

RESEARCH ARTICLE

Four-dimensional computed tomography of the left ventricle, Part I: Motion artifact reduction

Jed D. Pack¹ | Ashish Manohar² | Sathish Ramani¹ | Bernhard Claus¹ |
 Zhye Yin¹ | Francisco J. Contijoch^{3,5} | Andrew J. Schluchter³ |
 Elliot R. McVeigh^{3,4,5}

¹Radiation Systems Lab, GE Global Research, Niskayuna, New York, USA

²Department of Mechanical and Aerospace Engineering, UC San Diego School of Engineering, La Jolla, California, USA

³Department of Bioengineering, UC San Diego School of Engineering, La Jolla, California, USA

⁴Department of Medicine, Division of Cardiology, UC San Diego School of Medicine, La Jolla, California, USA

⁵Department of Radiology, UC San Diego School of Medicine, La Jolla, California, USA

Correspondence

Jed D. Pack, Radiation Systems Lab, GE Global Research, KW C1300A, One Research Circle, Niskayuna, NY 12309, USA.
 Email: pack@ge.com

Taxonomy: Motion management (3.22); CT Development (3.10); 4DCT (3.14), Registration (3.16)

Funding information

NHLBI; National Institutes of Health, Grant/Award Number: R01 HL144678; American Heart Association, Grant/Award Number: 20PRE35210261

Abstract

Purpose: Standard four-dimensional computed tomography (4DCT) cardiac reconstructions typically include spiraling artifacts that depend not only on the motion of the heart but also on the gantry angle range over which the data was acquired. We seek to reduce these motion artifacts and, thereby, improve the accuracy of left ventricular wall positions in 4DCT image series.

Methods: We use a motion artifact reduction approach (ResyncCT) that is based largely on conjugate pairs of partial angle reconstruction (PAR) images. After identifying the key locations where motion artifacts exist in the uncorrected images, paired subvolumes within the PAR images are analyzed with a modified cross-correlation function in order to estimate 3D velocity and acceleration vectors at these locations. A subsequent motion compensation process (also based on PAR images) includes the creation of a dense motion field, followed by a backproject-and-warp style compensation. The algorithm was tested on a 3D printed phantom, which represents the left ventricle (LV) and on challenging clinical cases corrupted by severe artifacts.

Results: The results from our preliminary phantom test as well as from clinical cardiac scans show crisp endocardial edges and resolved double-wall artifacts. When viewed as a temporal series, the corrected images exhibit a much smoother motion of the LV endocardial boundary as compared to the uncorrected images. In addition, quantitative results from our phantom studies show that ResyncCT processing reduces endocardial surface distance errors from 0.9 ± 0.8 to 0.2 ± 0.1 mm.

Conclusions: The ResyncCT algorithm was shown to be effective in reducing motion artifacts and restoring accurate wall positions. Some perspectives on the use of conjugate-PAR images and on techniques for CT motion artifact reduction more generally are also given.

KEYWORDS

4DCT, cardiac CT, motion correction

1 | INTRODUCTION

This is the first in a series of two papers exploring the feasibility of using computed tomography (CT) to produce temporally accurate 4D images of the moving heart. This first paper describes the motion artifact reduction algorithm we have developed, whereas

the second paper¹ focuses on the clinical aspects of dyssynchrony analysis^{2–5} and evaluates the performance of this clinical task.

Modern CT scanners can provide a full 4D dataset from a low-dose scan of a single heartbeat. Unfortunately, during phases of rapid ventricular motion, CT images of the myocardium can be distorted by motion

artifacts. If left unchecked, these distortions can produce the appearance of a dyssynchronous contraction even for a healthy, functioning heart.^{6,7} This paper describes a general cardiac CT motion artifact reduction algorithm that leverages the power of conjugate pairs of partial angle reconstruction (PAR) images and illustrates the image quality improvement that can result from its use and the improved fidelity of left ventricle (LV) boundary localization. We use the term “conjugate” to mean “separated by 180 degrees,” that is, ignoring the cone angle; this means the same rays are measured in the opposite direction.

The remainder of the paper is organized as follows: In Section 2, a survey of recent developments in motion artifact reduction technology for CT is given. An introduction to and a root-cause analysis of gantry-angle-dependent motion artifacts around the LV is also given. In Section 3, the motion artifact reduction method is described in detail. Section 4 describes the experimental setup used to test the applicability of our algorithm to the motion of the LV and reports on our results. In Section 5, a few concluding remarks and perspectives for future research are given.

2 | BACKGROUND

2.1 | Recent progress in motion-compensated reconstruction

Motion artifact reduction is a critical technology for CT. Much of the cost of CT scanners can be traced to the need for robust, low-noise, artifact-free cardiac imaging. With motion artifact reduction, the cost/quality trade-off for cardiac imaging can be greatly improved. Specifically, coronary CT angiography (CCTA), heart valve imaging, calcium scoring, synthetic fractional flow reserve estimation, myocardial perfusion imaging, pediatric imaging of congenital heart defects, and cardiac wall motion analysis are all most effective when the whole heart can be imaged (in some cases repeatedly) during a fraction of a heartbeat. In order to perform such scans effectively, a third-generation CT scanner must have a high-power X-ray source and a large detector mounted on a gantry that is engineered to withstand the g -forces associated with rotating at a rate of three to five times every second. Furthermore, dual-source CT scanners have two source/detector imaging subsystems mounted on the same gantry in order to further improve temporal resolution by a factor of 2. All this hardware is quite expensive. Robust motion artifact reduction algorithms can provide superior image quality even when several of these design considerations are relaxed. For example, a slower gantry rotation can be tolerated on a single-source system. Using a slower gantry rotation not only relaxes the engineering considerations for the gantry but also reduces the tube power specifications needed to

reach a prescribed mAs level during a scan (alternatively, it allows more flux to be delivered for high BMI patients, thus reducing image noise).

The most widely prescribed cardiac CT imaging protocol is the CCTA, which is being used extensively to avoid invasive conventional coronary angiography (CA) scans that are diagnostic in nature (billions of dollars are spent annually on diagnostic CA procedures). However, there is growing interest in several other cardiac CT imaging capabilities. In particular, wall motion analysis can be done using CT by scanning over an entire heart cycle. Such a scan can produce a high-quality 4D dataset for analyzing dyssynchrony in the LV. Although the resulting CT image volumes have very low noise and clear detail compared to echo images that are often used for the same purpose, four-dimensional CT (4DCT) images of a healthy beating LV can (falsely) appear to have dyssynchrony due only to motion artifacts, which rotate in a gantry angle-dependent way in a 4D dataset. Such false dyssynchrony can corrupt measurements of the local timing of LV wall contraction.⁷ In this paper, a method for motion artifact reduction of the heart within the context of wall motion analysis is described and evaluated.

Most motion artifact reduction methods can be split into two major steps: motion estimation and motion compensation. The motion compensation problem has at least two solutions that work remarkably well. The first is the warp-then-backproject approach^{8,9} and the second is the backproject-then-warp approach.^{10–13} The backproject-then-warp approach also has a variation that does not require explicit backprojection of each view subset but instead uses 2D image filters to produce PARs.^{14,15}

The motion estimation problem is more challenging and has been an active area of research for some time. As detailed in Ref. [16], a number of different estimation approaches have been used over the last several years. Motion can be estimated by iteratively updating motion fields until the warped, forward projected volume matches the measured projection data.¹⁷ Other algorithms (e.g., MAM and PAMoCo) use various image metrics to iteratively refine motion estimates.^{18,19} Some approaches use (image and/or feature) registration between a series of half-scan reconstructions.^{11,20} Recently, deep learning networks have also been used effectively to estimate and correct for motion in CT.^{16,21,22} Finally, the concept of using PARs has been used to facilitate robust motion estimation, with correspondences being found in conjugate pairs of PAR images.^{15,23}

It is interesting to note that backproject-then-warp motion compensation as well as PAR-based motion estimation both rely on essentially breaking up the sinogram data into smaller sections in time and representing these in the image domain as PAR images. The key insight driving both of these innovations is that the CT motion artifact problem is not necessarily a hardware

problem. The data coming from a CT scanner actually has excellent temporal resolution as each sinogram view is acquired in a fraction of a millisecond. Instead, the CT motion artifact problem can be considered a reconstruction software problem because the traditional reconstruction does not account for the fact that the various views are acquired at different motion states, yielding a full sinogram dataset that is inconsistent. By subdividing the sinogram data used in reconstruction into smaller chunks and handling each correctly (e.g., warping according to the relevant object motion state), motion artifacts can be dramatically reduced. PAR-based motion estimation and compensation have been successfully used in commercially available motion artifact reduction technology for nearly a decade now.^{24,25}

There are at least three different ways that PAR images can be generated. In one approach,²³ the actual sinogram data is segmented into chunks (the chunks usually overlap and a weighting function is used to compensate for the overlap) and these chunks are reconstructed by filtered backprojection with band-pass filtering. In a second approach,¹⁹ the reconstructed image volume can be forward projected to produce synthetic sinogram data, and this synthetic sinogram can then be reconstructed, thus avoiding the need for the original projection data altogether. This second approach (as well as the third approach below) has a potential limitation due to the fact that with a Parker-style view weighting function, some lines through the origin of 2D Fourier space contain data from two different views that are separated in time by $\sim 1/2$ the gantry period, thus limiting the angular range of views that can be accurately generated to less than 180 degrees. This can be overcome by using multiple image volumes reconstructed from overlapping view angle ranges.¹⁴ The amount of overlap depends on the effective view weighting used in the reconstruction, but for a typical cardiac reconstruction, the temporal sampling interval should not be much larger than the time it takes to rotate ~ 130 degrees. The third approach¹⁴ (which also only requires images) is to apply 2D image filters representing wedges in the Fourier domain. Any of these three approaches can work well, though the choice can have important implications for data requirements and computational speed.

2.2 | CT motion artifact characteristics

Artifacts that result from motion during a CT scan have different characteristics than artifacts that result from motion during other imaging acquisitions. These unique characteristics can be better understood by analyzing the order in which Fourier data is acquired in CT as a result of the Fourier slice theorem. A good understanding of these unique characteristics can help to guide the development of motion artifact reduction algorithms and

can also help to understand the clinically relevant problem of false dyssynchrony⁷ that can arise in the analysis of non-motion-corrected 4DCT images of the LV.

To motivate this discussion, consider Figure 1. The images in the left column show snapshots of four dynamically changing objects with the corresponding motion fields overlaid as vectors. The top two rows contain a small disc that is moving with constant velocity to the right (similar to a moving coronary artery). The third row is a much larger disc (perhaps representing a heart chamber or aorta) that is also moving to the right. The fourth row is a large disc that is contracting (like the LV during systole). The second column shows a snapshot of each of the four images from the first column that is later in time. The time elapsed between the image pair in the top row corresponds to the time it takes a CT gantry to complete a full rotation, whereas the time interval for the image pairs in rows 2–4 is the time it takes the gantry to rotate by 210 degrees (similar to the data range used in many cardiac reconstruction algorithms). The third column shows an average of multiple time frames within the temporal window bounded by the times of the first two sets of images. The blurring in column 3 is very severe—the motion either severely attenuates or even eliminates some of the higher frequency structures (e.g., edges). An attempt at deblurring with a linear filter would require a strong amplification of these frequencies and would, therefore, yield very high noise amplification. In contrast, the images in the fourth column, though they are also spatially distorted retain all high-frequency information. The third column represents what would happen if all spatial frequencies were measured continuously and then averaged over a time window.

The fourth column of images is representative of what happens in CT. To produce each of these images, the corresponding time series of images were each transformed to the Fourier domain and only the Fourier data that would have been acquired during the small temporal window of that particular snapshot was preserved. This was done by applying a wedge-shaped filter in the Fourier domain as the Fourier slice theorem indicates that a small set of adjacent views (those acquired in a small time window) correspond to a set of angularly adjacent diameters in 2D Fourier space. This process is illustrated in Figure 2 (only the real part of the Fourier data is displayed). By accumulating these images, the Fourier data of the images in the original image series is preserved, provided that the time window corresponds to less than 180 degrees of gantry rotation. For a typical “half-scan” acquisition used in cardiac imaging, a small portion of the Fourier data in the image obtained from the summation of the temporal views is corrupted by mixing with information from conjugate views (because the scan typically uses views from 180 degrees + a relatively small angle). In regions of Fourier space that are sampled only once, each Fourier transform of the object

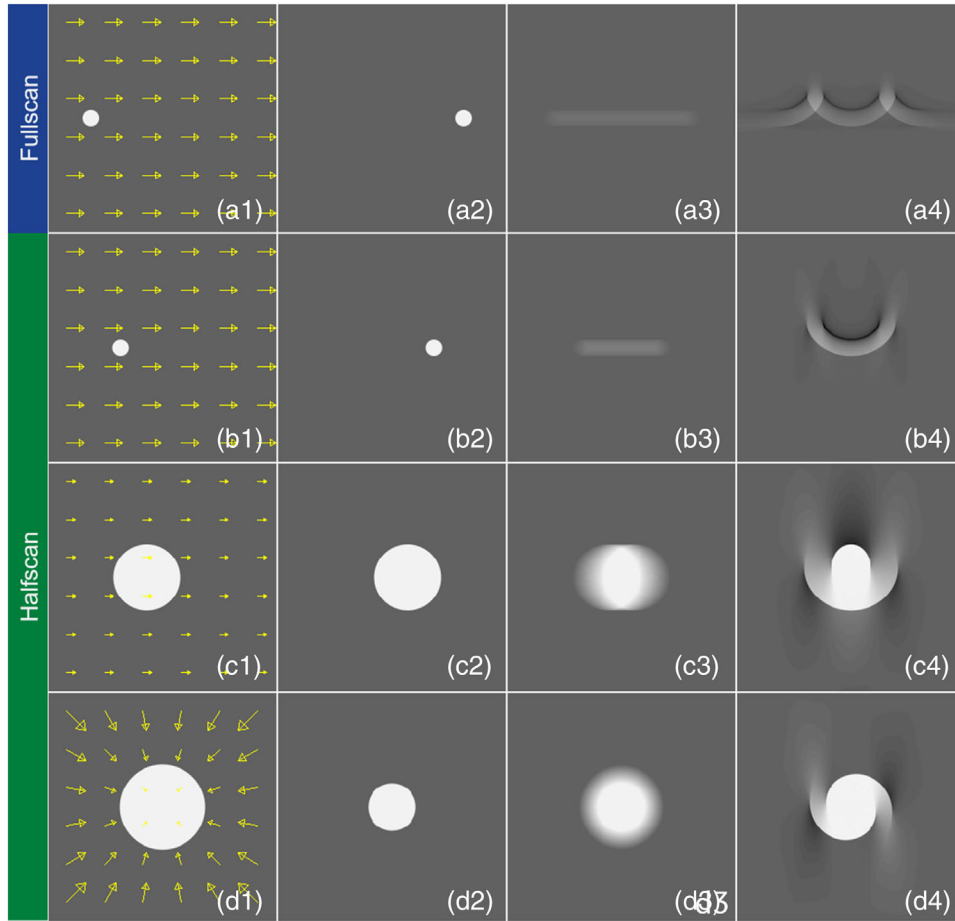


FIGURE 1 CT motion artifact characteristics are illustrated here with four simple dynamic objects, including three discs moving with constant velocity (first three rows) and one disc that is contracting (bottom row). The first two columns (from left) show snapshots in time at the beginning and end of the scan window, respectively. The third column is a simple average of the images representing each motion state within the time window. Due to the Fourier slice theorem and rotational nature of CT data acquisition, the fourth column images are more representative of motion artifacts in practice. Further details on the synthesis of these (fourth column) images are given in Figure 2. CT, computed tomography




FIGURE 2 The process used to emulate CT motion artifacts for Figure 1 is illustrated here using a subset of five time frames. The images that represent snapshots in time are converted to Fourier space, where a very small wedge (exaggerated here for clarity) is extracted. These are summed together (with an optional weighting function) and then converted back to image space using a 2D inverse discrete Fourier transform. CT, computed tomography




FIGURE 3 The four images in this figure are generated from the same object motion and the same data acquisition duration. The starting gantry angle of the emulated X-ray source is progressively shifted by 45 degrees for each image. Gantry angle dependency like this can introduce errors in the estimated local wall motion of the left ventricle

is merely phase shifted (the magnitude is unchanged from the original). This means that a linear process can deblur the image successfully without noise amplification (unlike in the case of a continuous Fourier measurement described previously for column three in Figure 1).

There are several observations that can be made from the motion artifacts seen in the fourth column of Figure 1. First, as noted by others,²⁶ the small disc does not appear to be moving in a straight line. In fact, it appears to be moving on a cycloid. This fact will be revisited later in the paper as motivation for using non-adjacent (180-degree-conjugate) time frames for motion estimation. A second observation is that the motion artifacts of the large disc in rows 3 and 4 are not obvious. For a local region of a large (low-curvature) surface, the Fourier information that is relevant for localizing that region of the surface is acquired by X-ray paths tangent to that surface. This means that there is effectively a temporal skew in the final reconstructed image that is dependent not only on the orientation of the surface but also on the gantry angle of the scanner relative to a selected motion state. Consider, for example, what happens if we regenerate image d4 (Figure 1) for a series of different starting gantry angles (SGAs) (while keeping the motion and time window identical). Such images are shown in Figure 3 for four SGAs with a 45-degree increment. For each location on the boundary of the object, we find that the reconstructed position of the edge of the boundary is dependent on the SGA. We refer to this effect as the SGA dependence effect.²⁷

An illustration of the SGA dependence of motion artifacts is given in Figure 4. As (as mentioned earlier) the edge locations in a CT image are defined based on the time at which the tangent X-ray paths are acquired, this figure shows these surface-tangent ray paths at a series of times for two different SGAs separated by 45 degrees. There is a clear resemblance between the union of these tangential ray paths and the shape of the corresponding motion artifact images in Figure 3.

Clinically, the SGA-dependent spatial skew produces prominent dynamic spiraling artifacts when a 4DCT dataset is viewed in a cine loop (as a movie). Furthermore, the positional errors that result near the

endocardial contour can cause errors in the estimation of the time to onset of shortening in functional analysis of the LV. Our motivation for applying motion artifact reduction in this paper is to resolve these SGA-dependent motion artifacts in order to produce more reliable metrics for measuring the relative timing of local LV contraction, leading to an improvement in cardiac resynchronization therapy (CRT) procedure planning.

3 | RESYNCCT METHOD

In order to perform local cardiac functional analysis of the LV, motion-corrected images are required at a series of phases throughout the heart cycle. For a heart rate of 60 bpm and a temporal reconstruction spacing of 70 ms, for example, it takes 15 image volumes to cover the heart cycle. This is different from what is typically done for a CCTA image volume, where the goal is to produce a single high-quality “frozen” snapshot of the heart for the purpose of visualizing the coronary arteries. In particular, motion must be estimated over a much longer time interval for functional analysis. A motion model (e.g., polynomial in time), which is appropriate for estimating local cardiac motion over the short time period required for CCTA (<200 ms), will not be ideal to model the local motion of the heart over an entire heart cycle. One option for dealing with this condition is to estimate motion for the entire temporal window with a complex model by considering data from all time windows. Another option (chosen for the work presented here) is to process each phase separately with lower order models. Not only does the second option allow the reuse of existing motion models, but it also enables parallelization of the processing for computational performance purposes since each output image volume is computed separately. For the remainder of the paper, the method that is used to produce a single output volume in the time sequence is described, bearing in mind that the full algorithm requires repeating this process for each prescribed output phase.

For each output phase, the motion artifact reduction algorithm (ResyncCT) estimates motion over a field of image locations (points) and then interpolates these




FIGURE 4 This figure illustrates why a contracting disc results in a spiral-shaped object from the reconstruction of a ~ 180 -degree CT acquisition. Each view defines the position of the boundaries that are tangent to rays composing that particular view. The composition of the red tangent lines from all views traces out the spiral pattern. CT, computed tomography

motions to build a dense 4D motion field, which is then used for motion compensation. Both the motion estimation and the motion compensation are based on partial angle images as described in the background previously. The steps of the method are listed as follows:

1. generation of PAR images,
2. automated positioning of motion estimation points,
3. motion estimation with conjugate cross-correlation,
4. construction of dense motion field from point motion estimates,
5. PAR-based motion compensation.

3.1 | Generation of PAR images

As methods for PAR image generation (step 1 mentioned earlier) and PAR-based motion compensation (step 5) have been published previously as described in Section 2.1,^{10,13–15,19,23} these are addressed only very briefly here so as to focus on the implementation of steps 2–4. Three different PAR image generation approaches were described in Section 2.1. For ResyncCT, an image domain filtering method¹⁴ for PAR generation is used whenever PAR images are needed (this includes steps 2, 3, and 5).

3.2 | Automated positioning of motion estimation points

The ResyncCT partial angle-based motion compensation method requires a dense 4D motion field in order to warp each of a temporal series of PAR volumes. The full 4D motion field has a large number of parameters, and it is not possible to robustly estimate them all. Thus, if each voxel moved independently, motion estimation

would be hopeless. In reality, the motion fields are quite smooth spatially, which allows the motion field to be modeled with a smaller set of parameters. Hence, for ResyncCT, motion is estimated at a somewhat sparse series of points throughout the cardiac volume. Results are best when the set of points are customized to the patient anatomy, rather than being predetermined (e.g., on a fixed 3D grid). In particular, points can be placed at various features throughout the heart.

Both the strategy for finding the points at which motion is estimated and the motion estimation itself are based on conjugate-PAR image pairs. Conjugate-PAR image pairs are separated from one another in time by one half gantry rotation. To be more precise, PAR images that are reconstructed from parallel rebinned data or which are generated by Fourier domain wedge filters represent a snapshot in ray angle rather than in time (these two only perfectly line up along the central axis of the scanner). As a result, all motion estimates are actually computed as a function of the local ray angle rather than as a function of time. Conjugate-PAR image pairs are dissimilar only where there is motion (ignoring temporal changes in contrast); therefore, taking a difference of two conjugate-PAR volumes will highlight regions of interim motion. Due to high motion gradients near the pericardial boundary, before generating the PAR images, the magnitude of the HU drop at the heart–lung boundary is reduced by applying a simple threshold on low voxel intensities. Three conjugate-PAR image pairs are generated, with the middle pair being centered around the angle of the X-ray source at the current reconstruction phase. We call this angle β_c . One image of the middle pair is advanced by 90 degrees from β_c , whereas the other precedes β_c by 90 degrees. The other two pairs are shifted in either direction from this middle pair by a spacing of roughly 1/8 rotation (see Figure 5). For each pair, a high-pass filter is applied to remove some




FIGURE 5 Placement of motion estimation points (as well as motion estimation) is based on the analysis of conjugate-PAR image pairs, such as those illustrated earlier. PAR, partial angle reconstruction




FIGURE 6 After masking the region outside the heart, a series of filters on two separate reconstructions can produce a mapping of the regions where features are moving. These regions (shown with a red overlay) are good candidate locations for motion estimation

of the lowest spatial frequencies (these low frequencies can be influenced strongly by any changes in contrast that may occur during the half-rotation) and the difference is taken. After filtration, the mean of the absolute value of the three different images is computed. Finally, this mean image is smoothed with a Gaussian kernel. The smoothed image is referred to as a difference map and highlights areas of moving features (e.g., chamber boundaries, moving vessels, or valves). Figure 6 shows a

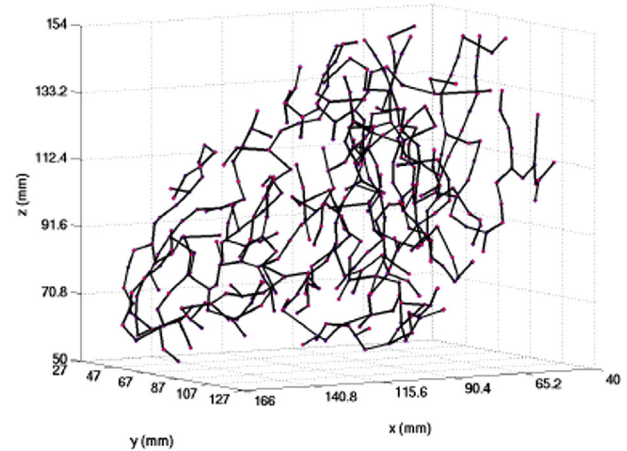


FIGURE 7 One example of a coarsely sampled tree built of points that can be used in the ResyncCT motion estimation process

sample 2D axial slice with the corresponding difference map overlaid.

In order to prevent points from being placed outside the heart, a fully automated heart segmentation routine is used and the difference map is masked to the resulting cardiac region (this masking is reflected in Figure 6). Points are placed by repeatedly selecting the brightest point in the difference map volume and zeroing out a small neighborhood until the global maximum falls below a threshold. The selected points are then formed into a minimum-spanning tree using Kruskal's algorithm,²⁸ which provides connectivity between neighboring points. Finally, thinning along the tree ensures neighboring points are separated by a distance that is close to a specified value (e.g., 7 mm), while reducing the number of points from several thousand to a few hundred for a typical adult heart. Figure 7 shows a sample 3D tree (with a fairly large point spacing) built from the initial points.

3.3 | Motion estimation with conjugate cross-correlation

Once points are identified, the next step is to estimate the motion of each point (displacement from its position at a central reference time) as a function of time. For ResyncCT, the motion at each point is modeled with a 3-element velocity vector and a 3-element acceleration vector. In order to estimate these parameters, a set of N conjugate-PAR image pairs are produced, similar to what was done for the positioning of motion estimation points earlier. Details of the PAR image generation process, including the angular "width" of the estimation PAR images, impact PAR spatial resolution in the direction of the central view rays, temporal sensitivity, noise, and robustness to contrast dynamics. Although a wide range of values works well, we suggest using an angular

width of roughly 20–60 degrees. We denote the central angle of each PAR image pair as β_i , with i being an index from 1 to N and $\beta_i = \beta_c + D_\beta(i - (N + 1)/2)$ for all i from 1, ..., N . Here, D_β is the angular PAR sampling and is chosen such that the $(N - 1)D_\beta$ covers the desired angular range. For a quadratic motion model, the angular range can be chosen between 50 and 170 degrees, with N between 2 and 15 ($N = 3$ and $D_\beta = 56$ degrees were used to produce results in this paper). A brief study illustrating some of the benefits that come from using multiple PAR pairs is detailed in a supplement to this paper. Only the linear component of the motion can be estimated when N is one. The two PAR images in each pair are interpolated onto a uniform 3D grid centered at the estimation point to produce a pair of subvolumes. The nominal size of the subvolumes was 47 mm in each dimension. We denote the subvolume built from data near $\beta_i - 90$ as A_i and the subvolume built from data near $\beta_i + 90$ as B_i . For each subvolume pair, a “modified cross-correlation map” is produced that indicates how well the local region of A_i matches B_i for each shift in a 3D grid. A_i is limited to a local region by applying a spherically symmetric weighting function to the source volume that decays with radius. The weighting function used is very smooth and decreases monotonically from 1 to 0 with radius, reaching a value of 0.5 at a distance of roughly 11 mm. Each modified cross-correlation is computed as follows:

$$\chi_{AB}[m] := \frac{\langle w \odot \mathcal{M}(A), S(B, m) \rangle}{\max \{ \|\sqrt{w} \odot \mathcal{M}(A)\|, \|\sqrt{w} \odot \mathcal{M}(S(B, m))\| \}} \quad (1)$$

This can be computed simultaneously for all shift vectors (m) on a centered 3D grid using a series of 3D FFT operations. In the previous equation, the circled dot represents element-wise multiplication. Also, w is the local weighting function mentioned previously, $S(B, m)$ indicates the result of shifting B by a 3-element index vector m , and the script \mathcal{M} function means that a constant is subtracted so as to make the DC value zero after weighting with w . In addition, a symmetric operation to that mentioned previously is done by determining how well the local region of B matches with A under the opposing shifts ($-m$) and the two results are averaged. Applying both symmetric operations effectively elongates the weighting function w in the direction of each candidate motion while preserving the convolutional nature of the computation. The result is a correlation map: a quantitative indicator showing how well the conjugate-PAR images match at the subvolume A under any shift (up to a limit high enough to capture any realistic cardiac motion magnitude). Once the N cross-correlation maps for each estimation point are computed, a velocity vector at each point is found by averaging the N shift vectors (one from the peak of each correlation map). Provided that $N > 1$, an acceleration vector at each point can also be found by quantifying how the correlation changes

with the angle (time) index i (we used the simple difference between the first and last of the three shift vectors). The motion estimation process is summarized in Figure 8.

3.4 | Construction of the dense motion field from point motion estimates

After motion estimation using conjugate-pair PAR images, an estimate of the velocity and acceleration vectors at all selected points in the volume is in hand. The next step is to produce a dense, voxelized estimate of the motion, such that each voxel can be warped appropriately for each time frame of the PAR-based motion compensation. This is accomplished by interpolating the six scalar parameters (3D velocity and acceleration) that describe the motion onto a dense grid and then using these parameters to build the motion path at each voxel. A 3D interpolation weighting function is defined in the neighborhood of each point that defines how strongly each nearby voxel will be influenced by the corresponding motion estimate. There are many ways to define these functions (e.g., using Voronoi maps, triangulated irregular networks, splines, and inverse distance weighting). The method used in ResyncCT involves defining an effective radius for each unit vector direction of each point. The value of the weighting function at each voxel is then uniquely defined by a 1D tapering function that drops smoothly from 1 to 0 and takes as its input the distance of the voxel from the selected point, normalized by the effective radius that is defined for the direction from the point to the voxel. The effective radius defines the distance at which the tapering function reaches 0.5 and is modulated in a smooth way based on the locations of other nearby points. We used a maximum (pre-modulation) value of about 15 mm. Once every point has a weighting function defined in this way, the functions are normalized such that the sum of all weighting functions across the 3D volume is equal to a single 3D function we call the compensation weight, which is defined for each voxel. The compensation weight drops smoothly and somewhat quickly to zero outside the heart, which helps avoid contamination of the image in areas (such as the liver and chest wall) that border the heart, but move much less. Through most of the heart, the compensation weight is very close to 1, except in rare areas where all estimation points are too far away to have a reliable influence.

3.5 | PAR-based motion compensation

After the six motion parameters are interpolated to a voxel grid using the weighting functions described earlier, the final step of ResyncCT (PAR-based motion compensation) can be performed. In the compensation

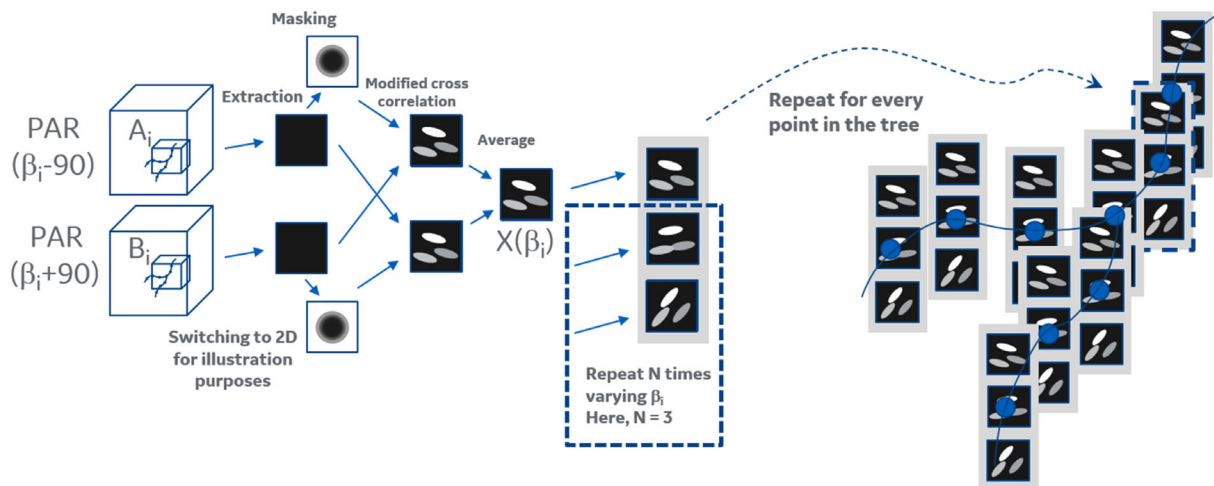


FIGURE 8 Illustration of the motion estimation process. Conjugate-PAR images are processed with a modified 3D cross-correlation process (see Equation 1) to identify spatial correspondences. This is repeated for N pairs of images that are offset in view angle. All of the preceding is again repeated for subvolumes centered on each motion estimation point. Finally, a 3D velocity and a 3D acceleration vector are chosen for each point that are most consistent with the cross-correlation maps. PAR, partial angle reconstruction

step, the six interpolated parameters (3D velocity and acceleration vectors) are used to compute the voxel displacement ($\Delta \underline{x}$) at all required times based on the prescribed angular (temporal) sampling using the quadratic motion model:

$$\Delta \underline{x} = \underline{v}t + \frac{1}{2} \underline{a}t^2 \quad (2)$$

As mentioned in Section 3.2, we used the ray angle θ (relative to a reference θ_{ref}) for t rather than time. For simplicity, we set θ_{ref} to a constant value (β_c) for each output volume, but it can alternatively be defined in a voxel-dependent way based on the ray angles measured at a single time instance. The PAR images are then warped according to these motion fields, weighted, and summed. For results herein, the weighting was chosen to match the (half-scan) data window of the input images. It turns out that the temporal PAR sampling required for good image quality is not high,¹³ so the motion compensation can be done quickly. We used a sampling of one PAR volume every 9 degrees in order to accommodate very fast motions and recommend sampling at least every 30 degrees.

As mentioned previously, the entire ResyncCT process is repeated at each desired output phase to produce a 4D motion-corrected dataset. Despite the fact that the processing is done independently for each target time frame, very good temporal consistency in the quality of motion artifact reduction is normally achieved.

4 | EXPERIMENTAL RESULTS

The ResyncCT motion artifact reduction algorithm described previously has been tested on both a moving

3D printed phantom and on clinical data. For phantom tests, a dynamic model was created that represents the temporal evolution of a single 2D short-axis slice of a dyssynchronous LV. This model was 3D printed in such a way that each z-slice of the printed phantom is a snapshot of the 2D model in time. The physical design ensured that when the phantom was moved perpendicular to the scan plane at a rate of 30 mm/s, the effect within the scan plane was the same as a dyssynchronous ventricle beating at 60 beats per minute. The physical phantom as printed is a hollow shell with just one material (proprietary clear photopolymer resin FLGPCL04, Formlabs Inc., Somerville, MA), representing the myocardial tissue. The phantom was then filled with a solution of water mixed with an iodinated contrast agent (10% by the volume solution of Visipaque 320, General Electric Healthcare, Chicago, IL) in order to represent the contrast-enhanced blood that fills the LV during a contrast-enhanced cardiac CT scan. A very small bubble of air is visible in the phantom (see Figure 9f), despite care taken to fill the cavity as full as possible.

The phantom described earlier was scanned on a Revolution CT (General Electric Healthcare, Chicago, IL) scanner. The scans were done in wide-cone axial mode with 160-mm z coverage acquired in a small fraction of a second (the gantry rotation time was set to 0.28 s). Figure 9a shows a reformatted slice of a reconstruction of the phantom from a static (no motion) scan. The horizontal (z) axis corresponds to time and certain temporal features of the dyssynchronous contraction are labeled. The phantom was also scanned in dynamic mode with a constant linear z-motion at 30 mm/s and a gantry rotation rate of 0.28 s. The dynamic data was reconstructed at three overlapping phases with a temporal spacing of 62 ms. The second panel (b) shows

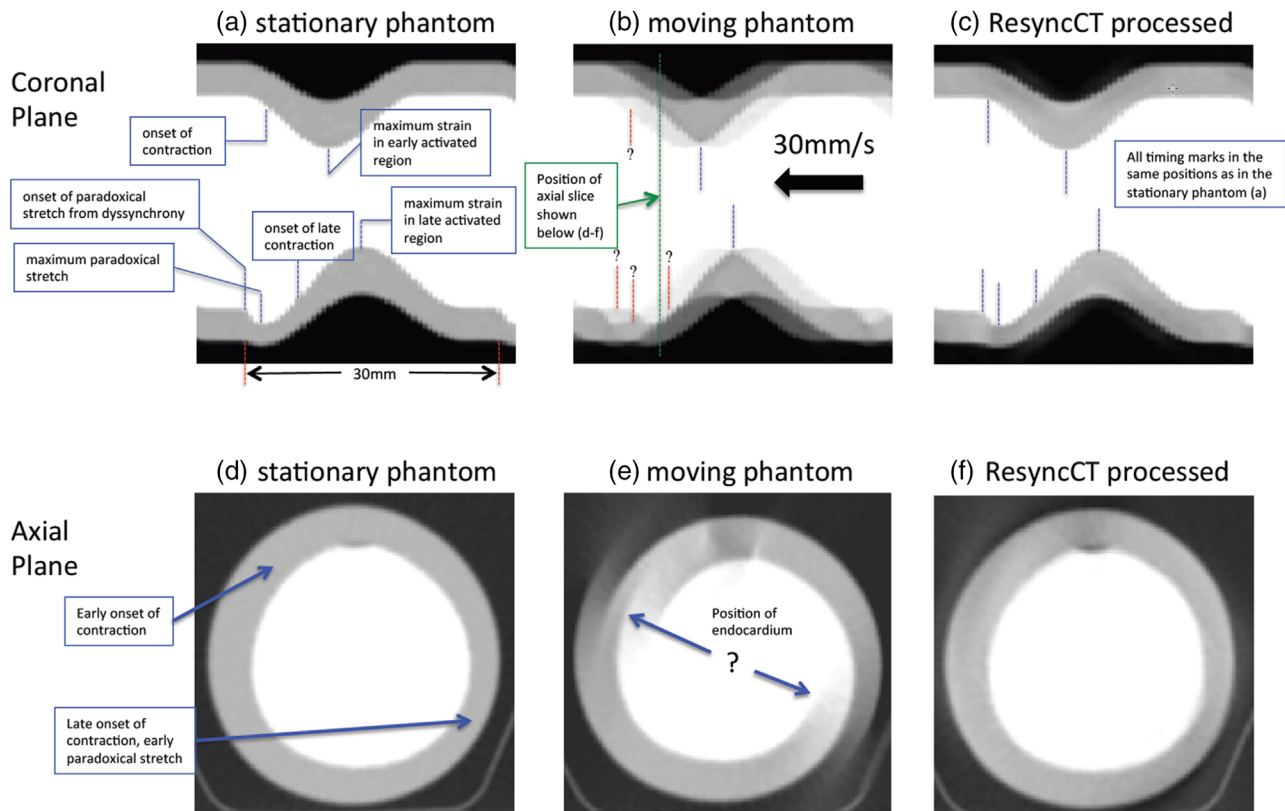


FIGURE 9 A 3D printed phantom was designed to emulate a beating left ventricle in a central 2D slice when the phantom is rigidly translated along its axis. Reconstructions of two scans of this phantom are mentioned previously: one without motion (left) and the other with motion (center). Sagittal images are shown on the top (a–c) and axial images on the bottom (d–f). The x-axis on the top row images (scanner z-axis) is the representative of time. It is difficult to identify key contraction features from the reconstruction of the moving phantom. After ResyncCT processing, however, the image quality is largely restored (right column) and these features are easily discernible

the central phase reconstruction from the dynamic data. In this image, an identification of the key temporal landmarks that were clearly captured in the static image was very difficult due to poor image quality from the motion artifacts. Next, ResyncCT motion artifact reduction was performed on the dynamic data. The resulting reformatted image shows restoration of image quality to the point that the temporal landmarks identified in panel (a) are once again readily identifiable. The improvement relative to the original motion corrupted image is notable. An axial slice through this phantom for all three cases is also shown in panels (d)–(f). Again, the morphology of the stationary image is recovered by the ResyncCT processing. In addition, the morphology of the motion-corrupted axial image is (as expected) reminiscent of the motion artifact simulated images of the contracting disk from Figures 1 and 3.

A quantitative analysis of the LV endocardial boundary was also performed. First, a finely triangulated endocardial surface mesh was generated from each of the three reconstructed volumes (ground truth, uncorrected, and ResyncCT) using the marching cubes algorithm (isosurface in Matlab). Next, for each point on the motion-corrupted mesh, the distance of the nearest point on the ground truth mesh was recorded.

These distance errors are given in histogram form in Figure 10(left). This process was repeated with the ResyncCT mesh instead of the uncorrected mesh, with the resulting (much improved) histogram being given in Figure 10(right).

In addition to the phantom experiment previously, the ResyncCT algorithm has been tested on clinical cardiac scans. In each case, reconstruction volumes were generated for every 90 degrees of gantry rotation of an entire heart cycle. This allows the local motion of the LV wall to be visualized at any point in a view similar to M-mode echocardiography, wherein each volume is resampled on a line segment that is roughly perpendicular to the endocardial boundary, and these resampled 1D functions are stacked side-by-side to form an image, as shown in Figure 11a. The y-axis on this image is spatial position along the sampled line and extends from a point that is always inside the blood pool to a point that is always outside the blood pool. The x-axis is time and is sampled every 70 ms; that is the time for the gantry to rotate 90 degrees. Artifact-free “M-mode” images show the position of the endocardial boundary through time with a smooth, narrow edge, whereas motion artifacts can cause distortions and make the LV boundary appear diffuse and the motion appear jagged or uneven. The

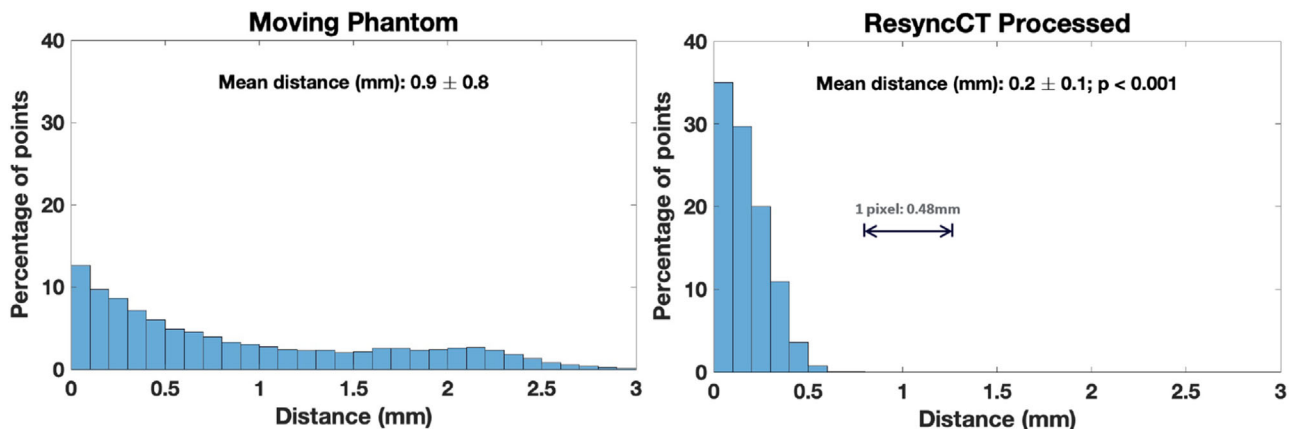


FIGURE 10 Comparison between two histograms of endocardial surface distance errors: (left) the standard uncorrected reconstruction from dynamic data; (right) the ResyncCT reconstruction from the same data. After ResyncCT processing, all errors are within ~ 1 pixel

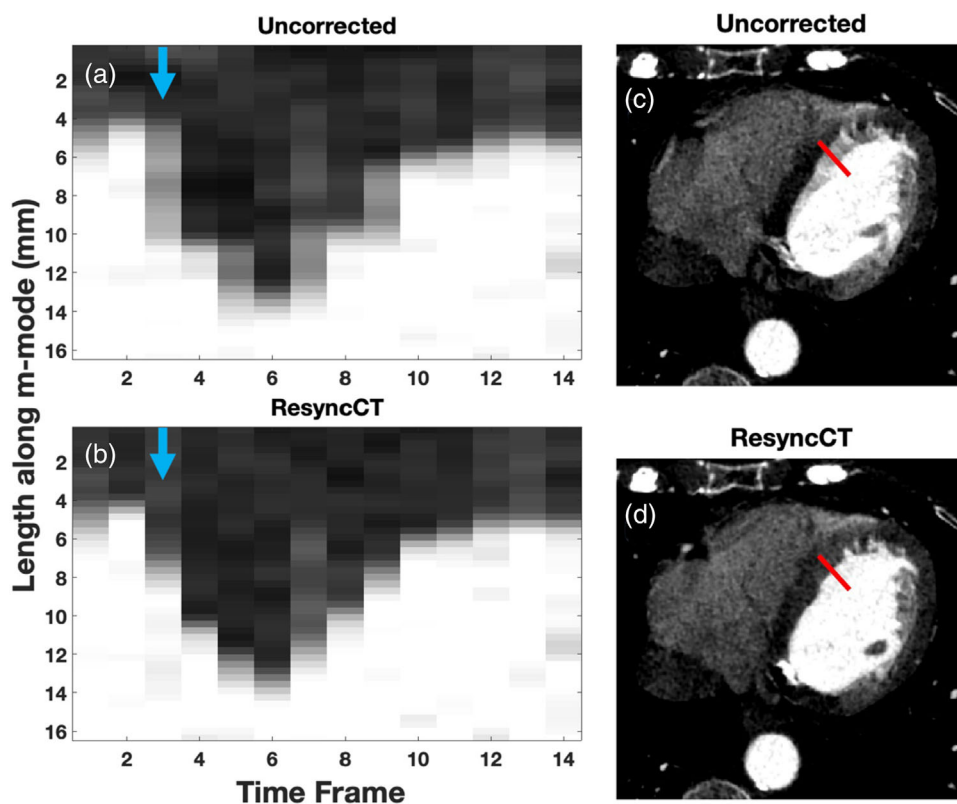


FIGURE 11 A clinical case showing the impact of motion artifact reduction on improving the clarity of the left ventricle endocardial boundary. The left images (a) and (b) are similar to “M-mode” echo images with each column showing a spatial profile at a given time in the heart cycle. The right images (c) and (d) show the location of the spatial profile in an axial slice. Images on the top are uncorrected, whereas images on the bottom are corrected with ResyncCT. A severe “double-wall” artifact is visible on the top right at the position of the M-mode profile (red line). This artifact is shown in the M-mode plots as a blurred boundary between myocardium and blood pool (blue arrow)

ResyncCT algorithm removes much of the motion artifacts. Detailed results on three illustrative cases are given in the following, with scan parameters given in Table 1.

Figure 11 shows a case in which there is a severe “double-wall” artifact at the septal wall in selected time frames. Double-wall artifacts occur when the rays at

the central projection angle of the reconstruction are nearly perpendicular to the wall boundary. For these time frames, the measured rays near the beginning and the end of the half-scan are parallel to the LV boundary and, thus, contribute strongly to the boundary localization; however, these views are temporally delayed from one another. The uncorrected M-mode plot on the left

TABLE 1 Scan details for the clinical cases shown in Figures 11–13

Figure	Protocol	HR	Kernel	Rot (s)	DFOV	mA	kVp
11	CTA cardiac morphology, PV	73	Std	0.28	178	460	80
12	CTA cardiac morphology, TAVR	53	Std	0.28	170	595	80
13	CCTA Coronary	62	Std	0.28	200	719	100

Abbreviations: CCTA, coronary computed tomography angiography. DFOV, Display field of view. PV, pulmonary veins. TAVR, Trans-Aortic Valve Replacement.

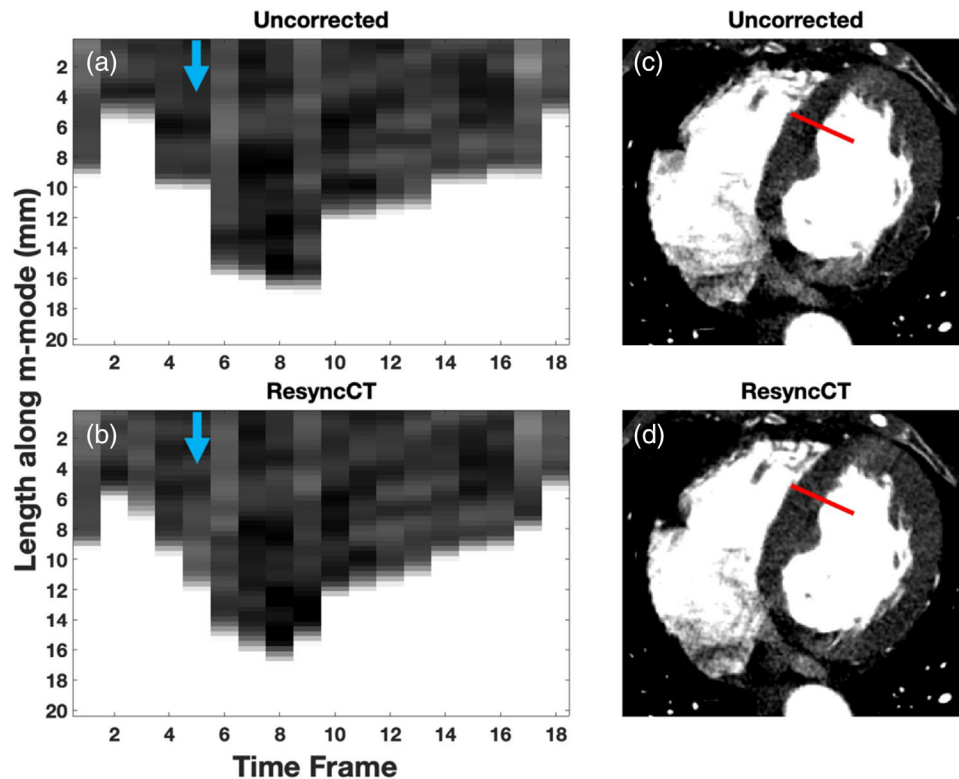


FIGURE 12 A clinical case showing distinct “stationary-wall” artifacts in the motion of the endocardial wall. Parts (a) and (b) show the M-mode data from the septal LV endocardium. Note that in (a) the wall appears to remain stationary for every two time frames; after ResyncCT correction, the motion has a smoother progression as shown in (b). The apparent motion from (a) is stair-step like as rays tangent to the surface are only measured twice per gantry rotation. LV, left ventricle

(Figure 11a) shows an indistinct boundary between the contrast-enhanced blood and the septal wall, particularly at odd-numbered time frames, which in this case were oriented as described earlier. These poorly defined boundaries make it difficult to precisely determine the activation time of wall motion. Axial images showing the location of the M-mode profile are shown in Figure 11c,d for the third time point. There is a severe double-wall motion artifact in this image (Figure 11c) that is corrected by ResyncCT in Figure 11d. Other severe motion artifacts in Figure 11c such as the blurring of the right coronary artery have also been corrected, leading to a clearer, higher fidelity image.

Figure 12 illustrates a clinical case in which the boundaries are sharp but “stationary-wall” artifacts corrupt the data. In this case, the temporal sampling happens to be phase shifted by about 35 ms relative to the sampling in the clinical case described in Figure 11. The

M-mode images prior to ResyncCT processing show a distinctive stair-step pattern in which the wall moves every other time frame, which is obviously an artifact. The apparent position of the septal wall is once again mostly determined by its position at the time that X-ray paths passed tangent to the boundary. In this case, adjacent time frames, (2,3), (4,5), (6,7), ..., are reconstructed from projection views that include the same views that sample the edge of the septal wall at a single time point. By estimating the true motion and correcting for it, the ResyncCT method is able to restore the smooth motion profile shown in Figure 12b. As in the previous example of Figure 11, the axial images in Figure 12c,d also show a reduction in other motion artifacts.

Finally, in Figure 13, a third clinical example is shown. In this case, the phase offset of the sampling is at a point that is intermediate (between that shown in the previous two figures). Here there is evidence of both

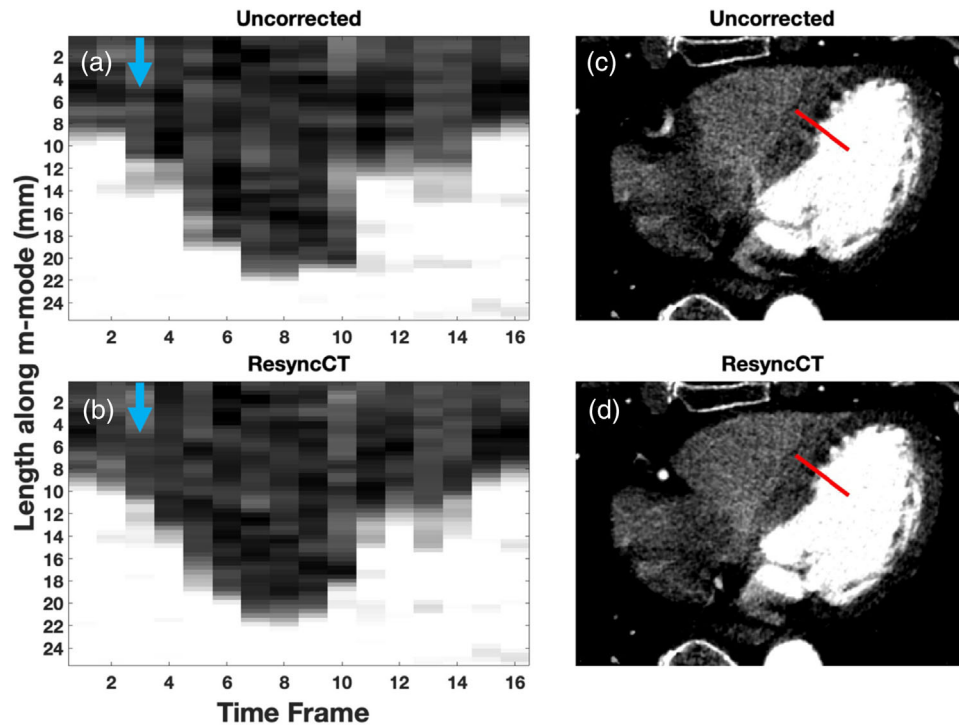


FIGURE 13 (a–d) A clinical case in which both “double-wall” and “stationary-wall” artifacts are present. Image layout matches that of Figures 11 and 12

double- and stationary-wall artifacts in the uncorrected M-mode image. As before, the axial image also shows a clear reduction of motion artifacts after ResyncCT processing. It should be noted that the phase offset sampling referred to in the discussion earlier (and the resultant artifact characteristics such as the appearance of double- and stationary walls) is dependent on the orientation of the surface and the orientation of the gantry during each scan. As a result, if one were to create an M-mode image at another wall location with a different relative orientation for any of the previous clinical cases, one should expect to see the artifact change. The key to controlling the artifacts is estimating and correcting for the motion over the entire field as is done by the ResyncCT method.

Although the previous figures are illustrative, the improvement in image quality from ResyncCT for clinical cases is, perhaps, most easily appreciated when the volumes are viewed dynamically in a cine-loop. When compared to the uncorrected images, the ResyncCT images show much more even and consistent LV contraction with a dramatic reduction in spiraling/swirling gantry angle motion artifacts.

5 | DISCUSSIONS AND FUTURE WORK

There are two key advantages to finding correspondences in conjugate pairs of images rather than trying to track individual features through a temporal series of

PAR images. First, conjugate pairs of PAR images share the same spatial frequencies. It is very difficult to try to track any but the simplest features (e.g., small, compact, high-contrast objects like a localized calcification) with a series of PAR images that do not have conjugate redundancy. For example, a cylindrical in-plane object (e.g., a vessel) will nearly disappear in PAR images that do not contain the most relevant spatial frequencies (those corresponding to rays that are parallel to the cylinder axis), thereby making motion estimation very difficult. As a result, motion artifact reduction is most robust when the measured scan data extends beyond 180 degrees. In fact, if only 180 degrees of parallel data are available, there is no conjugate information at all, the reconstructed image is consistent with the measured data (except very near DC), and there is no reliable, objective evidence that the true object was moving at all. Any motion estimates, in this case, must be based on a priori assumptions about the characteristics of the image. Having redundant (conjugate) data is, therefore, important for robust motion estimation.

The second key advantage to using correspondences in conjugate-pair images is that such correspondences are robust to position distortions that occur in PAR images due to (at least) linear motion. To illustrate this point, consider Figure 14: three source positions (S_n) are shown at the top of the scan along with the corresponding positions of a small object (O_n) in the image. Backprojecting the data corresponding to the three connecting rays will produce an estimate of the location of

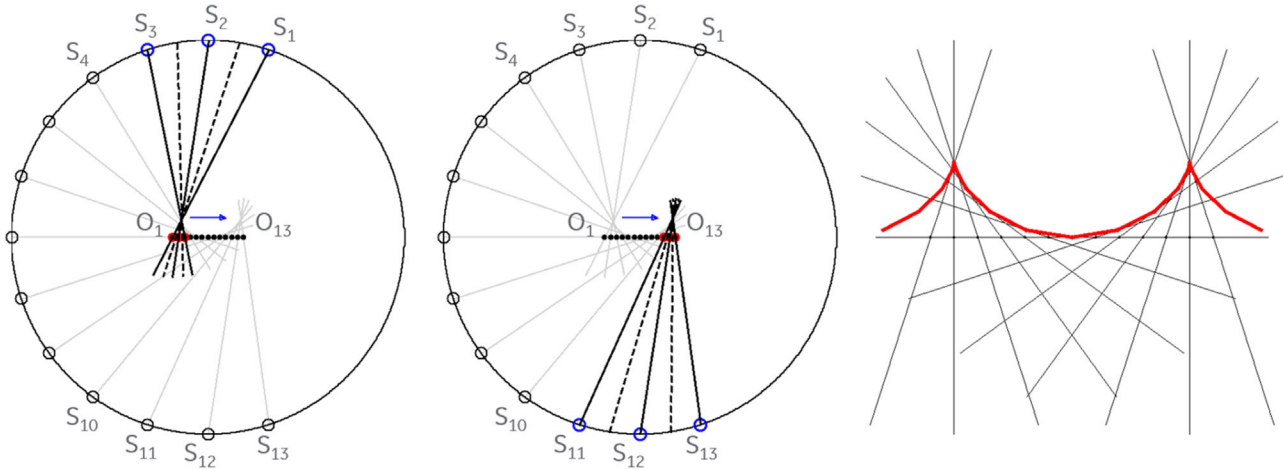


FIGURE 14 Small objects moving with a constant velocity appear to move on a cycloid path in CT images. Fortunately, cross-correlations from conjugate-PAR images are immune to these spatial distortions because they are periodic with a 180-degree gantry rotation. CT, computed tomography; PAR, partial angle reconstruction

the small object that does not match with reality. As this does not occur for a static object, one might assume that cutting the temporal window between S_1 and S_3 by a factor of 2 might reduce the magnitude of this distortion. Unfortunately, this is not the case as can be seen by the convergence of the two intermediate dashed lines in the same (distorted) location. In fact, no matter how small you make your PAR window, the distortion does not change as its magnitude is related only to the motion magnitude, the gantry rotation rate, and the sine of the angle between the X-ray paths and the motion direction. Fortunately, as the sine function is pi-periodic, the distortion is the same in the conjugate-PAR image (see Figure 14(center)) except for any changes due to variations in the motion direction or magnitude. Tracking the position of the small object moving linearly over a scan results in a cycloid path rather than the true linear path (see Figure 14(right)). Fortunately, taking the difference between points that are 180 degrees apart on a cycloid always yields the same velocity estimate. It turns out that such positional distortions are not much of a problem when it comes to motion compensation. This is because the difference between the true (linear) path and the apparent (cycloid) path is completely constrained to the component of motion that is parallel to the ray path. Errors in motion estimates along this direction are largely irrelevant as the PAR images are nearly invariant in that direction. Understanding the distortion effect, however, can be relevant for improving motion estimation models.

In this paper, a detailed description of the ResyncCT motion artifact reduction algorithm was given and placed in context with the many recent scientific advances in the field of motion artifact reduction for CT. In addition, the method was discussed in the context of correcting motion around the LV, with the goal of reducing “false dyssynchrony” and improving the quality of information

leading up to clinical decisions about whether and how to perform CRT. This represents an exciting clinical application of 4DCT. Unfortunately, 4DCT is not recognized generally for use in myocardial function and there are some technical challenges that need to be addressed to make this capability and other 4DCT applications more broadly available. In particular, most large, heavily used PACS systems do not currently support cine capability, meaning that the review of 4DCT datasets requires a workstation with additional review software.

Though beyond the scope of this paper, additional experiments have been performed to quantify the improvement in the accuracy of timing measurements that results from the application of ResyncCT using a 3D printed LV phantom.¹ In addition, plans are in place to quantify the ability of 4DCT with motion artifact reduction to better select patients for CRT procedures. Finally, additional work on the topic of accelerating motion artifact reduction by leveraging deep convolutional neural networks is also in progress.

ACKNOWLEDGMENTS

Research reported in this publication was supported by NHLBI of the National Institutes of Health under award number R01 HL144678 (McVeigh). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

This work was done during the term of an Award from the American Heart Association (AHA 20PRE35210261).

CONFLICT OF INTEREST

Dr. McVeigh holds founder shares in ClearPoint Neuro, Inc. and receives research funding from GE Healthcare, Abbott Medical, and Pacesetter Inc.

REFERENCES

1. Manohar A, Pack JD, Schluchter AJ, McVeigh ER. Four-dimensional computed tomography of the left ventricle, Part II: Estimation of mechanical activation times. *Med Phys*. 2022;49(4):2309-2323. Under Concurrent Review Medical Physics.
2. Vernooy K, Van Deursen CJM, Strik M, Prinzen FW. Strategies to improve cardiac resynchronization therapy. *Nat Rev Cardiol*. 2014;11(8):481-493. doi:10.1038/nrcardio.2014.67
3. Mirea O, Pagourelis ED, Duchenne J, et al. Variability and reproducibility of segmental longitudinal strain measurement: a report from the EACVI-ASE Strain Standardization Task Force. *JACC Cardiovasc Imaging*. 2018;11(1):15-24.
4. Daubert J-C, Saxon L, Adamson PB, et al. 2012 EHRA/HRS expert consensus statement on cardiac resynchronization therapy in heart failure: implant and follow-up recommendations and management. *Heart Rhythm (United States)*. 2012;9(9):1524-1576.
5. Pourmorteza A, Chen MY, van der Pals J, Arai AE, McVeigh ER. Correlation of CT-based regional cardiac function (SQUEEZ) with myocardial strain calculated from tagged MRI: an experimental study. *Int J Cardiovasc Imaging*. 2016;32(5):817-823.
6. Kidoh M, Shen Z, Suzuki Y, et al. False dyssynchrony: problem with image-based cardiac functional analysis using x-ray computed tomography. *Prog Biomed Opt Imaging – Proc SPIE*. 2017;10132:1-7.
7. Kidoh M, Shen Z, Suzuki Y, et al. False dyssynchrony: problem with image-based cardiac functional analysis using x-ray computed tomography. *Proc. SPIE 10132, Medical Imaging 2017: Physics of Medical Imaging*, 101321U. 9 March 2017. doi:10.1117/12.2250257
8. Schäfer D, Borgert J, Rasche V, Grass M. Motion-compensated and gated cone beam filtered back-projection for 3-D rotational X-ray angiography. *IEEE Trans Med Imaging*. 2006;25(7):898-906. doi:10.1109/tmi.2006.876147
9. Ritchie CJ, Crawford CR, Godwin JD, King KF, Kim Y. Correction of computed tomography motion artifacts using pixel-specific back-projection. *IEEE Trans Med Imaging*. 1996;15(3):333-342. doi:10.1109/42.500142
10. Iatrou M, Pack JD, Bhagalia R, Beque D, Seamans J. Coronary artery motion estimation and compensation: a feasibility study. *IEEE Nuclear Science Symposium & Medical Imaging Conference*. IEEE; November 2010:2819-2821.
11. Bhagalia R, Pack JD, Miller JV, Iatrou M. Nonrigid registration-based coronary artery motion correction for cardiac computed tomography. *Med Phys*. 2012;39(7 pt 1):4245-4254.
12. Pack JD, Claus B. An analysis of motion artifacts in CT and implications for motion compensation. *Proceedings of the 2nd CT Meeting*. 2012:322. <https://www.ct-meeting.org/data/ProceedingsCTMeeting2012.pdf>
13. Brendel B, Bippus R, Kabus S, Grass M. Motion compensated back-projection versus back-project-then-warp for motion compensated reconstruction. *Proceedings of the 3rd CT Meeting*. 2014:169. <https://www.ct-meeting.org/data/ProceedingsCTMeeting2014.pdf>
14. Tang Q, Matthews J, Razeto M, Linde JJ, Nakanishi S. Motion estimation and compensation for coronary artery and myocardium in cardiac CT. In: *Medical Imaging 2015: Physics of Medical Imaging*. International Society for Optics and Photonics; Mar 2015:94120Q.
15. Pack JD, Edic PM, Claus BEH, Iatrou M, Miller JV. Method for computed tomography motion estimation and compensation. U.S. patent no. 8224056. Washington, DC: U.S. Patent and Trademark Office. 2012.
16. Lossau T, Nickisch H, Wissel T, et al. Motion estimation and correction in cardiac CT angiography images using convolutional neural networks. *Comput Med Imaging Graph*. 2019;76:101640.
17. Tang Q, Cammin J, Srivastava S, Taguchi K. A fully four-dimensional, iterative motion estimation and compensation method for cardiac CT. *Med Phys*. 2012;39(7):4291-4305. doi:10.1118/1.4725754
18. Rohkohl C, Bruder H, Stierstorfer K, Flohr T. Improving best-phase image quality in cardiac CT by motion correction with MAM optimization. *Med Phys*. 2013;40(3):031901.
19. Hahn J, Bruder H, Rohkohl C, et al. Motion compensation in the region of the coronary arteries based on partial angle reconstructions from short-scan CT data. *Med Phys*. 2017;44(11):5795-5813. doi:10.1002/mp.12514
20. Fung GSK, Ciuffo L, Ashikaga H, Taguchi K. Motion estimation for cardiac functional analysis using two x-ray computed tomography scans. *Med Phys*. 2017;44(9):4677-4686.
21. Holmes T, Ma K, Pourmorteza A. Combination of CT motion simulation and deep convolutional neural networks with transfer learning to recover Agatston scores, The 15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine. 2019. doi:10.1117/12.2534882
22. Maier J, Lebedev S, Erath J, et al. Deep learning-based coronary artery motion estimation and compensation for short-scan cardiac CT. *Med Phys*. 2021;48:3559-3571. doi:10.1002/mp.14927
23. Kim S, Chang Y, Ra JB. Cardiac motion correction based on partial angle reconstructed images in x-ray CT. *Med Phys*. 2015;42(5):2560-2571. doi:10.1118/1.4918580
24. Okerlund D, Pack JD, Jackson J. *SnapShot Freeze Motion Correction in Coronary CT Angiography*. Dec. 2011. A GE Healthcare white paper.
25. Leipsic J, Labounty TM, Hague CJ, et al. Effect of a novel vendor-specific motion-correction algorithm on image quality and diagnostic accuracy in persons undergoing coronary CT angiography without rate-control medications. *J Cardiovasc Comput Tomogr*. 2012;6(3):164-171.
26. Liu F, Cuevas C, Moss AA, Kolokythas O, Dubinsky TJ, Kinahan PE. Gas bubble motion artifact in MDCT. *Am J Roentgenol*. 2008;190(2):294-299.
27. Contijoch F, Stayman JW, McVeigh ER. The impact of small motion on the visualization of coronary vessels and lesions in cardiac CT: a simulation study. *Med Phys*. 2017;44(7):3512-3524. doi:10.1002/mp.12295
28. Kruskal J. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc Am Math Soc*. 1956;7:48-50.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Pack JD, Manohar A, Ramani S, et al. Four-dimensional computed tomography of the left ventricle, Part I: Motion artifact reduction. *Med Phys*. 2022;1-15. <https://doi.org/10.1002/mp.15709>