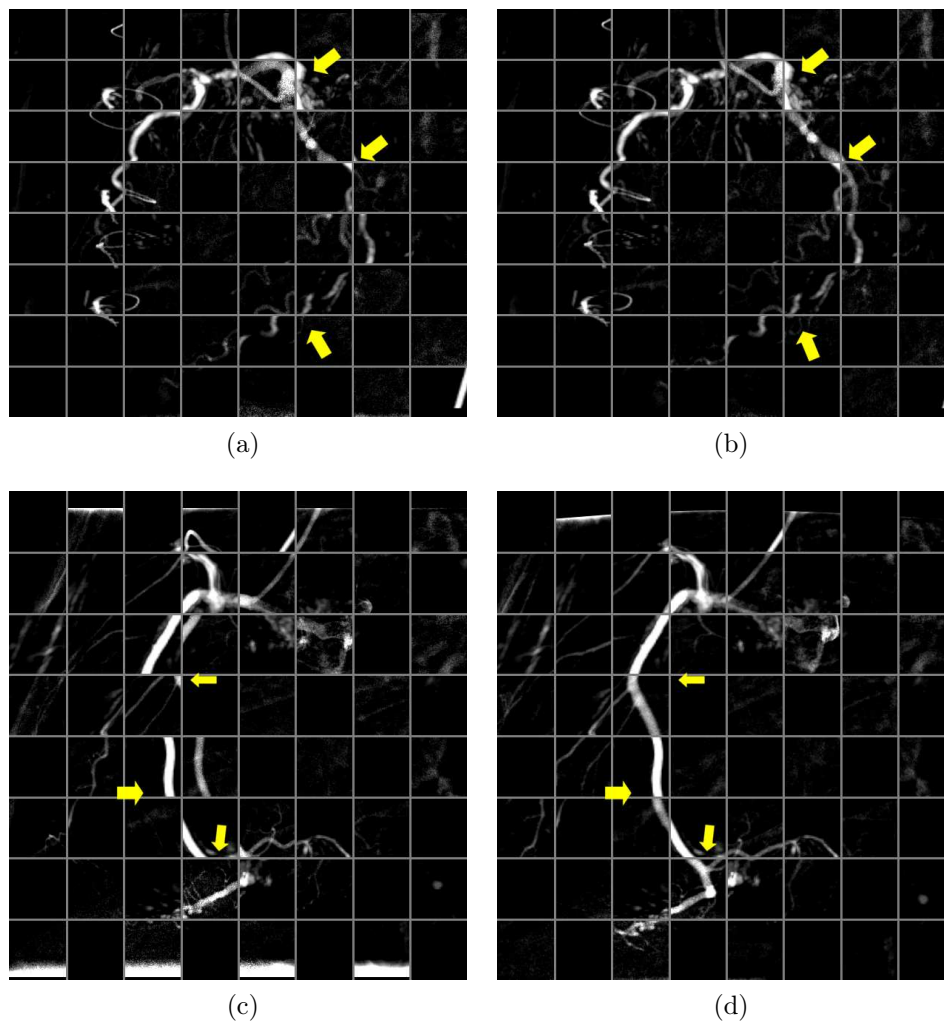


Figure 8: Checkerboard visualisation of $p_{\text{bgr}}(i, \mathbf{u})$ and $p_{\text{fwip}}(i, \mathbf{u})$ before (left) and after (right) registration. Top row: LCA1. Bottom row: RCA3. Arrows indicate locations with large displacement before registration.

that the deformable iterations are responsible for 96% of the total registration time over all datasets.

Between datasets, we observe that the initial reconstructions of the left coronary artery datasets had better image quality than the right coronary artery datasets. We expect less residual motion in the LCA cases to be the cause of this. This results in an easier and faster registration. Within each dataset, a similar observation can be made. Images in the middle of the gating window, i.e. with a heart phase very close or identical to h_r , only contain little or no residual motion ($h_r = 0.75$ is within a resting phase of the heart cycle). Images close to the ends of the window contain more, increasing the effort for registration. This effect can be seen in Figure 7, where the average NCC value before registration is shown with respect to the heart phase. It can also be seen that for the RCA datasets, the initial NCC value for $h(i) > h_r$ decreased faster than for the LCA datasets, which is also an indication for an increased registration effort. The second iteration took longer than the first for all datasets. We attribute this to a better resolution of fine details in the forward projection, which increases the effort for registration.

Figure 8 shows an overlay of $p_{\text{bgr}}(i, \mathbf{u})$ and $p_{\text{fwp}}(i, \mathbf{u})$ in a checkerboard pattern before and after two iterations of our algorithm. It can be seen that the displacement of vessel sections between $p_{\text{bgr}}(i, \mathbf{u})$ and $p_{\text{fwp}}(i, \mathbf{u})$ is drastically reduced after registration. Also, the large displacement in Figure 8c (indicated by arrows) is completely compensated, as seen in Figure 8d.

4.3. Qualitative Image Evaluation

4.3.1. CAVAREV. Figure 9 shows the resulting reconstructions before and after one iteration of our algorithm. No visual difference between one and two iterations could be found for the CAVAREV dataset. A good reduction of motion-related, low-frequency artefacts can be observed after motion compensation. In addition, both homogeneity and sharpness of the vascular structures are increased.

4.3.2. Human Clinical Datasets. Figures 10 and 11 show the resulting reconstructions before, after one and after two iterations of our algorithm. A strong reduction of image artefacts can be observed for all datasets after the first iteration. At the same time, depiction and sharpness of vascular structures are increased. This is also true for the artificial valve in dataset LCA3, although this was not explicitly




Figure 9: Reconstruction results of the CAVAREV dataset. The grey scale window was 1000 HU.

part of this study. Note that the segments 1 and 5 are corrupted in all reconstructions. This image degradation is caused by data inconsistency. Parts of the contrast agent leak into the aortic root and accumulate in this region. The resulting non-static clouds reduce the achievable reconstruction quality.

A second iteration of the algorithm does not reduce image artefacts noticeably anymore. Instead, the visibility and sharpness of distal vessels is increased. Arrows in Figures 10 and 11 indicate where the most improvement after the second iteration can be seen. For example, the posterolateral branch of LCA1 (top row of Figure 10) is only partially visible in the initial reconstruction, whereas it is almost fully recovered after two iterations. Another example for the right coronary arteries is the first right ventricular branch of RCA3 (bottom row of Figure 11). It is almost not visible in the initial reconstruction, but recovered after two iterations.

4.4. Quantitative Image Evaluation

4.4.1. CAVAREV. We achieved $Q_{3D} = 0.776$ for both one and two iterations of our algorithm. For reference, the published reconstruction quality of standard FDK is $Q_{3D} = 0.431$. For streak-reduced, ECG-gated reconstruction at $h_r = 0.9$, which is equivalent to our initial reconstruction, $Q_{3D} = 0.744$ (Rohkohl 2012). The latter algorithm is the best result in the public CAVAREV ranking at the time of this writing.




Figure 10: Reconstruction results of left coronary arteries. Left: Initial reconstruction. Middle: Compensated (1 iteration). Right: Compensated (2 iterations).

Top row: LCA1 (left posterior oblique view). Middle row: LCA2 (left posterior oblique view). Bottom row: LCA3 (left anterior oblique view).

The grey scale window was 1000 HU. Arrows indicate locations with major reconstruction improvements.




Figure 11: Reconstruction results of right coronary arteries. Left: Initial reconstruction. Middle: Compensated (1 iteration). Right: Compensated (2 iterations). Top row: RCA1 (left anterior oblique view). Middle row: RCA2 (left posterior oblique view). Bottom row: RCA3 (left sagittal view).

The grey scale window was 1000 HU. Arrows indicate locations with major reconstruction improvements.

Table 3: Estimated low-frequency artefact level around the reconstructed vasculature in HU.

	LCA1	LCA2	LCA3	RCA1	RCA2	RCA3
Initial	14	20	14	18	13	13
1 Iter.	12	18	12	14	12	12
2 Iter.	12	18	12	14	12	12

Table 4: Average vessel sharpness in 1/mm.

	LCA1		LCA2		LCA3		RCA1	RCA2	RCA3
	LAD	LCX	LAD	LCX	LAD	LCX			
Initial	0.41	0.368	0.324	0.354	0.453	0.423	0.358	0.375	0.447
1 Iter.	0.482	0.446	0.443	0.479	0.512	0.511	0.5	0.457	0.483
2 Iter.	0.486	0.464	0.451	0.484	0.524	0.516	0.515	0.46	0.484

Therefore, our algorithm achieved an improvement of 80% compared to standard FDK and 4% compared to the best public result.

Since the cardiac motion of the dataset is strictly periodic, ECG-gating already performs very well. Due to this, the initial reconstruction is already good enough so that all data inconsistencies are removed in the first iteration. A second iteration of our algorithm does not improve the quality anymore.

4.4.2. Human Clinical Datasets. Table 3 lists the estimated low-frequency artefact level for each dataset and number of iterations. The first iteration yields an average artefact reduction of about 13%, the second iteration does not result in a noticeable artefact reduction.

Vessel sharpness values are presented both as mean values for a complete vessel branch (Table 4) and as mean values for individual segments (Table 5). Over all datasets, one iteration increases vessel sharpness by 23%, two iterations by 25%. After motion compensation, the overall vessel sharpness is more similar between datasets than before (standard deviation between datasets of 0.026 after compensation vs. 0.045 before). An interesting example is dataset RCA1 (cf. Table

Table 5: Vessel sharpness of individual segments along the vessel tree in 1/mm.

(a) LCA1							
Segment	5	6	7	8	11	13	14
Initial	0.296	0.300	0.358	0.467	0.318	0.368	0.343
1 Iter.	0.323	0.385	0.430	0.558	0.372	0.442	0.447
2 Iter.	0.319	0.381	0.433	0.569	0.373	0.445	0.488

(b) LCA2							
Segment	5	6	7	8	11	13	14
Initial	0.319	0.260	N/A	N/A	0.236	0.368	0.375
1 Iter.	0.344	0.303	N/A	N/A	0.250	0.515	0.438
2 Iter.	0.337	0.306	N/A	N/A	0.245	0.521	0.443

(c) LCA3							
Segment	5	6	7	8	11	13	14
Initial	0.238	0.355	0.464	0.484	0.402	0.415	0.493
1 Iter.	0.266	0.352	0.522	0.548	0.454	0.538	0.573
2 Iter.	0.276	0.352	0.514	0.569	0.461	0.534	0.584

(d) RCA1				(e) RCA2					
Segment	1	2	3	4	Segment	1	2	3	4
Initial	0.219	0.387	0.348	0.330	Initial	0.325	0.354	0.429	0.391
1 Iter.	0.330	0.499	0.486	0.519	1 Iter.	0.374	0.469	0.494	0.449
2 Iter.	0.379	0.501	0.519	0.541	2 Iter.	0.382	0.465	0.495	0.453

(f) RCA3				
Segment	1	2	3	4
Initial	0.295	0.448	0.445	0.449
1 Iter.	0.345	0.500	0.490	0.450
2 Iter.	0.338	0.497	0.495	0.452

5d and Figure 11a–11c). Segments 1 and 4 are heavily affected by motion artefacts. After two iterations, their sharpness is improved by 73% and 64%. Although the initial sharpness of segment 4 is inferior to segments 2 and 3, after motion compensation the sharpness of segment 4 is actually the highest.

Overall, our observation from the qualitative evaluation, that a second iteration especially improves distal vessels, is confirmed here.

5. Discussion

The parameters of our algorithm were chosen heuristically. They achieve high reconstruction quality while limiting the registration time. By visual inspection, we observed no trapping in local minima during registration with the presented settings. At the same time, an increase in the number of resolution levels did not decrease computation time substantially, because the coarse alignment is carried out very quickly in contrast to the deformable registration on the finest level. Also, the use of the deformable motion model on coarser resolution levels did not improve convergence speed noticeably. We assume that the main computational cost of the deformable motion estimation is within the highest level for the data and parameters presented. Since the parameter set performed equally well for all datasets, we do not expect it needs to be adjusted for new datasets as long as the type of data (i.e. contrasted coronary arteries) and acquisition protocol stay the same.

Using the smallest possible ROI for each projection instead of the union of all individual ROIs could further increase efficiency. On the other hand, this has an influence on the stability of the registration: The ROI detection partly worked so well, because the union of all individual ROIs eliminates possible misdetections in single projection images.

The benefit of the approach in Equation 9 is that the application of the motion model is a simple coordinate transform. The only interpolation of pixel values is during the evaluation of $p_F(i, \mathbf{u})$, which is necessary in any case.

We found that the achievable quality after motion compensation is relatively robust against the quality of the initial reconstruction. That means that cases with strong residual motion artefacts benefit the most. Of course, the lower bound for this compensation is that the initial reconstruction needs to contain any structure at all. If, for example, strong breathing motion is present, data will have such a high inconsistency that the initial reconstruction does not contain any structure that can

be registered at all.

Our presented method and evaluation are tailored to the properties and needs of interventional imaging of coronary vasculature. This allows the use of preprocessing methods like top-hat filtering and thresholding that utilise the thin, sparse and high-contrast nature of the object of interest. Different cardiac structures, like heart chambers, would need other imaging protocols and motion estimation algorithms (Müller et al. 2012).

6. Conclusion

We presented an algorithm to correct for residual motion in ECG-gated coronary vasculature reconstruction. This residual motion is caused by an imperfect gating, which leads to the usage of data from different motion states. A smaller gating window is not desirable to avoid an increase of undersampling artefacts. Motion compensation can help regarding the trade-off between undersampling and motion artefacts. Our method does a projection-based 2-D motion estimation using image registration. The forward projection of an initial, ECG-gated reconstruction is registered to the pre-processed original projection data in a multi-resolution scheme without requiring any complex segmentation or user interaction. The region of interest for the image registration is determined automatically for each dataset. The resulting motion vector field is then used for a subsequent, motion-compensated reconstruction.

We evaluated the algorithm using the publicly available CAVAREV platform and on six clinical datasets, both qualitatively and quantitatively. After the first iteration, the artefact level is reduced and the sharpness and visibility of vascular structures is enhanced. A second iteration does not noticeably change the artefact level anymore, but increases sharpness and visibility of distal vessels in the clinical datasets. For the CAVAREV dataset, a Q_{3D} of 0.776 could be achieved, which is better than the leading result of the public ranking list at the time of this writing. For the clinical datasets and compared to the initial reconstruction, motion artefacts could be reduced by $13\pm 6\%$ on average. Vessel sharpness could be improved by $25\pm 12\%$ on average.

Further work may include the optimisation of the registration implementation to further reduce the computation time needed for this step. To this end, an implementation of the similarity measure and motion model on a GPU are expected to yield the highest gain.

Acknowledgments

The authors gratefully acknowledge funding of the Erlangen Graduate School in Advanced Optical Technologies (SAOT) by the German Research Foundation (DFG) in the framework of the German excellence initiative. We also would like to thank Prof. Dr. D. Böcker and Dr. P. Skurzewski, St. Marienhospital Hamm, Germany for providing clinical data, and our colleague Dr.-Ing. Andreas Maier for his valuable suggestions during the writing of this article.

Disclaimer: The concepts and information presented in this paper are based on research and are not commercially available.

References

- Austen W, Edwards J, Frye R, Gensini G, Gott V, Griffith L, McGoon D, Murphy M & Roe B 1975 A reporting system on patients evaluated for coronary artery disease *Circulation* **51**(4), 5–40.
- Blondel C, Malandain G, Vaillant R & Ayache N 2006 Reconstruction of coronary arteries from a single rotational X-ray projection sequence *IEEE Trans. Med. Imaging* **25**(5), 653–663.
- Desjardins B & Kazerooni E 2004 ECG-gated cardiac CT *Am. J. Roentgenol.* **182**(4), 993–1010.
- Dougherty E & Lotufo R 2003 *Hands-on Morphological Image Processing* Vol. TT59 of *Tutorial Texts in Optical Engineering* SPIE Bellingham, WA, USA.
- Friman O, Hindennach M & Peitgen H O 2008 Template-based Multiple Hypotheses Tracking of Small Vessels in ‘5th International Symposium on Biomedical Imaging: From Nano to Macro (ISBI 2008)’ IEEE Paris, France pp. 1047–1050.
- Galigekere R, Wiesent K & Holdsworth D 2003 Cone-beam Reprojection Using Projection-Matrices *IEEE Trans. Med. Imaging* **22**(10), 1202–1214.
- Gutleben K J, Nölker G, Ritscher G, Rittger H, Rohkohl C, Lauritsch G, Brachmann J & Sina A M 2011 Three-dimensional coronary sinus reconstruction-guided left ventricular lead implantation based on intraprocedural rotational angiography: a novel imaging modality in cardiac resynchronisation device implantation *Europace* **13**(5), 675–682.
- Hansis E, Schäfer D, Dössel O & Grass M 2008 Projection-based motion compensation for gated coronary artery reconstruction from rotational x-ray angiograms *Phys. Med. Biol.* **53**(14), 3807–3820.
- Hansis E, Schäfer D, Dössel O & Grass M 2008 Evaluation of iterative sparse object reconstruction from few projections for 3-D rotational coronary angiography *IEEE Trans. Med. Imaging* **27**(11), 1548–1555.
- Hansis E, Schomberg H, Erhard K, Dössel O & Grass M 2009 Four-Dimensional Cardiac Reconstruction from Rotational X-ray Sequences: First Results for 4D Coronary Angiography in E Samei & J Hsieh, eds, ‘SPIE Medical Imaging: Physics of Medical Imaging’ Vol. 7258 of *Proceedings of SPIE* SPIE Lake Buena Vista, FL, USA pp. 72580B–72580B–11.
- Hetterich H, Redel T, Lauritsch G, Rohkohl C & Rieber J 2010 New X-ray imaging modalities

- and their integration with intravascular imaging and interventions *Int. J. Cardiovas. Imag.* **26**(7), 797–808.
- Hofmann H G, Keck B, Rohkohl C & Hornegger J 2011 Comparing performance of many-core CPUs and GPUs for static and motion compensated reconstruction of C-arm CT data *Med. Phys.* **38**(1), 468–473.
- Jandt U, Schäfer D, Grass M & Rasche V 2009 Automatic generation of time resolved motion vector fields of coronary arteries and 4D surface extraction using rotational x-ray angiography *Phys. Med. Biol.* **54**(1), 45–64.
- Keil A, Vogel J, Lauritsch G & Navab N 2009 Dynamic Cone Beam Reconstruction Using a New Level Set Formulation *in* G.-Z Yang, D Hawkes, D Rueckert, A Noble & C Taylor, eds, ‘Medical Image Computing and Computer-Assisted Intervention – MICCAI 2009’ Vol. 5762 of *Lecture Notes in Computer Science* Springer London, UK pp. 389–397.
- Lebois A, Florent R & Auvray V 2011 Geometry-constrained coronary arteries motion estimation from 2D angiograms – Application to injection side recognition *in* S Wright, X Pan & M Liebling, eds, ‘8th International Symposium on Biomedical Imaging: From Nano to Macro (ISBI 2011)’ IEEE Chicago, IL, USA pp. 541–544.
- Li D, Carr J, Shea S, Zheng J, Deshpande V, Wielopolski P & Finn J 2001 Coronary Arteries: Magnetization-prepared Contrast-enhanced Threedimensional Volume-targeted Breath-hold MR Angiography *Radiology* **219**(1), 270–277.
- Müller K, Rohkohl C, Lauritsch G, Schwemmer C, Heidbüchel H, De Buck S, Nuyens D, Kyriakou Y, Köhler C & Hornegger J 2012 4-D Motion Field Estimation by Combined Multiple Heart Phase Registration (CMHPR) for Cardiac C-arm Data *in* ‘IEEE NSS/MIC Record’ Anaheim, CA, USA pp. 3707–3712.
- Perrenot B, Vaillant R, Prost R, Finet G, Douek P & Peyrin F 2007 Motion correction for coronary stent reconstruction from rotational X-ray projection sequences *IEEE Trans. Med. Imaging* **26**(10), 1412–1423.
- Prümmer M, Hornegger J, Lauritsch G, Wigström L, Girard-Hughes E & Fahrig R 2009 Cardiac C-Arm CT: A Unified Framework for Motion Estimation and Dynamic CT *IEEE Trans. Med. Imaging* **28**(11), 1836–49.
- Rohkohl C 2012 ‘CAVAREV’.
http://www.cavarev.com/evaluation_area?pub=1
- Rohkohl C, Lauritsch G, Keil A & Hornegger J 2010 CAVAREV – An Open Platform for Evaluating 3D and 4D Cardiac Vasculature Reconstruction *Phys. Med. Biol.* **55**(10), 2905–2915.
- Rohkohl C, Lauritsch G, Nöttling A, Prümmer M & Hornegger J 2008 C-Arm CT: Reconstruction of Dynamic High Contrast Objects Applied to the Coronary Sinus *in* ‘IEEE NSS/MIC Record’ Dresden, Germany pp. M10–328.
- Ruijters D, ter Haar Romeny B & Suetens P 2008 Efficient GPU-Based Texture Interpolation using Uniform B-Splines *Journal of Graphics, GPU, & Game Tools* **13**(4), 61–69.
- Russakoff D, Rohlfing T, Ho A, Kim D, Shahidi R, Adler J & Maurer C 2003 Evaluation of Intensity-Based 2D–3D Spine Image Registration Using Clinical Gold-Standard Data *Biomed. Image Registration* **2717**, 151–160.
- Schäfer D, Borgert J, Rasche V & Grass M 2006 Motion-compensated and gated cone beam filtered

- back-projection for 3-D rotational X-ray angiography *IEEE Trans. Med. Imaging* **25**(7), 898–906.
- Schomberg H 2007 Time-Resolved Cardiac Cone Beam CT *in* ‘9th Intl. Mtg. on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine’ Lindau, Germany pp. 362–365.
- Schwemmer C 2012 ‘Image Noise Estimation Software’.
<http://www5.cs.fau.de/our-team/schwemmer-chris/software/>
- Schwemmer C, Rohkohl C, Lauritsch G, Müller K & Hornegger J 2012 Residual Motion Compensation in ECG-Gated Cardiac Vasculature Reconstruction *in* F Noo, ed., ‘Proc. 2nd Intl. Mtg. on image formation in X-ray CT’ Salt Lake City, UT, USA pp. 259–262.
- Taguchi K, Sun Z, Segars W P, Fishman E K & Tsui B M W 2007 Image-domain motion compensated time resolved 4D cardiac CT *in* J Hsieh & M Flynn, eds, ‘SPIE Medical Imaging: Physics of Medical Imaging’ Vol. 6510 of *Proceedings of SPIE* SPIE San Diego, CA, USA pp. 651016–651016–9.
- Tang Q, Cammin J, Srivastava S & Taguchi K 2012 A fully four-dimensional, iterative motion estimation and compensation method for cardiac CT *Med. Phys.* **39**(7), 4291–4305.
- Unser M 1999 Splines: A Perfect Fit for Signal and Image Processing *IEEE Signal Proc. Mag.* **16**(6), 22–38.
- Weinlich A, Keck B, Scherl H, Kowarschik M & Hornegger J 2008 Comparison of High-Speed Ray Casting on GPU using CUDA and OpenGL *in* ‘High-performance and Hardware-aware Computing (HipHaC 2008)’ Como, Italy pp. 25–30.
- Wiesent K, Barth K, Navab N, Durlak P, Brunner T, Schuetz O & Seissler W 2000 Enhanced 3-D-reconstruction algorithm for C-arm systems suitable for interventional procedures *IEEE Trans. Med. Imaging* **19**(5), 391–403.
- Yang S M & Tai S C 2010 Fast and reliable image-noise estimation using a hybrid approach *J. Electron. Imaging* **19**(3), 033007–1–15.
- Zeng R, Fessler J & Balter J 2005 Respiratory motion estimation from slowly rotating x-ray projections: Theory and simulation *Med. Phys.* **32**(4), 984–991.