

Coping with Real World Data: Artifact Reduction and Denoising for Motion-Compensated Cardiac C-arm CT

Oliver Taubmann,^{a)} Andreas Maier, and Joachim Hornegger

Pattern Recognition Lab, Computer Science Department, Friedrich-Alexander-University Erlangen-Nuremberg and Erlangen Graduate School in Advanced Optical Technologies (SAOT), Germany

Günter Lauritsch

Siemens Healthcare GmbH, Forchheim, Germany

Rebecca Fahrig

Radiological Sciences Laboratory, Stanford University, California, USA and Siemens Healthcare GmbH, Forchheim, Germany

(Dated: 28 January 2016)

Purpose: Detailed analysis of cardiac motion would be helpful for supporting clinical workflow in the interventional suite. With an angiographic C-arm system, multiple heart phases can be reconstructed using electrocardiogram gating. However, the resulting angular undersampling is highly detrimental to the quality of the reconstructed images, especially in non-ideal intraprocedural imaging conditions. Motion-compensated reconstruction has previously been shown to alleviate this problem, but it heavily relies on a preliminary reconstruction suitable for motion estimation. In this work, we propose a processing pipeline tailored to augment these initial images for the purpose of motion estimation, and assess how it affects the final images after motion compensation.

Methods: The following combination of simple, direct methods inspired by the core ideas of existing approaches proved beneficial: **(a)** Streak reduction by masking high-intensity components in projection domain after filtering. **(b)** Streak reduction by subtraction of estimated artifact volumes in reconstruction domain. **(c)** Denoising in spatial domain using a joint bilateral filter guided by an uncompensated reconstruction. **(d)** Denoising in temporal domain using an adaptive Gaussian smoothing based on a novel motion detection scheme.

Results: Experiments on a numerical heart phantom yield a reduction of the relative root-mean-square error from 89.9% to 3.6% and an increase of correlation with the ground truth from 95.763% to 99.995% for the motion-compensated reconstruction when our processing is applied to the initial images. In three clinical patient data sets, the signal-to-noise ratio measured in an ideally homogeneous region is increased by 37.7% on average. Overall visual appearance is improved notably and some anatomical features are more readily discernible.

Conclusions: Our findings suggest that the proposed sequence of steps provides a clear advantage over an arbitrary sequence of individual image enhancement methods and is fit to overcome the issue of lacking image quality in motion-compensated C-arm imaging of the heart. As for future work, the obtained results pave the way for investigating how accurately cardiac functional motion parameters can be determined with this modality.

I. INTRODUCTION

Comprehensive functional analysis of heart motion during catheter-based interventions could prove a valuable asset. For instance, in case of asynchronous motion, it would allow cardiologists to identify a ventricle's region of latest contraction¹. Presently, in clinical practice, functional analysis in the cath lab is usually based on echocardiography, which is inconvenient to perform interventionally, or 2-D X-ray ventriculography using an angiographic C-

arm device, which inherently offers only a limited view of the organ motion.

To overcome this limitation, such a device can also be used to perform a volumetric reconstruction of multiple heart phases by means of rotational angiography and cone-beam computed tomography (CBCT) techniques. As the C-arm rotation is not fast enough to achieve the temporal resolution required for imaging the beating heart directly, multi-segment electrocardiogram (ECG) gating is intro-

duced to select or weight the projection data for each cardiac phase to be reconstructed. This leads to a strong undersampling that constitutes the main challenge for the practical application of this technique.

Early approaches used multi-sweep protocols with as many as four rotations of the C-arm to acquire sufficiently sampled data^{2,3}. However, the long acquisition time results in an increased radiation dose and the need for an extended breath-hold period in addition to a longer delay of the actual procedure. A more recent single-sweep protocol alleviates these problems by shortening the acquisition duration to about 14s, using electrophysiologic pacing to ensure a sufficiently high heart rate so that enough distinct views are observed for each cardiac phase⁴. It has been successfully applied in porcine models^{5,6} as well as in humans⁷. In the work of Müller et al.⁵, a motion-compensated reconstruction technique was employed in order to achieve good image quality.

In this paper, we propose to enhance this state-of-the-art motion compensation framework with several additional processing steps applied to the initial, ECG-gated images in order to cope with the challenges posed by clinical patient data sets. These challenges consist mainly in a dramatically increased amount of artifacts caused by high-density objects, aggravated by the strong undersampling, as well as overall worse contrast due to the larger amount of tissue to traverse, especially in lateral direction due to acquisition of the patients with both arms resting beside the torso.

Countless methods have been employed in the past for the purpose of artifact reduction and denoising in (dynamic) C-arm CT reconstruction. Prevalent approaches to avoid streaking rely on the identification and removal of high-density objects^{8,9} and the use of iterative reconstruction methods with various types of regularization¹⁰⁻¹². Denoising is commonly performed with non-linear low-pass filters with edge-preserving capabilities¹³⁻¹⁵. The rationale behind some existing methods and how they relate to our processing steps is detailed in the corresponding parts below (sections II.B and II.C). We have found the proposed combination and order of steps to be favorable as the parts complement each other.

II. METHODS

This work is concerned with improving the quality of initial gated reconstructions needed for motion estimation within the motion-compensated reconstruction framework summarized in section II.A. This endeavor is worthwhile since it has been established that artifacts in the initial images propagate through the steps of motion estimation and compensation, thus strongly influencing the final images⁵. The focus lies on straight-forward direct methods addressing the major problems faced in rotational an-

diagrams strongly undersampled due to the use of ECG gating. Successful concepts are identified in established approaches and reduced to their essential ideas to design effective components which hardly require prior knowledge about the object or acquisition and do not rely on iterative optimization schemes. The four automated steps cover streak reduction in the projection (section II.B.1) and the reconstruction domains (II.B.2), as well as edge-preserving spatial (II.C.1) and adaptive temporal smoothing (II.C.2).

II.A. Motion Compensation Framework

Initial Image Reconstruction. A temporal sequence of initial images is generated using the Feldkamp-Davis-Kress (FDK) method¹⁶ combined with retrospective ECG gating¹⁷ using a rectangular window, i. e. for each heart phase to be reconstructed, a subset of the available projections is used depending on their position in the cardiac cycle as indicated by the ECG.

Motion Estimation. The motion between all pairs of reconstructed cardiac phases is then estimated by 3-D/3-D image registration using a uniform cubic B-spline motion model with an isotropic control point spacing of 8mm and a normalized cross-correlation similarity metric⁵. Optimization is restricted to a region of interest (ROI) Ω around the heart and performed on a multi-resolution pyramid with three levels. In each iteration, 10,000 random image samples are evaluated within Ω to compute a quasi-Newton update step using the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm¹⁸. Our implementation uses `elastix`, a toolbox for nonrigid registration of medical images¹⁹.

Motion-compensated Reconstruction. The resulting B-spline motion model is evaluated on the image grid to obtain a dense deformation field, which is incorporated into a voxel-driven FDK reconstruction that compensates for motion by shifting the currently considered voxel according to the deformation field during the back-projection step²⁰.

II.B. Reduction of Streak Artifacts

Streak artifacts in the initial images typically appear due to the angular undersampling. The most prominent streaks are caused by high-density objects not only inside, but also outside of the reconstruction field-of-view (FOV) as the imaged object is highly truncated. Additionally, many other effects can cause raw data inconsistency that leads to—usually slightly more subtle, blurred and less pronounced—streaks, such as motion, noise, photon starvation, beam hardening, overexposure, truncation and scatter. In our pipeline, a projection-domain method to alleviate the strong influence of high-density objects is followed by

a reconstruction-domain approach to further reduce streaking.

II.B.1. Projection-Domain Artifact Reduction

In many cases, artifacts caused by high-density objects are dealt with by segmenting the objects and removing them from the projection images, e.g. by replacing them with values interpolated from surrounding regions. If segmentation of these objects directly in the projection images turns out not to be feasible, a preliminary reconstruction may be performed in order to segment the objects in the target domain and then project the segmentation forward again⁸. However, this naturally does not work for objects outside of the FOV. In an interventional setting, cables, electrodes and other large objects outside the body such as scissors may be visible in some projections only. As shown by Haase et al., properties arising from the fact that these objects are far away from the isocenter can help identify them more reliably in the projections⁹.

Rohkohl et al. proposed a different approach to streak reduction²¹ that does not rely on segmentation or prior knowledge, but follows a rather pragmatic notion, starting from the undesired effect observed in the reconstructed images: During the back-projection step of FDK reconstruction, the extreme, i.e. the highest and lowest, contributions to each voxel are those that ultimately constitute streaking in the image when they are not canceled out by their corresponding counterparts from missing nearby angulations. Therefore, they simply ignore these contributions altogether, i.e. leave out a fixed number of values to be back-projected, determined on a per-voxel basis. This method, which has previously been used in coronary tree reconstruction²², has the downside that the implementation of the back-projection algorithm itself has to be modified and additional memory and run time burdens are introduced that increase with the desired level of reduction.

We employ a purely projection-domain based method similar to this approach as a preprocessing step immediately *before* back-projection. Let $\mathbf{P}_F(i, \mathbf{u})$ denote the ramp-filtered, redundancy- and cosine-weighted projection images, where $i \in \{1, \dots, N\}$ is the projection number and $\mathbf{u} \in \Omega_{\mathbf{u}} \subset \mathbb{R}^2$ a coordinate with $\Omega_{\mathbf{u}}$ the set of all pixel positions on the detector and $|\Omega_{\mathbf{u}}|$ its cardinality, i.e. the number of detector pixels. We apply a two-sided thresholding and set all outside values to zero,

$$\mathbf{P}_F^*(i, \mathbf{u}) = \begin{cases} 0 & \text{if } |\mathbf{P}_F(i, \mathbf{u}) - \mu^*| > 1.75 \sigma^*, \\ \mathbf{P}_F(i, \mathbf{u}) & \text{else,} \end{cases} \quad (1)$$

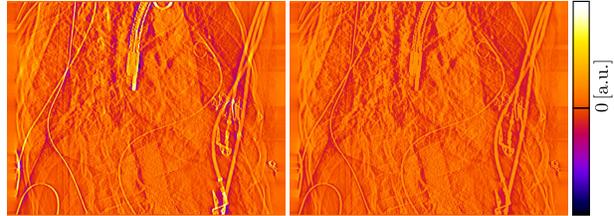


FIG. 1 Color-coded visualization of a ramp-filtered projection image before (left) and after (right) the threshold-based masking described in section II.B.1. Extreme contributions of several high-density objects have been removed, leading to a distinctly streak-reduced reconstruction as seen in Fig. 5a.

$$\mu^* = \frac{\sum_{i, \mathbf{u}} \mathbf{P}_F(i, \mathbf{u})}{N \cdot |\Omega_{\mathbf{u}}|},$$

$$\sigma^* = \frac{\sum_{i, \mathbf{u}} \left(\mathbf{P}_F(i, \mathbf{u}) - \frac{1}{|\Omega_{\mathbf{u}}|} \sum_{\mathbf{u}'} \mathbf{P}_F(i, \mathbf{u}') \right)^2}{N \cdot (|\Omega_{\mathbf{u}}| - 1)}.$$

This means that we also ignore extreme contributions. At the same time, it can be considered a simple, threshold-based segmentation of the critical object edges, which are the cause of the problem, in the relevant high-pass domain that further amplifies them, rendering their separation from the background easier (cf. Fig. 1). Note that μ^* and σ^* are estimated from all data to avoid inconsistency of thresholds between individual projections. The processed images $\mathbf{P}_F^*(i, \mathbf{u})$ can simply be reused should an additional reconstruction be required in the remainder of the pipeline.

II.B.2. Reconstruction-Domain Artifact Reduction

As mentioned above, any remaining streaks may be caused by various effects, with many dedicated methods available and in place to reduce them. However, the streaks have in common that they are vastly enhanced by undersampling, whereas the reconstruction from fully sampled data (“uncompensated reconstruction”) exhibits satisfactory image quality apart from the motion blur.

Therefore, they can be found and eliminated in reconstruction domain with a method based on an algorithm by McKinnon and Bates²³ that makes use of the information provided by the uncompensated reconstruction. As a slight variation to the original formulation, we reconstruct volumes containing only the artifacts caused by undersampling and subtract these from the original reconstructions. The procedure is illustrated in Fig. 2. In our notation, ungated FDK reconstruction reads,

$$\hat{\mathbf{I}}(\mathbf{x}) = \sum_i w(i, \mathbf{x}) \cdot \mathbf{P}_F^*(i, \mathbf{A}_i(\mathbf{x})), \quad (2)$$

with $w(i, \mathbf{x})$ the distance weight of $\mathbf{x} \in \mathbb{R}^3$ for projection image i and $\mathbf{A}_i(\mathbf{x})$ the projection of \mathbf{x} onto the detector plane. Similarly, the gated reconstructions are denoted as,

$$\mathbf{I}^t(\mathbf{x}) = \sum_i \frac{\lambda(i, t)}{\nu(t)} \cdot w(i, \mathbf{x}) \cdot \mathbf{P}_F^*(i, \mathbf{A}_i(\mathbf{x})), \quad (3)$$

where $\lambda(i, t)$ is the gating weight for projection i and cardiac phase t , and $\nu(t) = \sum_i \lambda(i, t)/N$, ensuring a common intensity scale. We forward-project $\hat{\mathbf{I}}(\mathbf{x})$ to all acquired views to obtain a set of consistent projection images $\mathbf{P}_{\text{cons}}(i, \mathbf{u})$ and perform weighting, ramp-filtering and thresholding as described above to get $\mathbf{P}_{F, \text{cons}}^*(i, \mathbf{u})$. We then perform another gated reconstruction for these images,

$$\mathbf{I}_{\text{cons}}^t(\mathbf{x}) = \sum_i \frac{\lambda(i, t)}{\nu(t)} \cdot w(i, \mathbf{x}) \cdot \mathbf{P}_{F, \text{cons}}^*(i, \mathbf{A}_i(\mathbf{x})). \quad (4)$$

The difference between $\mathbf{I}_{\text{cons}}^t(\mathbf{x})$ and $\hat{\mathbf{I}}(\mathbf{x})$ consists of undersampling artifact patterns that we can subtract from the original gated reconstructions $\mathbf{I}^t(\mathbf{x})$ to obtain the corrected images,

$$\mathbf{I}_{\text{corr}}^t(\mathbf{x}) = \mathbf{I}^t(\mathbf{x}) - (\mathbf{I}_{\text{cons}}^t(\mathbf{x}) - \hat{\mathbf{I}}(\mathbf{x})). \quad (5)$$

In the original paper, gated reconstruction is performed *after* finding the difference between the original and consistent (projection) data, which would be equivalent given an idealized linear reconstruction, but has led to concerns regarding susceptibility to noise²⁴. In our experience, computing the artifacts in volume domain after reconstruction behaves robustly. A similar idea has also been applied previously for finding 3-D artifact motion fields for better robustness of motion estimation²⁵.

II.C. Suppression of Noise and Temporal Inconsistency

After streak removal, noise and smaller artificial structures can be eliminated by smoothing both spatially and temporally. We limit ourselves to the reconstruction domain since noise in the projections is typically handled by an appropriately chosen apodization window for the ramp-filter. Also, non-linear filtering in projection images may introduce data inconsistency and was previously shown to yield no major benefits compared to similar volume-domain filtering²⁶.

II.C.1. Edge-Preserving Spatial Denoising

A popular choice for denoising is edge-preserving smoothing via bilateral filtering²⁷. However, the standard bilateral filter struggles in case of a low signal-to-noise ratio (SNR) as the intensity-based (range) component can hardly discriminate between structure and noise. Again, the key idea is to make

use of the fact that we have the fully sampled reconstruction $\hat{\mathbf{I}}(\mathbf{x})$ which has a much better SNR. Using an image to guide filtering of another is the eponymous idea of the guided filter proposed by He et al.¹⁵, but the core idea has earlier been introduced as joint bilateral filtering²⁸ and found applications in medical imaging, for instance in perfusion imaging¹³. Our bilateral filter kernel $k_{\mathbf{x}}(\mathbf{x}')$ has two Gaussian components based on proximity in space and intensity distance in the guidance image $\hat{\mathbf{I}}(\mathbf{x})$,

$$\mathbf{I}_S^t(\mathbf{x}) = \frac{\sum_{\mathbf{x}' \in N(\mathbf{x})} \mathbf{I}_{\text{corr}}^t(\mathbf{x}') \cdot k_{\mathbf{x}}(\mathbf{x}')}{\sum_{\mathbf{x}' \in N(\mathbf{x})} k_{\mathbf{x}}(\mathbf{x}')}, \quad (6)$$

$$k_{\mathbf{x}}(\mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2\sigma_S^2} - \frac{(\hat{\mathbf{I}}(\mathbf{x}) - \hat{\mathbf{I}}(\mathbf{x}'))^2}{2\sigma_I^2}\right),$$

with σ_S and σ_I the standard deviations in the spatial and intensity domain, respectively. $N(\mathbf{x})$ is the local neighborhood around \mathbf{x} with $|N(\mathbf{x})| = 7^3$ voxels. In our experiments, even though it mostly circumvents the weighting of spatial proximity, we choose $\sigma_S \gg \sqrt[3]{|N(\mathbf{x})|}$ to increase the smoothing effect without having to choose a larger $N(\mathbf{x})$, and use a narrow intensity window ($\sigma_I = 3\text{HU}$) to preserve the distinction between contrasted blood and surrounding tissue even in the presence of motion blur. Note the conceptual similarity to applications of the widely known prior image constrained compressed sensing (PICCS) method in which the uncompensated image guides iterative reconstruction^{10,29}, as well as the fact that it is commonly combined with total variation regularization³⁰ which in turn has been shown to yield results comparable to bilateral filtering in some cases¹⁴.

II.C.2. Adaptive Temporal Smoothing

Although each individual phase now has a smooth appearance, remaining artifact patterns are usually inconsistent in the temporal domain, i. e. they appear to be moving due to their occurrence at different positions. The joint bilateral smoothing introduced in the previous section, however, cannot be extended to the temporal domain trivially as by definition, the guidance image $\hat{\mathbf{I}}(\mathbf{x})$ has no temporal resolution. For this reason, other information is needed to determine where smoothing can be performed safely and where we risk blurring structures in motion, in order to adapt the smoothing filter accordingly.

Previously, we presented a novel 3-D motion detection method based on frequency analysis of the acquired projection images³¹. It can be used to estimate the amount of periodic motion corresponding to the heart rate for each voxel, resulting in a volume image $\mathbf{M}(\mathbf{x})$ that holds the relative spatial distribution of expected motion magnitudes (“motion map”). Fig. 3 gives a schematic overview of the procedure.

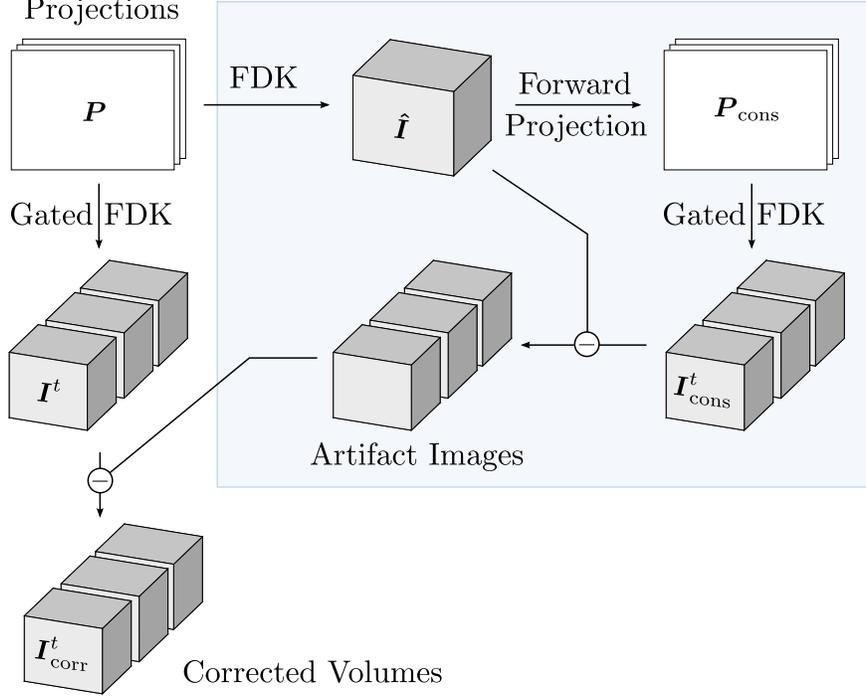


FIG. 2 Reconstruction-domain few-view artifact reduction for cardiac imaging in analogy to the approach by McKinnon and Bates²³.

We follow the forward-projection of \mathbf{x} over the sequence of acquired projection images, $\mathbf{P}(i, \mathbf{A}_i(\mathbf{x}))$, and consider the sequence of line integral values as a temporal profile. We are interested in a Fourier analysis of this profile w.r.t. a single frequency component, the heart rate f_H . The spectral power for f_H is found using the Goertzel algorithm³², which employs a two-stage recursive filter to efficiently compute the discrete Fourier transform (DFT) for one frequency,

$$\mathbf{S}_i(\mathbf{x}) = \mathbf{S}_{i-1}(\mathbf{x}) \cdot 2 \cos(2\pi f_H) - \mathbf{S}_{i-2}(\mathbf{x}) + \mathbf{W}(i) \cdot \mathbf{P}(i, \mathbf{A}_i(\mathbf{x})), \quad (7)$$

$$\mathbf{M}(\mathbf{x}) = |\mathbf{S}_N(\mathbf{x}) - e^{-2\pi i f_H} \mathbf{S}_{N-1}(\mathbf{x})|, \quad (8)$$

where $\mathbf{W}(i)$ is the “exact Blackman” apodization window³³ and the filter stages are initialized as $\mathbf{S}_k(\mathbf{x}) = 0 \quad \forall k \leq 0$. The value range of the projection images $\mathbf{P}(i, \mathbf{A}_i(\mathbf{x}))$ has been clamped to reduce the influence of high-density objects, similar to the thresholding described in section II.B.1. $\mathbf{M}(\mathbf{x})$ is subsequently cropped to the ROI Ω , i.e. outside regions are set to zero, effectively assuming no motion, and is rid of outliers and noise by a 3×3 median filter and a blur filter with 1.5 mm standard deviation, respectively. Finally, its values are normalized to weights,

$$\mathbf{M}_w(\mathbf{x}) = \begin{cases} \frac{\mathbf{M}(\mathbf{x})}{\overline{\mathbf{M}}} & \text{if } \mathbf{M}(\mathbf{x}) \leq \overline{\mathbf{M}}, \\ 1 & \text{else,} \end{cases} \quad (9)$$

where $\overline{\mathbf{M}} = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} \mathbf{M}(\mathbf{x})$, thereby linearly mapping $[0; \overline{\mathbf{M}}]$ to the range $[0; 1]$ clamping higher values

to 1. $\mathbf{M}_w(\mathbf{x})$ is in turn used to adapt the width $\sigma(\mathbf{x})$ of a Gaussian kernel for temporal smoothing,

$$\mathbf{I}_T^t(\mathbf{x}) = \sum_{t'} \frac{\mathbf{I}_S^{t'}(\mathbf{x})}{\sigma(\mathbf{x})\sqrt{2\pi}} \exp\left(-\frac{\text{dist}^2(t, t')}{2\sigma^2(\mathbf{x})}\right), \quad (10)$$

where $\text{dist}(t, t')$ denotes the distance of cardiac phases t and t' , taking periodicity into account, and $\sigma(\mathbf{x})$ is found by linear interpolation between $\sigma_{\text{max}} = 50\%$ and $\sigma_{\text{min}} = 1\%$ of the cardiac cycle,

$$\sigma(\mathbf{x}) = \sigma_{\text{min}} \cdot \mathbf{M}_w(\mathbf{x}) + \sigma_{\text{max}} \cdot (1 - \mathbf{M}_w(\mathbf{x})). \quad (11)$$

Where little motion occurs, artifacts are further reduced by stronger smoothing, while motion-intensive areas are preserved due to a very narrow smoothing kernel. To achieve a stronger effect, the filtering step can be repeated; in our experiments, we used three iterations.

For the purpose of registration, Ω is narrowed down by removing voxels for which $\mathbf{M}(\mathbf{x}) \leq 0.5 \overline{\mathbf{M}}$. This improves convergence speed as the random image samples drawn in each iteration will be more concentrated in regions with motion.

In recent literature, temporal smoothing is often used as a regularization term during algebraic reconstruction¹². Similarly, we have previously proposed iterating motion estimation, motion-compensated (analytical) reconstruction and spatio-temporal filtering to improve convergence³⁴. Such iterative methods, while fairly successful, are several orders of magnitude slower than the adaptive smoothing described above.

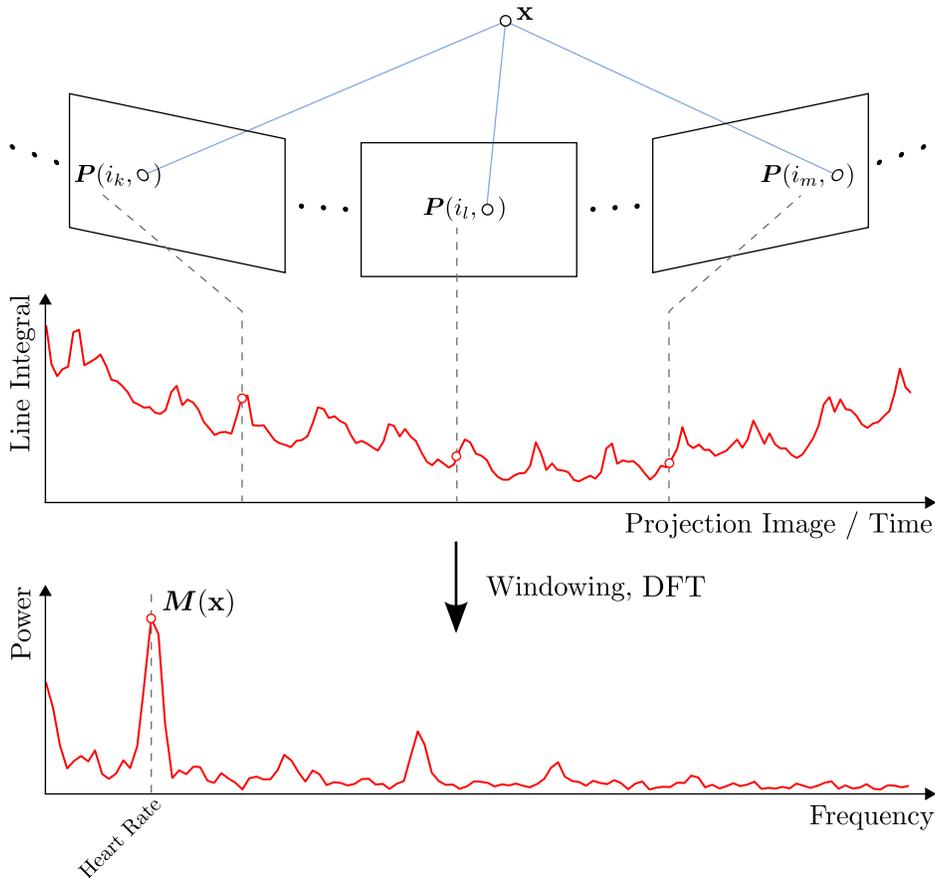


FIG. 3 For each 3-D position \mathbf{x} , the energy $M(\mathbf{x})$ of the heart rate is computed from the line integral images.

III. EXPERIMENTS

III.A. Data

III.A.1. Numerical Phantom Model

We utilize a 4-D numerical phantom model based on the XCAT phantom^{35,36} for quantitative evaluation. It is generated from segmentations of human acquisitions and highly flexible. For instance, physiological parameters such as the heart rate can readily be adapted. We generated 381 projection images using a discretized polychromatic spectrum with energy bins of width 2.5 keV from 25 keV to 120 keV (peak energy) and a time-current product of 2.5 mAs per X-ray pulse. Mass attenuation coefficients of bones and bone marrow were taken from the NIST X-ray table³⁷. Material properties of other structures were modeled with the absorption behavior of water at modified densities. The contrasted left ventricular blood pool, the contrasted blood in the aorta and the myocardium had densities of 2.5 g/cm³, 2.0 g/cm³ and 1.5 g/cm³, respectively. For reconstruction of a ground truth image, denoted as $\tilde{\mathbf{I}}(\mathbf{x})$, a complete set of projection images for a single phase at 30% of the cardiac cycle w.r.t. end-diastole was simulated

(static phantom). Other properties related to the acquisition, such as size and resolution of the images and the number of heart beats observed during the scan, were chosen identical to the clinical data sets detailed below. An archive of all data required to reproduce our phantom study for other reconstruction algorithms is available on our website³⁸.

III.A.2. Clinical Patients

Three data sets of clinical patients were acquired with an Artis zeego system (Siemens Healthcare GmbH, Forchheim, Germany). In a single C-arm sweep of 14 s duration, 381 projection images were captured at approx. 30 Hz with an angular increment of 0.52°. A detector size of 1240 × 960 pixels with an isotropic resolution of 0.31 mm/pixel (0.21 mm in isocenter) was used. Based on prior dose measurements carried out for comparable protocols³⁹, an estimate of the effective dose applied during our acquisition can be given as 2.3 mSv. Reconstruction was performed on a grid of 256³ voxels with an isotropic resolution of 1.0 mm/voxel. At a speed of 7 ml/s, 91 ml of undiluted contrast agent was administered in the pulmonary artery. The appropriate X-ray de-

TABLE I Labeling scheme used in the experiments. For any pipeline step, $\diamond = 1 / 0$ signifies that it is activated / deactivated in a given configuration.

Label	Section	Method
$\langle \diamond, \cdot, \cdot, \cdot \rangle$	II.B.1	Streak red. (projections)
$\langle \cdot, \diamond, \cdot, \cdot \rangle$	II.B.2	Streak red. (reconstruction)
$\langle \cdot, \cdot, \diamond, \cdot \rangle$	II.C.1	Spatial denoising
$\langle \cdot, \cdot, \cdot, \diamond \rangle$	II.C.2	Temporal smoothing

lay was determined using a prior test bolus injection. A heart rate of 115 bpm was enforced by right ventricular pacing. Ten gating windows were used, each covering 10% of the cardiac cycle, resulting in ten volumes per data set for the initial images. Motion-compensated reconstruction was subsequently performed for the same phases.

III.B. Experimental Setup

For the sake of conciseness, we restrict our evaluation to the below-listed combinations of the four proposed steps corresponding to sections II.B.1, II.B.2, II.C.1 and II.C.2. To show the progression through the processing pipeline, initial reconstructions for all data sets were performed without any of the processing steps (label: $\langle 0, 0, 0, 0 \rangle$), with only the first one, two, and three steps applied ($\langle 1, 0, 0, 0 \rangle$, $\langle 1, 1, 0, 0 \rangle$, $\langle 1, 1, 1, 0 \rangle$), as well as with all four ($\langle 1, 1, 1, 1 \rangle$). To highlight the effect of each individual step, additional variants in which just this particular step is missing were generated ($\langle 0, 1, 1, 1 \rangle$, $\langle 1, 0, 1, 1 \rangle$, etc.). The labeling scheme is shown in Tab. I for quick reference. For all variants, motion estimation and compensation as summarized in section II.A were performed based on the respective initial images to obtain the final reconstruction. We use the shorthands ‘‘Initial’’ and ‘‘MoCo’’ to distinguish between initial and motion-compensated reconstructions.

Our implementation was written in C++ and has not been specifically optimized for high performance apart from GPU-based bilateral filtering as well as projection and back-projection operators, as they typically constitute the computational bottlenecks. We measured the run times of all proposed steps on a machine equipped with a quad-core 2.93 GHz CPU, 12 GB of RAM and an Nvidia Quadro FX 5800 GPU.

III.C. Evaluation Measures

For the phantom experiment, we assess the root-mean-square error (RMSE) in HU to the ground truth reconstruction $\tilde{\mathbf{I}}(\mathbf{x})$ over the ROI Ω ,

$$\varepsilon_{\text{RMS}} = \sqrt{\frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} \left(\mathbf{I}_C^{t_{\text{ref}}}(\mathbf{x}) - \tilde{\mathbf{I}}(\mathbf{x}) \right)^2}, \quad (12)$$

where t_{ref} is the cardiac phase of the static ground truth and $\mathbf{I}_C^{t_{\text{ref}}}$ denotes the corresponding motion-compensated final image. To obtain a relative RMSE, ε_{RMS} is divided by the maximum voxel intensity occurring in Ω for the reference image. Additionally, the correlation of both images is measured using the Pearson correlation coefficient (CC),

$$\rho_{\text{CC}} = \frac{\sum_{\mathbf{x} \in \Omega} (\mathbf{I}_C^{t_{\text{ref}}}(\mathbf{x}) - \mu_{\Omega}) \cdot (\tilde{\mathbf{I}}(\mathbf{x}) - \tilde{\mu}_{\Omega})}{\sqrt{\zeta_{\Omega} \cdot \tilde{\zeta}_{\Omega}}}, \quad (13)$$

$$\zeta_{\Omega} = \sum_{\mathbf{x} \in \Omega} (\mathbf{I}_C^{t_{\text{ref}}}(\mathbf{x}) - \mu_{\Omega})^2,$$

$$\tilde{\zeta}_{\Omega} = \sum_{\mathbf{x} \in \Omega} (\tilde{\mathbf{I}}(\mathbf{x}) - \tilde{\mu}_{\Omega})^2,$$

$$\mu_{\Omega} = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} \mathbf{I}_C^{t_{\text{ref}}}(\mathbf{x}), \quad \tilde{\mu}_{\Omega} = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} \tilde{\mathbf{I}}(\mathbf{x}),$$

as well as the universal image quality index (UQI) proposed by Wang and Bovik⁴⁰,

$$\rho_{\text{UQI}} = \frac{1}{N_B} \sum_{i=1}^{N_B} \frac{4 \cdot \mu_{B_i} \cdot \tilde{\mu}_{B_i} \cdot \zeta_{B_i}}{(\mu_{B_i}^2 + \tilde{\mu}_{B_i}^2) \cdot \zeta_{B_i}}, \quad (14)$$

$$\zeta_{B_i} = \sum_{\mathbf{x} \in B_i} (\mathbf{I}_C^{t_{\text{ref}}}(\mathbf{x}) - \mu_{B_i}) \cdot (\tilde{\mathbf{I}}(\mathbf{x}) - \tilde{\mu}_{B_i}),$$

$$\zeta_{B_i} = \sum_{\mathbf{x} \in B_i} (\mathbf{I}_C^{t_{\text{ref}}}(\mathbf{x}) - \mu_{B_i})^2 + (\tilde{\mathbf{I}}(\mathbf{x}) - \tilde{\mu}_{B_i})^2,$$

$$\mu_{B_i} = \frac{1}{|B_i|} \sum_{\mathbf{x} \in B_i} \mathbf{I}_C^{t_{\text{ref}}}(\mathbf{x}), \quad \tilde{\mu}_{B_i} = \frac{1}{|B_i|} \sum_{\mathbf{x} \in B_i} \tilde{\mathbf{I}}(\mathbf{x}),$$

where the $B_i \subset \Omega$, $i \in \{1, \dots, N_B\}$, are all regular blocks of size $(16 \text{ mm})^3$ that completely fit inside Ω . In our case, this resulted in $N_B = 9$ blocks in total. Computation was performed blockwise as it makes the measure more sensitive than the global correlation coefficient⁴⁰. For both CC and UQI, the obtained values are in the range $[-1, 1]$, with the maximum value of 1 corresponding to perfect correlation.

For both phantom and clinical patient data sets, SNR values are estimated to assess the magnitude of artifact patterns. For this purpose, a region $\Omega_S \subset \Omega$ in the *uncontrasted* right heart, which should appear mostly homogeneous and static, is identified in each data set. The spatial SNR is then calculated as

$$\text{SNR}_{\text{spat}} = \frac{1}{N_T} \sum_{t=1}^{N_T} \frac{\mu_{\Omega_S}^t}{\sqrt{\sigma_{\Omega_S}^{2t}}}, \quad (15)$$

$$\sigma_{\Omega_S}^{2t} = \frac{1}{|\Omega_S| - 1} \sum_{\mathbf{x} \in \Omega_S} (\mathbf{I}_C^t(\mathbf{x}) - \mu_{\Omega_S}^t)^2,$$

$$\mu_{\Omega_S}^t = \frac{1}{|\Omega_S|} \sum_{\mathbf{x} \in \Omega_S} \mathbf{I}_C^t(\mathbf{x}).$$

Similarly, we compute a temporal SNR, averaged over Ω_S , by looking at intensity statistics w. r. t. the

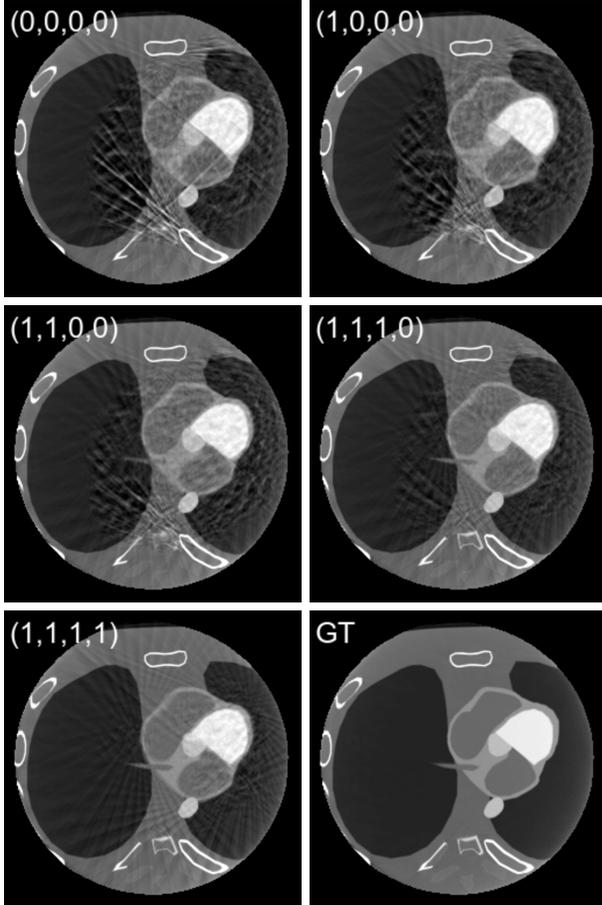


FIG. 4 Motion-compensated reconstructions of the phantom data set for configurations demonstrating the progression through the processing pipeline (cf. section III.B), and the static ground truth reconstruction (GT). Corresponding quality measures are listed in Tab. II. All images are displayed with the same window at [center, width] = [340, 2680] HU.

N_T cardiac phases,

$$\begin{aligned} \text{SNR}_{\text{temp}} &= \frac{1}{|\Omega_S|} \sum_{\mathbf{x} \in \Omega_S} \frac{\mu(\mathbf{x})}{\sqrt{\sigma^2(\mathbf{x})}}, \\ \sigma^2(\mathbf{x}) &= \frac{1}{N_T - 1} \sum_{t=1}^{N_T} (\mathbf{I}_C^t(\mathbf{x}) - \mu(\mathbf{x}))^2, \\ \mu(\mathbf{x}) &= \frac{1}{N_T} \sum_{t=1}^{N_T} \mathbf{I}_C^t(\mathbf{x}). \end{aligned} \quad (16)$$

IV. RESULTS

Quantitative results for the phantom study are listed in Tab. II, with some corresponding slice images displayed in Fig. 4. Please note that all of the compared images are reconstructed from the same raw data with a motion-compensated filtered back-

projection. Differences in the images are caused solely by the varying quality of the motion estimate for different configurations.

A comprehensive visual overview of the various configurations for one of the clinical data sets is found in Fig. 5. Figure 6 shows a comparison of configurations $\langle 0, 0, 0, 0 \rangle$ (no processing) and $\langle 1, 1, 1, 1 \rangle$ (full pipeline) for motion-compensated reconstructions of all three data sets. We highly recommend the reader to take a look at the animated versions of these figures on our website⁴¹ where differences in the temporal domain become evident. Corresponding SNR measurements are displayed in Fig. 7.

With our implementation, projection-domain streak reduction takes 0.42s for all projection images. For ten heart phases, reconstruction-domain streak reduction and spatial denoising are run in 98s and 17s, respectively. Temporal smoothing takes 66s in total, of which 34s are spent on motion detection and 32s on the actual smoothing.

V. DISCUSSION

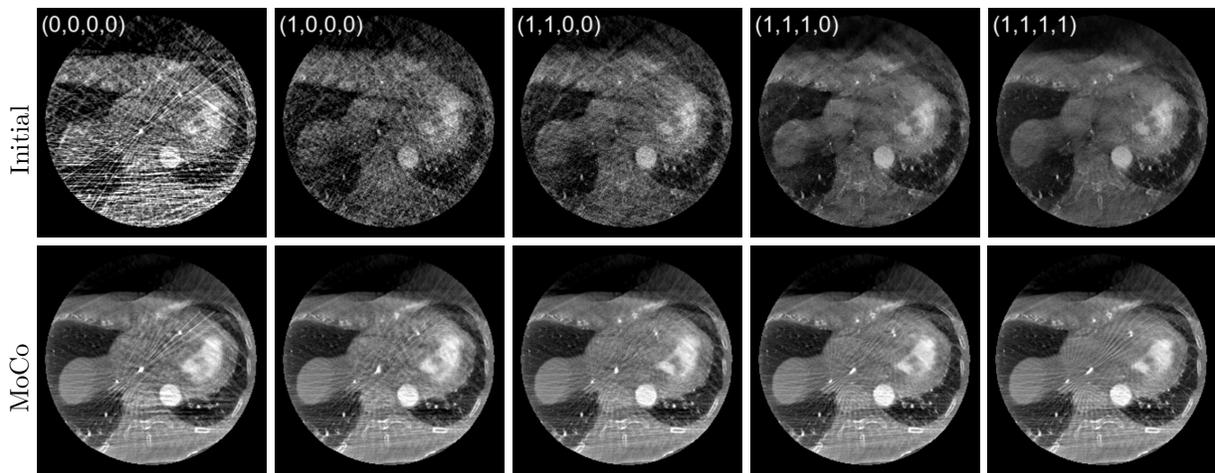
V.A. Discussion of the Results

In the phantom study, each measure is consistently improved with every pipeline step apart from a single exception in the case of UQI for $\langle 1, 1, 0, 0 \rangle$. From the configurations where singular steps are omitted, it can be seen that if either of both complementing streak reduction methods or the spatial denoising is missing, the error does not decrease much more than if we were to skip the remainder of the steps entirely. More precisely, without projection-domain streak reduction applied beforehand, the McKinnon-Bates approach alone fails to remove the artifacts reliably enough to keep them from misleading the motion estimation. If applied consecutively, however, it nearly halves the error. Similarly, spatial denoising may not improve the images by much as long as the streaks are still so prominent that they are mistaken for structure by the edge-stopping term of the kernel, yet reduces the error substantially otherwise (35.3% to 4.1%). The effect of the temporal smoothing is minor here. However, the UQI, which is much more sensitive to differences in the low-error range, still reports an increase of about 14%.

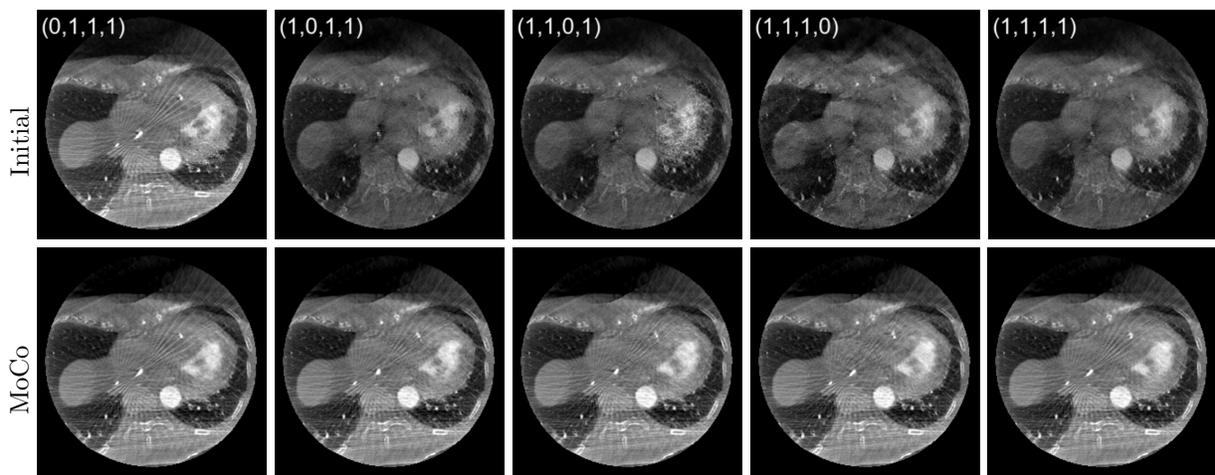
Visually, it can be confirmed that artifacts are strongly reduced. Some anatomical features such as the papillary muscles appear sharper, especially in patients 1 and 2 (cf. Fig. 6). Artifacts in the initial images, even when they only occupy nearby static regions, influence how accurately the dynamics of moving parts can be captured by typically non-local motion estimation methods such as the B-spline model based algorithm used in our framework. By reducing these artifacts, a more robust motion estimate is obtained and motion blur is eliminated more effectively.

TABLE II Quality measures computed for motion-compensated reconstructions of the phantom data set for all tested configurations: Root-mean-square error (RMSE), correlation coefficient (CC), and universal image quality index (UQI) w. r. t. the reference reconstruction, as well as spatial and temporal SNR; cf. section III.C.

Config.	RMSE [HU] (%)	CC [%]	UQI	SNR _{spat}	SNR _{temp}
$\langle 0, 0, 0, 0 \rangle$	6685.0 (89.8)	95.736	0.0055	8.96	1.08
$\langle 1, 0, 0, 0 \rangle$	3125.3 (66.1)	98.991	0.0139	10.18	1.64
$\langle 1, 1, 0, 0 \rangle$	1132.5 (35.3)	99.854	0.0097	10.99	2.39
$\langle 1, 1, 1, 0 \rangle$	96.0 (4.1)	99.994	0.8231	14.80	19.44
$\langle 1, 1, 1, 1 \rangle$	85.0 (3.6)	99.995	0.9392	20.46	229.31
$\langle 0, 1, 1, 1 \rangle$	6684.2 (89.7)	95.725	0.0048	17.43	71.36
$\langle 1, 0, 1, 1 \rangle$	3118.9 (65.9)	98.993	0.0073	17.63	105.92
$\langle 1, 1, 0, 1 \rangle$	1126.7 (35.1)	99.857	-0.0057	14.94	265.14



(a) From left to right, the progression through the processing pipeline is demonstrated.



(b) Each individual pipeline step is omitted to show its individual effect.

FIG. 5 Initial and motion-compensated reconstructions of clinical patient data set 1 for various configurations (see section III.B for details). Corresponding SNR measurements are shown in Fig. 7. All images are displayed with the same window at [center, width] = [210, 2420] HU. Note that initial images generated with projection-domain streak reduction appear darker as we do not compensate for the object mass lost due to the thresholding operation (cf. sec. II.B.1).

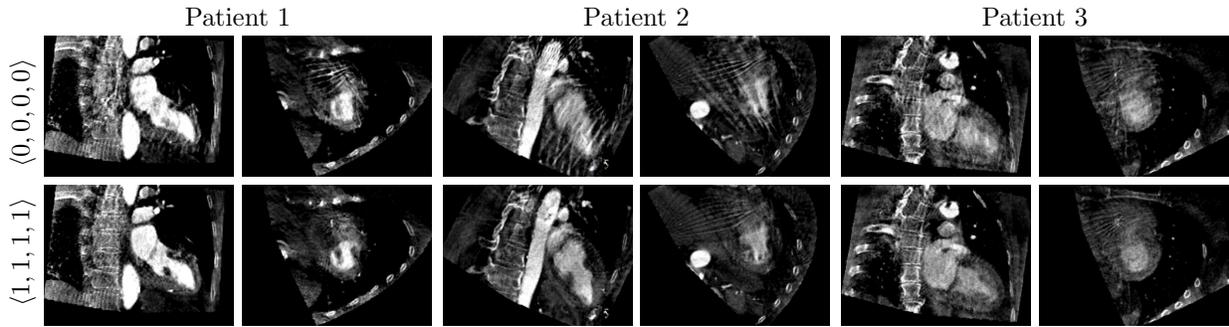
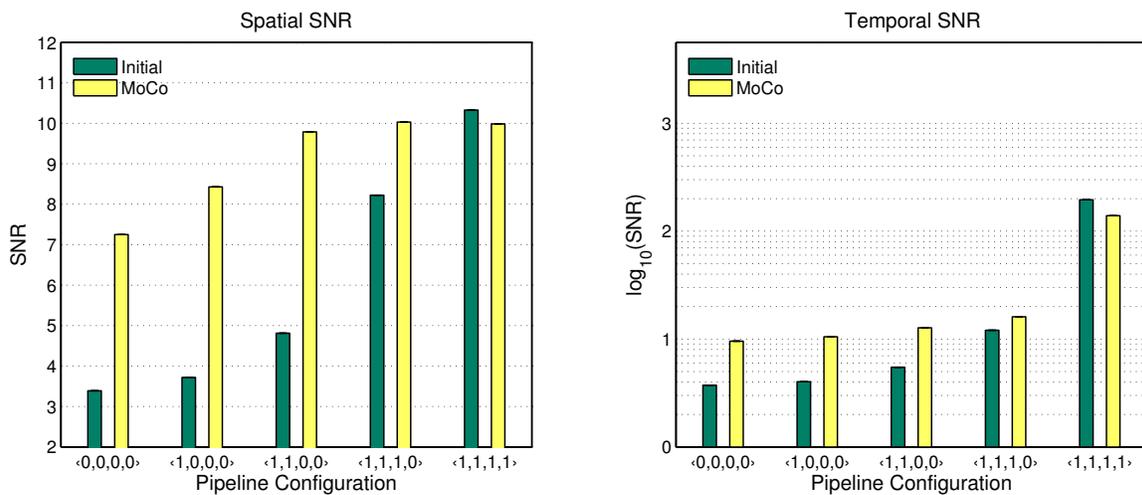
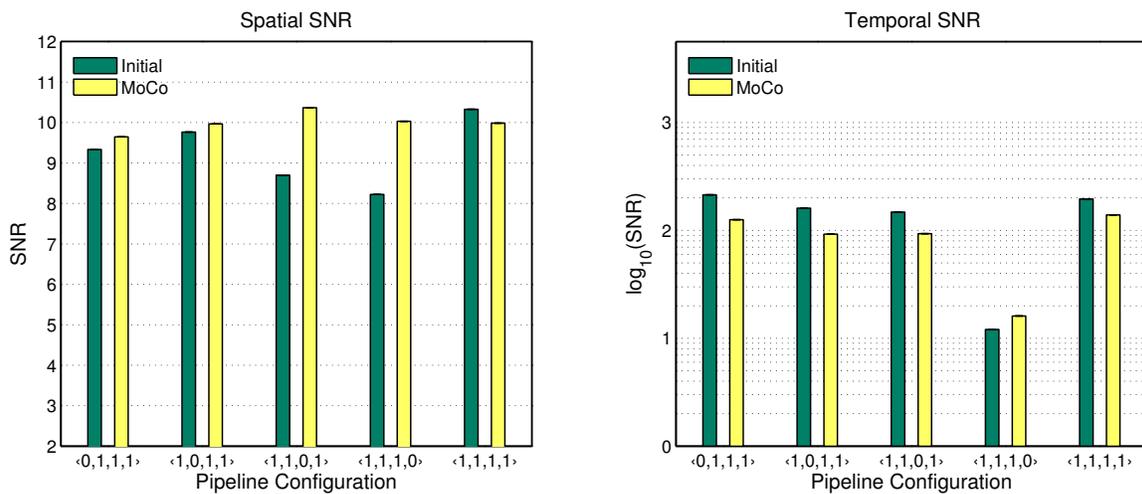


FIG. 6 Long-axis and short-axis views in motion-compensated reconstructions of three clinical patient data sets. Reconstructions were performed with (bottom row) and without (top row) the processing pipeline. All images are displayed with the same window at $[\text{center}, \text{width}] = [500, 1000]$ HU.



(a) Configurations demonstrating the progression through the processing pipeline.



(b) Omission of each pipeline step to show its individual effect.

FIG. 7 Spatial and temporal SNR measurements as described in section III.C, averaged over the evaluated clinical data sets. The temporal SNR values are shown on a logarithmic scale due to their large range.

In Fig. 7a, a rise in SNR values can be observed when progressing through the pipeline. The relative improvement is less pronounced for the final images as the motion compensation by itself helps achieve a high SNR, which is then further improved by more robust motion estimates, whereas the initial images benefit from the processing directly. Judging solely from the SNR measurements for configuration $\langle 1, 1, 1, 1 \rangle$, one may be tempted to assume that with all of the processing steps performed, motion compensation provides no further benefit or may even be detrimental. However, it has to be kept in mind that some details may be removed in the initial images and, albeit well suited for motion estimation, their rather unnatural appearance is inferior in terms of visual impression. Activating the last step increases the temporal SNR by about an order of magnitude—regardless of the remaining configuration (cf. Fig. 7b). This is mainly due to motion detection correctly identifying the region selected for evaluation as showing only little movement.

V.B. Caveats and Limitations

Regarding projection-domain streak reduction, the threshold in Eq. 1 in units of σ^* is chosen heuristically such that predominantly non-organic high-density objects are affected (cf. Fig. 1). In our experience, this is feasible without a sophisticated parameter search as the attenuation of such objects is much stronger than that of the contrasted ventricle. However, with an inappropriately chosen threshold, the ventricle boundary could potentially become blurred.

For reconstruction-domain streak reduction, no parameters had to be adjusted. A limitation of the algorithm is that not all streaks might be found: If an object edge is strongly affected by motion blur in the uncompensated reconstruction, the streaks it causes are blurred as well and are not detected as accurately in the artifact images, leading to an undercorrection.

Such as with any smoothing filter, there is a potential risk for blurring anatomical details in the denoising step. By design, the conservatively parameterized range component of the joint bilateral filter reduces this risk substantially. In addition, as it is only applied to the initial images, no blur is introduced in the final images other than that caused by motion which could not be compensated for perfectly.

In its current form, the motion detection step relies on the periodicity of the heart motion for frequency analysis. For applications where data are acquired without pacing, one may need to adapt this step to cope with a change in heart rate. One conceivable approach would be a short-time Fourier analysis, i. e. applying the Goertzel algorithm on several smaller windows with different associated heart rates.

Pacing is frequently used during therapeutic procedures in the catheter laboratory treating structural

heart disease. Traversing the venous blood system, the pacing electrode is easily placed into the right heart chambers. Whereas extremely high frequencies above 200 bpm entail the risks of circulation collapse and tachycardia, a frequency of 115 bpm as used in our experiments appears harmless. We employ pacing to obtain a sufficient number of distinct projection angles for each heart phase while avoiding impracticably long scan times. It also ensures periodicity of the motion which is helpful for gated reconstruction. With continued development of dynamic reconstruction algorithms, the number of gated projection views required may be further reduced.

VI. CONCLUSIONS AND OUTLOOK

In summary, we propose a sequence of four straight-forward, complementary processing steps to improve the quality of initial images for motion-compensation in cardiac CBCT. In combination they achieve a high SNR, which may be helpful for segmentation, and visual inspection reveals a distinct reduction of artificial motion patterns. This is reflected by the lower RMSE and better correlation with the ground truth found in our phantom study. Evaluated over various configurations, these figures of merit also suggest that the proposed order of steps provides a clear advantage over an arbitrary sequence of individual image enhancement methods.

It may be worthwhile to compare our method to sophisticated algebraic reconstruction algorithms with spatial and temporal regularization in the future. However, the fact that ideal parameterization of such a method is rather complex renders a fair comparison challenging. Encouraged by promising results accomplished for both phantom and patient data, we consider a task-based evaluation of cardiac function analysis based on C-arm CT feasible.

ACKNOWLEDGMENTS

The authors gratefully acknowledge funding by Siemens Healthcare GmbH as well as 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 would like to thank Prof. Dr. A. Rastan, Dr. B. Abt and Dr. H. Köhler from the Centre of Cardiovascular Diseases in Rotenburg an der Fulda for providing the image data. *Disclaimer:* The concepts and information presented in this paper are based on research and are not commercially available.

^{a)}E-mail: oliver.taubmann@fau.de; web: www5.cs.fau.de/~taubmann

- ¹M. 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* **15**, 718–727 (2013).
- ²G. Lauritsch, J. Boese, L. Wigstrom, H. Kemeth, and R. Fahrig, “Towards cardiac C-arm computed tomography,” *IEEE Trans. Med. Imag.* **25**, 922–934 (2006).
- ³M. Prummer, J. Hornegger, G. Lauritsch, L. Wigstrom, E. Girard-Hughes, and R. Fahrig, “Cardiac C-arm CT: A unified framework for motion estimation and dynamic CT,” *Medical Imaging, IEEE Transactions on* **28**, 1836–1849 (2009).
- ⁴S. De Buck, D. Dauwe, J.-Y. Wielandts, P. Claus, S. Janssens, H. Heidbuchel, and D. Nuyens, “A new approach for prospectively gated cardiac rotational angiography,” *Proc. SPIE* **8668** (2013), 10.1117/12.2007904.
- ⁵K. 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.* **59**, 3121–3138 (2014).
- ⁶J.-Y. Wielandts, S. De Buck, J. Ector, D. Nuyens, F. Maes, and H. Heidbuchel, “Left ventricular four-dimensional rotational angiography with low radiation dose through interphase registration,” *Europace* **17**, 152–159 (2015).
- ⁷J.-Y. Wielandts, S. De Buck, K. Michielsen, R. Louw, C. Garweg, J. Nuyts, J. Ector, F. Maes, and H. Heidbuchel, “Multi-phase rotational angiography of the left ventricle to assist ablations: feasibility and accuracy of novel imaging,” *European Heart Journal Cardiovascular Imaging* (2015).
- ⁸K. 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. Intl. Conf. on Image Formation in X-ray CT* (2014) pp. 418–421.
- ⁹C. Haase, D. Schfer, M. Kim, S. J. Chen, J. D. Carroll, P. Eshuis, O. Dssel, and M. Grass, “First pass cable artefact correction for cardiac C-arm CT imaging,” *Phys. Med. Biol.* **59**, 3861 (2014).
- ¹⁰G.-H. Chen, J. Tang, and S. Leng, “Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets,” *Med. Phys.* **35**, 660–663 (2008).
- ¹¹C. Mory, V. Auvray, B. Zhang, M. Grass, D. Schäfer, S. Rit, F. Peyrin, P. Douek, and L. Bussel, “Removing streak artifacts from ECG-gated reconstructions using deconvolution.” *Journal of X-ray science and technology* **22**, 253–70 (2014).
- ¹²C. Mory, V. Auvray, B. Zhang, M. Grass, D. Schäfer, S. Chen, J. Carroll, S. Rit, F. Peyrin, P. Douek, and L. Bussel, “Cardiac C-arm computed tomography using a 3D + time ROI reconstruction method with spatial and temporal regularization,” *Med. Phys.* **41**, 021903 (2014).
- ¹³M. 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 *Nuclear Science Symposium and Medical Imaging Conference Record*, edited by IEEE (2012) pp. 2304–2311.
- ¹⁴C. Riess, M. Berger, H. Wu, M. Manhart, R. Fahrig, and A. Maier, “TV or not TV? That is the Question,” in *Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, edited by R. M. Leahy and J. Qi (2013) pp. 341–344.
- ¹⁵K. He, J. Sun, and X. Tang, “Guided image filtering,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **35**, 1397–1409 (2013).
- ¹⁶L. A. Feldkamp, L. C. Davis, and J. W. Kress, “Practical cone-beam algorithm,” *J Opt Soc Am* , 612–619 (1984).
- ¹⁷B. Desjardins and E. A. Kazerooni, “ECG-gated cardiac CT,” *Am. J. Roentgenol.* **182**, 993–1010 (2004).
- ¹⁸D. Liu and J. Nocedal, “On the limited memory BFGS method for large scale optimization,” *Mathematical Programming* **45**, 503–528 (1989).
- ¹⁹S. Klein, M. Staring, K. Murphy, M. Viergever, and J. Pluim, “elastix: a toolbox for intensity based medical image registration,” *IEEE Trans. Med. Imag.* **29**, 196–205 (2010).
- ²⁰D. Schäfer, J. Borgert, V. Rasche, and M. Grass, “Motion-compensated and gated cone beam fil-

- tered back-projection for 3-D rotational X-ray angiography,” *IEEE Trans. Med. Imag.* **25**, 898–906 (2006).
- ²¹C. Rohkohl, G. Lauritsch, A. Nöttling, M. Prümmer, and J. Hornegger, “C-Arm CT: Reconstruction of Dynamic High Contrast Objects Applied to the Coronary Sinus,” in *Nuclear Science Symposium and Medical Imaging Conference Record*, edited by IEEE (2008).
- ²²C. Schwemmer, C. Rohkohl, G. Lauritsch, K. Müller, and J. Hornegger, “Residual Motion Compensation in ECG-Gated Cardiac Vasculature Reconstruction,” in *Proceedings of the second international conference on image formation in x-ray computed tomography*, edited by F. Noo (2012) pp. 259–262.
- ²³G. C. Mc Kinnon and R. Bates, “Towards imaging the beating heart usefully with a conventional CT scanner,” *IEEE Trans. Biomed. Eng.* **BME-28**, 123–127 (1981).
- ²⁴Z. Zheng, M. Sun, J. Pavkovich, and J. Star-Lack, “Fast 4D cone-beam reconstruction using the McKinnon-Bates algorithm with truncation correction and nonlinear filtering,” *Proc. SPIE* **7961**, 79612U–79612U–8 (2011).
- ²⁵M. Brehm, P. Paysan, M. Oelhafen, and M. Kachelrieß, “Artifact-resistant motion estimation with a patient-specific artifact model for motion-compensated cone-beam CT,” *Med. Phys.* **40**, 101913 (2013).
- ²⁶B. Lorch, M. Berger, J. Hornegger, and A. Maier, “Projection and Reconstruction-Based Noise Filtering Methods in Cone Beam CT,” in *Bildverarbeitung für die Medizin 2015*, edited by H. Handels (2015) pp. 59–64.
- ²⁷C. Tomasi and R. Manduchi, “Bilateral filtering for gray and color images,” in *Computer Vision, Sixth International Conference on* (1998) pp. 839–846.
- ²⁸G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama, “Digital photography with flash and no-flash image pairs,” *ACM Trans. Graph.* **23**, 664–672 (2004).
- ²⁹G.-H. Chen, P. Theriault-Lauzier, J. Tang, B. Nett, S. Leng, J. Zambelli, Z. Qi, N. Bevins, A. Raval, S. Reeder, and H. Rowley, “Time-Resolved Interventional Cardiac C-arm Cone-Beam CT: An Application of the PICCS Algorithm,” *Medical Imaging*, *IEEE Transactions on* **31**, 907–923 (2012).
- ³⁰L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Journal of Phys. D* **60**, 259–268 (1992).
- ³¹O. Taubmann, G. Lauritsch, A. Maier, R. Fahrig, and J. Hornegger, “Keeping the Pace: Heart Rate Informed 3-D Motion Detection for Adaptive Temporal Smoothing,” in *Proc. 13th Intl. Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine* (2015) pp. 532–535.
- ³²G. Goertzel, “An algorithm for the evaluation of finite trigonometric series,” *American Mathematical Monthly* (1958), 10.2307/2310304.
- ³³F. Harris, “On the use of windows for harmonic analysis with the discrete Fourier transform,” *Proceedings of the IEEE* **66**, 51–83 (1978).
- ³⁴O. Taubmann, G. Lauritsch, A. Maier, R. Fahrig, and J. Hornegger, “Estimate, Compensate, Iterate: Joint Motion Estimation and Compensation in 4-D Cardiac C-arm Computed Tomography,” in *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)*, *Lecture Notes in Computer Science*, Vol. 9350, edited by N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi (2015) pp. 579–586.
- ³⁵W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, “4D XCAT phantom for multimodality imaging research,” *Med. Phys.* **37**, 4902–4915 (2010).
- ³⁶A. 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.* **40** (2013).
- ³⁷<http://physics.nist.gov/PhysRefData/Xcom/html/xcom1.html>.
- ³⁸www5.cs.fau.de/fileadmin/Persons/TaubmannOliver/supplementary_material/pipeline/phantom_data.zip.
- ³⁹C. Hohl, J. Boese, N. Strobel, R. Banckwitz, G. Lauritsch, G. Mühlenbruch, and R. Günther, “Angiographic CT: Measurement of patient dose,” in *European Congress of Radiology* (2008).
- ⁴⁰Z. Wang and A. Bovik, “A universal image quality index,” *IEEE Signal Proc Let* **9**, 81–84 (2002).
- ⁴¹www5.cs.fau.de/fileadmin/Persons/TaubmannOliver/supplementary_material/pipeline/animations.html.