

Keeping the Pace: Heart Rate Informed 3-D Motion Detection for Adaptive Temporal Smoothing

Oliver Taubmann^{*†}, Günter Lauritsch[‡], Andreas Maier^{*†}, Rebecca Fahrig[§], and Joachim Hornegger^{*†}

^{*}Pattern Recognition Lab, Computer Science Department, Friedrich-Alexander-University Erlangen-Nuremberg, Germany

[†]Graduate School in Advanced Optical Technologies (SAOT), Erlangen, Germany

[‡]Siemens AG, Healthcare, Forchheim, Germany

[§]Radiological Sciences Laboratory, Stanford University, California, USA

Abstract—Time-resolved 3-D cardiac imaging in the catheter laboratory using C-arm computed tomography (CT) could provide valuable information to physicians during interventions. However, for clinically reasonable acquisition protocols, electrocardiography (ECG) gated reconstructions of individual heart phases are severely degraded due to sparse view sampling artifacts. As they appear at different locations for each heart phase, these artifacts strongly influence motion estimation, thus impairing a subsequent motion-compensated reconstruction which has been shown to greatly improve image quality. We present a method for reducing such artifacts by an adaptive smoothing in the temporal domain, guided by a projection-based motion detection step which incorporates the heart rate as prior knowledge. In our experiments on clinical data and a numerical heart phantom, we show that the increased temporal consistency achieved for initial images propagates into the motion-compensated reconstructions, improving their quality.

I. INTRODUCTION

During interventions in the catheter laboratory, functional evaluation of the heart motion today is mostly limited to 2-D ventriculography. Modalities commonly used for wall motion analysis such as cardiac magnetic resonance imaging or 3-D echocardiography are inconvenient to apply intra-procedurally. Therefore, availability of interventional 3-D reconstructions of multiple heart phases (i. e., 4-D imaging) would be highly beneficial. For instance, they could provide valuable information regarding the area of latest contraction to cardiologists during cardiac resynchronization therapy procedures, and support them in finding the optimal lead position [1].

Using C-arm devices as they are commonly found in catheter labs, rotational angiography can, in principle, provide such images by means of multi-segment retrospective ECG-gating [2]. However, even when employing a specifically designed contrast and acquisition protocol for dynamic cardiac imaging [3], the gated reconstructions suffer from severe noise and streaking artifacts due to sparse view sampling.

Several contributions have been made to improve quality of 4-D imaging with C-arm systems explicitly. A notable one by Mory et al. employs 4-D iterative reconstruction incorporating both spatial and temporal regularization [4]. Another very promising approach combines reconstruction techniques based on filtered back-projection with motion-compensation [5]. It treats the ECG-gated reconstructions as initial images from which the motion between all phases is estimated, and subsequently performs a final reconstruction using all acquired

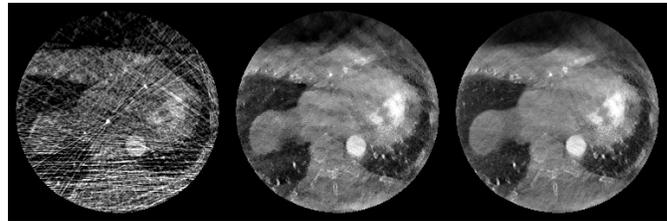


Fig. 1. Initial images for patient data set 1 reconstructed using ECG-gated filtered back-projection, with (middle) and without (left) the artifact reduction and denoising steps described in section II-A. On the right, the adaptive temporal smoothing described in section II-B was applied to the middle image.

projection data by taking the motion into account. However, the accuracy of motion estimation is highly sensitive to the quality of the initial images [5].

Artifacts appearing unpredictably over the whole spatio-temporal domain strongly disturb image registration, leading to false motion patterns—an effect we term temporal inconsistency. We present a method for increasing temporal consistency by an adaptive temporal smoothing which aims to reduce artifacts in the initial images while *keeping* motion corresponding to the heart *pace*. A novel, projection-based motion detection step is introduced to guide the smoothing.

II. MATERIALS AND METHODS

Below, we first introduce the workflow of the motion-compensated reconstruction technique used in this study (section II-A). We continue by describing the two main steps of our proposed method, which aims to improve image quality within this framework (section II-B). Afterwards, some details regarding our experimental data are provided (section II-C).

A. Motion-compensated Cardiac C-arm CT Reconstruction

A temporal sequence of initial images is generated using retrospective ECG-gating with a rectangular gating window and Feldkamp-Davis-Kress (FDK) filtered back-projection reconstruction. Few-view artifacts appear since only a subset of the projection images is used. Denoising and artifact reduction are performed by the following steps: **(i)** High-density objects within the reconstruction field of view (FOV) such as catheters are removed [6]. **(ii)** Artifacts caused by cables outside the FOV are reduced by a threshold-based masking in the filtered

projection images. **(iii)** To further reduce few-view artifacts, an approach by McKinnon and Bates [7] is employed in which object-dependent artifact images are estimated and afterwards subtracted from the original images. **(iv)** Spatial domain denoising is performed with a joint bilateral filter using a reconstruction from all available data as the guidance image [8], [9]. Fig. 1 exemplarily shows the combined effect of these steps on an initial image. A coarse region of interest (ROI) around the heart is sketched manually to which the subsequent motion estimation is constrained for increased efficiency.

Pairwise 3-D/3-D motion estimation is performed for all phases using a uniform cubic B-spline motion model and normalized cross-correlation as the image similarity metric [5]. Optimization is performed on a 3-level multi-resolution pyramid with a quasi-Newton method, the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm. 10,000 random image samples per iteration and level are evaluated within the ROI. Our implementation uses `elastix`, a toolbox for nonrigid registration of medical images [10]. The B-spline motion model, which has an isotropic control point spacing of 8mm, is evaluated on the image grid to obtain a dense deformation field. This is incorporated into a voxel-driven FDK reconstruction which compensates for the motion by shifting the currently considered voxel according to the deformation field during the back-projection step [11].

Our method consists in an adaptive temporal smoothing applied to the initial images prior to motion estimation in order to improve the quality of motion-compensated reconstruction.

B. Adaptive Temporal Smoothing

1) *3-D Motion Detection*: As a preparation for the subsequent smoothing, a motion detection step aims to determine how much each voxel position $\mathbf{x} = (x, y, z)^\top$ is affected by cardiac motion. In the initial images $\mathbf{I}^t(\mathbf{x})$, $t \in \{0, \dots, N_{\text{phases}}\}$, however, intensity changes over the cardiac phase t may either be due to actual heart motion or artifacts. Our key ideas to distinguish between both effects are the following ones: **(i)** The projection images show the heart motion, but no artifacts correlated with this motion. **(ii)** They also offer a very high temporal resolution, showing many individual heart beats as opposed to one “average” cycle obtained by ECG-gated reconstruction. **(iii)** The frequency of the motion, i. e., the heart rate, is known from the ECG.

Let us follow the forward-projection of \mathbf{x} over the sequence of acquired projection images \mathbf{P}_k , $k \in \{0, \dots, N_{\text{proj}}\}$ after line integral conversion. More precisely, forward-project \mathbf{x} to each image and consider the sequence of line integral values $p_k(\mathbf{x})$, $k \in \{0, \dots, N_{\text{proj}}\}$ as a temporal profile for \mathbf{x} . This profile will exhibit several effects. It will have a low-frequency component (“base-line drift”) as lateral projections at the beginning and end of the sweep, having traveled through more tissue, show larger integrals than the frontal ones. Locally, spikes will appear when objects of high density such as cables or catheters pass behind or in front of \mathbf{x} . If the tissue located at \mathbf{x} is affected by heart motion, we expect the profile to have a periodic component with a frequency equal to the heart rate for

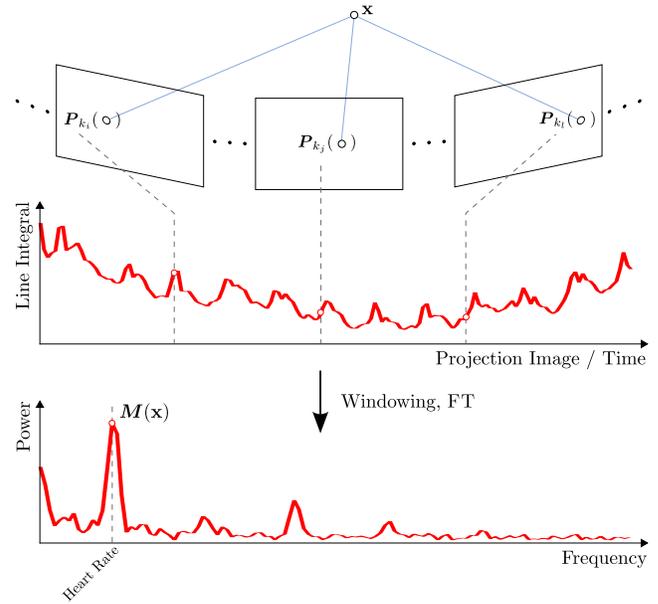


Fig. 2. For each 3-D position \mathbf{x} , the energy $M(\mathbf{x})$ of the heart rate is computed from the line integral images \mathbf{P}_k .

most of the scan. The magnitude of this component is what we are looking for. For this purpose, it seems natural to perform frequency analysis. We first apply a windowing function—in our experiments, we used the “exact Blackman” window—to the $p_k(\mathbf{x})$ to enforce periodicity, making the profile vanish at both ends. Then, we compute the Fourier transform to obtain the power spectrum, including the power (or energy) $M(\mathbf{x})$ associated with the heart rate for each voxel position \mathbf{x} . This approach is illustrated in Fig. 2. A visualization of the spatial distribution of the detected energies is shown in Fig. 3.

However, using the Fast Fourier Transform (FFT), this is rather inefficient for a large number of projections and voxels. If we were to analyze each voxel in parallel, memory complexity would be in $\mathcal{O}(N_{\text{proj}} \cdot N_{\text{voxels}})$. While it is possible to do it sequentially or in small blocks at the cost of increased runtime, there is a more elegant solution. The Discrete Fourier Transform (DFT) for a single frequency can be computed in $\mathcal{O}(n)$, where $n \equiv N_{\text{proj}}$ in our case, which is computationally less demanding than the $\mathcal{O}(n \log n)$ complexity of FFT. The Goertzel filter [12] is a simple recursive filter that performs this computation efficiently and with a constant memory footprint, independent of the sequence (profile) length.

2) *Temporal Gauss Filter*: In regions not affected by cardiac motion, we can apply strong temporal smoothing to eliminate temporally inconsistent artifacts. For this purpose, we use a simple 1-D Gauss filter in the temporal dimension, i. e. we blur the $\mathbf{I}^t(\mathbf{x})$ along t ,

$$\mathbf{I}_s^t(\mathbf{x}) = \sum_{t'=0}^{N_{\text{phases}}} \mathbf{I}^{t'}(\mathbf{x}) \cdot \frac{1}{\sigma(\mathbf{x})\sqrt{2\pi}} \exp\left(-\frac{\text{dist}^2(t, t')}{2\sigma^2(\mathbf{x})}\right), \quad (1)$$

where $\text{dist}(\cdot, \cdot)$ denotes the absolute distance of two phases in the cardiac cycle. In contrast, wherever we assume to have

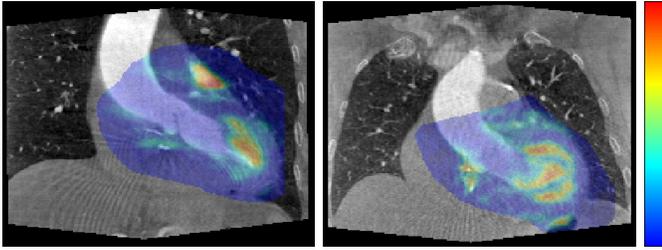


Fig. 3. Color-coded visualizations of the cardiac motion detected by our method inside the considered region of interest for patients 1 and 2. It is overlaid on a reconstruction from all data for anatomical orientation. Warmer hues correspond to larger motion.

cardiac motion, we need to take care so as not to lose actual motion in the process, and therefore narrow the smoothing kernel down by decreasing the standard deviation $\sigma(\mathbf{x})$.

For determining the appropriate $\sigma(\mathbf{x})$ for each \mathbf{x} , we make use of the heart rate energies $M(\mathbf{x})$, which are converted to a weight map $M_w(\mathbf{x})$: First, we remove outliers and denoise $M(\mathbf{x})$ by applying a 3×3 median filter as well as a blur filter with 1.5 mm standard deviation, respectively. In order to normalize the values, we perform a global linear scaling such that the mean \bar{M} of $M(\mathbf{x})$ over the considered ROI is mapped to 1. Values larger than this mean are clamped; thereby, we assign the maximum weight of 1 to all positions \mathbf{x} for which $M(\mathbf{x}) \geq \bar{M}$. For all other positions, the weights vary in the range $[0; 1]$ and can be used to interpolate between a σ_{\min} (in case of strong motion) and a σ_{\max} (no motion),

$$\sigma(\mathbf{x}) = \sigma_{\min} \cdot M_w(\mathbf{x}) + \sigma_{\max} \cdot (1 - M_w(\mathbf{x})). \quad (2)$$

In our experiments, we set σ_{\min} and σ_{\max} to 1% and 50% of the cardiac cycle, respectively. Additionally, for improved convergence of motion estimation, the ROI was narrowed down by removing voxels for which $M(\mathbf{x}) \leq 0.5 \bar{M}$.

C. Experiments

1) *Clinical Data*: Two clinical patient data sets were acquired using an Artis zeego system (Siemens AG, Healthcare, Forchheim, Germany). One C-arm rotation of 14s duration was performed, capturing 381 projection images at approx. 30 Hz with an angular increment of 0.52° . The isotropic pixel resolution was 0.31 mm/pixel (0.21 in isocenter), the detector size 1240×960 pixels. The heart was stimulated through external pacing to 115 bpm. The gating window width was chosen as 10% of the heart cycle. 10 cardiac phases were reconstructed at a phase increment of 10% for both the initial and motion-compensated images. 91 ml of undiluted contrast agent were administered in the pulmonary artery at a speed of 7 ml/s. The appropriate X-ray delay was determined by a prior test bolus injection. Image reconstruction was performed on a volume of size $(25.6 \text{ cm})^3$, distributed on a 256^3 voxel grid.

2) *Phantom Model*: A ventricle data set similar to the XCAT phantom was created [13], [14]. Projections of this phantom were simulated using a polychromatic X-ray spectrum with 56 energy bins from 10 keV to 150 keV and a

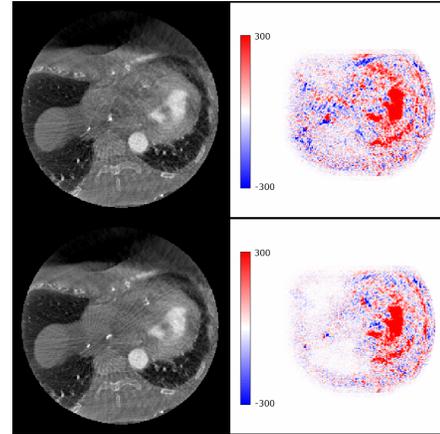


Fig. 4. Motion-compensated reconstructions of the end-diastolic phase for patient data set 1 and color-coded difference images between the end-diastolic and end-systolic phase. Motion has been estimated from initial reconstructions processed with (bottom) and without (top) temporal smoothing.

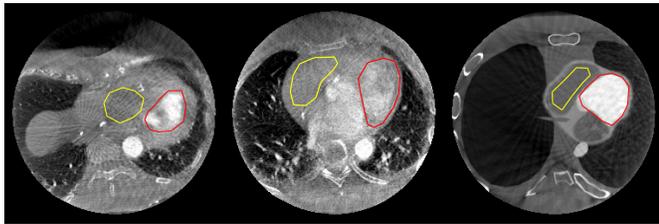
time-current product of 2.5 mAs per X-ray pulse. The material properties of bones and bone marrow were chosen according to the mass attenuation coefficients of the NIST X-ray table¹. All other structures were assumed to have the same absorption behavior as water with modified densities. The densities of the contrasted left ventricle blood pool, the myocardial wall and the contrasted blood in the aorta were set to 2.5 g/cm^3 , 1.5 g/cm^3 and 2.0 g/cm^3 , respectively. To obtain a gold standard reconstruction, projection images for a single cardiac phase (static phantom) were generated. Other properties of the data set, such as the heart rate, number of projection images, resolution and dimensions, were chosen equivalently to the clinical data sets described above.

III. RESULTS AND DISCUSSION

Fig. 4 shows that the intensity differences between end-systolic and end-diastolic phases in uncontrasted cardiac tissue and blood presumed—or in case of the phantom, known—to appear static decrease considerably while heart motion is preserved. For the quantitatively evaluated static regions depicted in Fig. 5, the average standard deviation w.r.t. the heart phases was reduced by more than 50%, whereas it was only reduced slightly in the contrasted left ventricular blood pool. In our phantom study, the root mean square error (RMSE) over the ROI, computed between the motion-compensated and the static ground truth reconstruction, is lowered from 771.4 to 689.1, a decrease of about 9%. As expected, this improvement is not as dramatic as that observed for the temporal variance, but it nonetheless indicates that enforcing consistency in time also reduces artifacts in the spatial domain. A visualization of the differences is found in Fig. 6. Errors at the heart walls can partially be attributed to residual motion within the gating window.

One might argue that the motion detection step could be replaced by a coarse segmentation of the heart. However,

¹<http://physics.nist.gov/PhysRefData/Xcom/html/xcom1.html>



Data set	TS	Uncontrasted	Contrasted
Patient 1	-	634 ± 70 [221]	964 ± 154 [446]
	✓	635 ± 32 [102]	959 ± 148 [434]
Patient 2	-	776 ± 49 [153]	870 ± 83 [252]
	✓	776 ± 23 [74]	866 ± 78 [236]
Phantom	-	593 ± 33 [100]	1232 ± 110 [298]
	✓	594 ± 12 [38]	1230 ± 107 [292]

Fig. 5. In all data sets (from left to right: patients 1 and 2, phantom), a region containing unconstrained cardiac tissue and blood (yellow), which should appear mostly static, as well as a region tightly surrounding the contrasted left ventricular blood pool in end-diastole (red), where the most prominent motion should be expected, were selected for quantitative evaluation. The table shows temporal statistics (i. e., with respect to the heart phases) of intensity values in motion-compensated reconstructions, averaged over all voxels in the respective region, in the format mean ± std [max – min]. The column TS indicates whether the proposed temporal smoothing was applied to the initial images used for motion estimation.

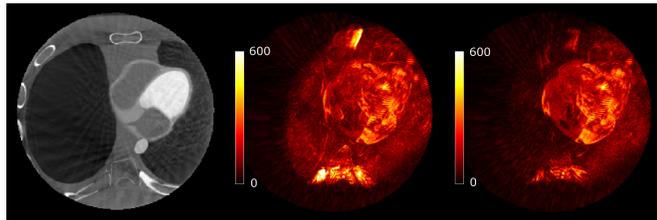


Fig. 6. Left: A motion-compensated reconstruction of our phantom model using the proposed method. For the two color-coded error images, computed as the maximum intensity projection along the z -axis of the absolute differences to the ground truth reconstruction, motion has been estimated from initial images processed with (right) and without (middle) temporal smoothing.

unlike most organ segmentation algorithms, it neither requires nor imposes any prior knowledge on the shape of the considered objects apart from the (known) frequency. It can identify regions with varying degrees of motion directly from the projection images without the need for a reconstruction of multiple phases or sophisticated registration methods. At the same time, it is computationally undemanding due to the use of the Goertzel filter. In our clinical data sets, projections acquired from lateral directions suffer from low contrast. Naturally, this affects both reconstruction as well as our motion detection method. However, we have observed that it mainly causes the motion mask to be spread out in anterior-posterior direction, i. e. motion is potentially overestimated. For the subsequent adaptive smoothing step, this is inconvenient, yet uncritical; it merely causes a less aggressive removal of artifacts.

IV. CONCLUSION

For the reduction of temporal inconsistency in motion-compensated image reconstruction, a voxel-based temporal smoothing of the initial image sequence used for motion estimation is proposed. The degree of smoothing depends inversely on the magnitude of cardiac motion detected by a novel scheme based on frequency analysis of projection data. For future work, other potential uses for the detected motion, e. g. as prior knowledge for regularization, as well as extending the method to irregular cardiac motion can be considered.

Acknowledgments and Disclaimer: 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, and thank Dr. Bernd Abt (Centre of Cardiovascular Diseases in Rotenburg a. d. Fulda) for providing the image data. The concepts and information presented in this paper are based on research and are not commercially available.

REFERENCES

- [1] M. Döring, F. Braunschweig, C. Eitel, T. Gaspar, U. Wetzel, B. Nitsche, G. Hindricks, and C. Piorkowski, "Individually tailored left ventricular lead placement: lessons from multimodality integration between three-dimensional echocardiography and coronary sinus angiogram," *Europace*, vol. 15, no. 5, pp. 718–727, May 2013.
- [2] G. Lauritsch, J. Boese, L. Wigstrom, H. Kemeth, and R. Fahrig, "Towards cardiac C-arm computed tomography," *IEEE Trans. Med. Imag.*, vol. 25, no. 7, pp. 922–934, July 2006.
- [3] S. De Buck, D. Dauwe, J.-Y. Wielandts, P. Claus, S. Janssens, H. Heidebuchel, and D. Nuyens, "A new approach for prospectively gated cardiac rotational angiography," *Proc. SPIE*, vol. 8668, 2013.
- [4] C. Mory, V. Auvray, B. Zhang, M. Grass, D. Schäfer, S. Chen, J. Carroll, S. Rit, F. Peyrin, P. Douek, and L. Boussel, "Cardiac C-arm computed tomography using a 3D + time ROI reconstruction method with spatial and temporal regularization," *Med. Phys.*, vol. 41, p. 021903, 2014.
- [5] K. Müller, A. Maier, C. Schwemmer, G. Lauritsch, S. D. Buck, J.-Y. Wielandts, J. Hornegger, and R. Fahrig, "Image artefact propagation in motion estimation and reconstruction in interventional cardiac C-arm CT," *Phys. Med. Biol.*, vol. 59, no. 12, pp. 3121–3138, 2014.
- [6] K. Müller, G. Lauritsch, C. Schwemmer, A. Maier, O. Taubmann, B. Abt, H. Köhler, A. Nöttling, J. Hornegger, and R. Fahrig, "Catheter artifact reduction (CAR) in dynamic cardiac chamber imaging with interventional C-arm CT," in *Proc. of 3rd Intl. Conf. on Image Formation in X-ray CT*, F. Noo, Ed., 2014, pp. 418–421.
- [7] G. C. Mc Kinnon and R. Bates, "Towards imaging the beating heart usefully with a conventional CT scanner," *IEEE Trans. Biomed. Eng.*, vol. BME-28, no. 2, pp. 123–127, Feb 1981.
- [8] M. Manhart, M. Kowarschik, A. Fieselmann, Y. Deuerling-Zheng, and J. Hornegger, "Fast Dynamic Reconstruction Algorithm with Joint Bilateral Filtering for Perfusion C-arm CT," in *Proceedings NSS/MIC*, 2012, pp. 2304–2311.
- [9] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *6th Intl. Conf. on Computer Vision*, Jan 1998, pp. 839–846.
- [10] S. Klein, M. Staring, K. Murphy, M. Viergever, and J. Pluim, "elastix: a toolbox for intensity based medical image registration," *IEEE Trans. Med. Imag.*, vol. 29, no. 1, pp. 196–205, 2010.
- [11] D. Schäfer, J. Borgert, V. Rasche, and M. Grass, "Motion-compensated and gated cone beam filtered back-projection for 3-D rotational X-ray angiography," *IEEE Trans. Med. Imag.*, vol. 25, pp. 898–906, 2006.
- [12] G. Goertzel, "An algorithm for the evaluation of finite trigonometric series," *American Mathematical Monthly*, 1958.
- [13] A. Maier, H. Hofmann, M. Berger, P. Fischer, C. Schwemmer, H. Wu, K. Müller, J. Hornegger, J.-H. Choi, C. Riess, A. Keil, and R. Fahrig, "CONRAD - A software framework for cone-beam imaging in radiology," *Med. Phys.*, vol. 40, no. 11, 2013.
- [14] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, "4D XCAT phantom for multimodality imaging research," *Med. Phys.*, vol. 37, pp. 4902–4915, 2010.