IOP Publishing | Institute of Physics and Engineering in Medicine Physics in Medicine & Biology

Phys. Med. Biol. **61** (2016) 6856–6877 doi:10.1088/0031-9155/61/18/6856

Motion-aware temporal regularization for improved 4D cone-beam computed tomography

Cyril Mory[1,2], Guillaume Janssens[3] and Simon Rit[2]

[1] iMagX Project, ICTEAM Institute, Université Catholique de Louvain, Louvain-la-Neuve, Belgium
[2] Université de Lyon, CREATIS; CNRS UMR5220; Inserm U1044; INSA-Lyon; Université Lyon 1; Centre Léon Bérard, France
[3] Ion Beam Applications SA, Louvain-La-Neuve, Belgium

E-mail: simon.rit@creatis.insa-lyon.fr

Received 23 February 2016
Accepted for publication 1 July 2016
Published 2 September 2016

Abstract

Four-dimensional cone-beam computed tomography (4D-CBCT) of the free-breathing thorax is a valuable tool in image-guided radiation therapy of the thorax and the upper abdomen. It allows the determination of the position of a tumor throughout the breathing cycle, while only its mean position can be extracted from three-dimensional CBCT. The classical approaches are not fully satisfactory: respiration-correlated methods allow one to accurately locate high-contrast structures in any frame, but contain strong streak artifacts unless the acquisition is significantly slowed down. Motion-compensated methods can yield streak-free, but static, reconstructions. This work proposes a 4D-CBCT method that can be seen as a trade-off between respiration-correlated and motion-compensated reconstruction. It builds upon the existing reconstruction using spatial and temporal regularization (ROOSTER) and is called motion-aware ROOSTER (MA-ROOSTER). It performs temporal regularization along curved trajectories, following the motion estimated on a prior 4D CT scan. MA-ROOSTER does not involve motion-compensated forward and back projections: the input motion is used only during temporal regularization. MA-ROOSTER is compared to ROOSTER, motion-compensated Feldkamp–Davis–Kress (MC-FDK), and two respiration-correlated methods, on CBCT acquisitions of one physical phantom and two patients. It yields streak-free reconstructions, visually similar to MC-FDK, and robust information on tumor location throughout the breathing cycle. MA-ROOSTER also allows a variation of the lung tissue density during the breathing cycle, similar to that of planning CT, which is required for quantitative post-processing.

Keywords: tomography, cone-beam, radiotherapy, 4D, motion, thorax

S Online supplementary data available from stacks.iop.org/PMB/61/6856/mmedia

(Some figures may appear in colour only in the online journal)

1. Introduction

State-of-the-art radiotherapy strategies for the treatment of thoracic and upper-abdominal tumors take the patient's breathing motion into account. The clinical workflow is currently the following: the patient first undergoes a four-dimensional (4D) computed tomography (CT) scan, from which doctors determine the treatment plan (Wolthaus *et al* 2008). This plan takes into account, among many parameters, the patient's breathing motion pattern. Treatment is then delivered with a radiotherapy device that can combine a high-energy x-ray beam for tumor treatment and cone-beam CT (CBCT) for pre-treatment imaging, both mounted on the same gantry (Jaffray *et al* 2002). On the treatment day, the patient lies down on the table of the radiotherapy device and a CBCT acquisition can be performed. The CBCT image is used to re-position the table, so as to match the patient position as closely as possible with that of the planning CT. When the CBCT image is reconstructed in three dimensions (3D), it is implicitly assumed that the patient's breathing motion does not change much throughout the treatment and remains close to what it was during the planning 4D CT. Unfortunately, this assumption can be wrong, e.g. if a large tumor shrinks under the effect of radiation, partly restoring the patient's respiratory function. Replacing the 3D with a 4D reconstruction from the same data would allow the clinicians to check whether the patient's breathing motion on the treatment day matches that of the planning CT. If the motions do not match, corrective actions, e.g. re-planning, could be taken, hence improving the radiotherapy.

The methods currently available to reconstruct a CBCT acquisition of a moving object can roughly be divided into four classes:

- Respiration-correlated reconstruction techniques, which reconstruct one 3D frame at a time and concatenate the results to obtain a 4D reconstruction. These techniques include the respiration-correlated versions of the Feldkamp–Davis–Kress (FDK) (Feldkamp *et al* 1984, Sonke *et al* 2005) and simultaneous algebraic reconstruction (Andersen and Kak 1984, Mory *et al* 2014) techniques, as well as 3D regularized reconstruction methods (Leng *et al* 2008, Sidky and Pan 2008, Bergner *et al* 2010). 4D reconstruction techniques that do not perform regularization along time also fall into this category. These methods allow one to accurately locate high-contrast structures throughout the breathing cycle, but each frame has low image quality, due to either streak artifacts, blurring or over-regularization, unless the scanner is slowed down to improve the sampling of cone-beam projections in each 3D frame (Sonke *et al* 2005).
- Classical motion-compensated reconstruction techniques, which use an *a priori* motion estimation (Li *et al* 2007, Rit *et al* 2009a, 2009b, 2011) to back-project along curved trajectories. These methods reconstruct one frame from all projections and yield a static reconstruction, which is only as good as the motion estimation used in the input.
- Joint motion-estimation and motion-compensated reconstruction methods, which estimate the motion directly from the CBCT data (Brehm *et al* 2012, Wang and Gu 2013a, 2013b, Liu *et al* 2015) and perform a motion-compensated reconstruction. The 3D static reconstruction obtained can then be animated with the estimated CBCT motion. These approaches are valid alternatives to the one we propose, although they share a limitation

with the classical motion-compensated reconstruction techniques: the variation of the linear attenuation of lung tissue along the breathing cycle cannot be estimated from such reconstructions.

- Regularized 4D reconstruction techniques, which reconstruct the whole cycle at once, making use of all the projection data, and enforce some similarity between consecutive frames by regularizing along time (Jia *et al* 2010, Ritschl *et al* 2012, Wu *et al* 2012, Mory *et al* 2014).

The proposed method, is halfway between the second and the fourth category. It builds upon the existing reconstruction using spatial and temporal regularization (ROOSTER) and is called motion-aware ROOSTER (MA-ROOSTER). The only change with respect to ROOSTER is that temporal regularization is performed along curved trajectories, following an *a priori* motion estimation computed from the planning 4D CT scan, instead of straight trajectories. In this work, we compare MA-ROOSTER to ROOSTER, to illustrate the benefits of motion-aware over straight regularization, and MA-ROOSTER to motion-compensated FDK (MC-FDK), both with an accurate and an inaccurate *a priori* motion estimation, to prove that MA-ROOSTER indeed shows some robustness to erroneous input motion, while MC-FDK does not. On phantom data, we evaluate the effect of the temporal regularization parameter on the robustness to erroneous input motion, using respiration-correlated reconstructions as references. Then we show reconstruction results with ROOSTER, MA-ROOSTER and MC-FDK on two clinical datasets. Finally, we compare the variations of lung tissue attenuation during the breathing cycle in a MA-ROOSTER reconstruction and in 4D planning CT, and show that they are consistent with each other.

2. Materials and methods

Throughout the paper, we shall use the word ‘frame’ to denote a 3D volume of a 4D sequence. In other contributions, such 3D volumes are often referred to as ‘phases’, but the ‘phase’ also denotes the real number in [0;1], defined as the relative position between two consecutive end-inhales. For simplicity, we shall also refer to ‘the frame representing the patient’s body at phase 50%’ as ‘frame 50%’.

2.1. Respiration-correlated FDK

A 4D reconstruction can be obtained by concatenating 3D respiration-correlated FDK reconstructions (Sonke *et al* 2005, Lu *et al* 2007, Bergner *et al* 2010). The 4D reconstruction contains severe streak artifacts but can serve as a reference to estimate the motion of a high-contrast object like a tumor. Throughout this paper, it is either called respiration-correlated FDK or simply 4D FDK.

2.2. Motion-compensated FDK

The most straightforward approach when a 4D displacement vector field (DVF) is available is to compute a 3D motion-compensated FDK (MC-FDK), performed in this paper as described in (Rit *et al* 2009a, 2009b). It belongs to the ‘classical motion-compensated reconstruction techniques’ described in the introduction. The DVF extracted from the planning CT allows one to warp all frames of the respiratory cycle to the end-exhale frame. By performing an MC-FDK using this DVF, we obtain a motion-compensated 3D reconstruction of the end-exhale frame. The more accurate the DVF, the sharper and better contrasted the reconstruction.

MC-FDK is a 3D reconstruction technique: it reconstructs a single frame, and does not provide any information on the way the tumor actually moves. Therefore, it cannot be used to check whether the patient's breathing motion of the treatment day matches that of the planning CT. We use it in this paper only as a reference for visual comparison to evaluate image quality. Note that with a null DVF, MC-FDK boils down to the blurry static FDK.

2.3. 4D conjugate gradient

The 4D conjugate gradient (CG) method reconstructs a sequence of volumes from a single stack of projections through a convex optimization approach. It consists in minimizing the single-term cost function $\sum_\alpha \left\| R_\alpha S_\alpha f - p_\alpha \right\|_2^2$ by the linear CG algorithm, where

- $\|.\|_2$ is the $\ell 2$ norm.
- α is the projection index.
- f is a 4D sequence of volumes.
- R_α is the forward projection operator at index α. It maps a 3D frame onto a 2D projection image.
- S_α is an interpolator along the time dimension. In this study, S is a linear interpolator. It maps a 3D + time sequence onto a 3D frame, using the respiratory phase of projection α. For example, if f contains ten frames ($f_0, f_1, f_2 ...$) and projection α has been acquired at phase 0.87, then $S_\alpha f = 0.3 f_8 + 0.7 f_9$.
- p_α is the measured projection with index α.

The 4D CG reconstructions are blurry and contain streak artifacts, but provide reliable information on the motion of objects that are not hidden by the streaks, e.g. large tumors or high-contrast structures. The ROOSTER method builds upon 4D CG, adding some regularization steps.

2.4. The ROOSTER method

ROOSTER is a recent iterative reconstruction method alternating between several optimization goals (Mory *et al* 2014). It assumes that a motion mask, i.e. a rough segmentation of the region where movement is expected to occur, is available. As motion can occur outside the lungs, since the rib cage and the abdomen move during breathing, we used the motion mask extracted from the 4D planning CT (Vandemeulebroucke *et al* 2012), dilated by morphological operations in order to include the ribs. ROOSTER consists in solving the following five subproblems at each iteration of the main loop:

- Minimizing the data-attachment term $\sum_\alpha \left\| R_\alpha S_\alpha f - p_\alpha \right\|_2^2$, by 4D CG.
- Enforcing positivity, by setting all negative voxels of f to zero.
- Removing motion where it is not expected to occur, by averaging along time outside the motion mask.
- Enforcing the spatial gradient's sparsity in each frame using 3D total variation (TV) denoising.
- Enforcing the temporal gradient's sparsity for each spatial position, by one-dimensional (1D) TV denoising along time.

Each supbroblem's output is used as the input for the next subproblem. This constitutes one iteration of the main loop, the output of which is fed back to the CG minimizer for the next iteration.

TV denoising is here intended in its convex optimization sense, i.e. applying TV denoising on f_{noisy} yields $f_{\text{denoised}} = \arg\min_u \left\|u - f_{\text{noisy}}\right\|_2^2 + \gamma TV(u)$, where parameter γ controls the strength of the TV regularization (the higher, the stronger the regularization). In convex optimization literature, TV denoising is also referred to as the 'proximal operator' of TV (Boyd and Ye 2011). ROOSTER makes use of 3D TV for spatial denoising of each frame (with strength parameter γ_{space} and of 1D temporal TV for temporal denoising between frames (with strength parameter γ_{time}). They are defined as follows:

$$TV_{\text{space}}(u) = \|\sqrt{(\nabla_x u)^2 + (\nabla_y u)^2 + (\nabla_z u)^2}\,\|_1 \tag{1}$$

and similarly

$$TV_{\text{time}}(u) = \|\nabla_t u\|_1 \tag{2}$$

where $\nabla_x, \nabla_y, \nabla_z$ and ∇_t are the gradient operators along the spatial axes x, y, z and along the time axis t, respectively, and $\|.\|_1$ is the $\ell 1$ norm. In both spatial 3D and temporal 1D TV denoising, the minimization is performed by the basis pursuit dequantization algorithm (Jacques *et al* 2010).

2.5. Motion-active ROOSTER

Through a breathing cycle, a given spatial location can contain tissues within a large range of linear attenuation coefficients, e.g. lung tissues and blood vessels. In this case, regularization along time as performed in ROOSTER enforces similarity between unrelated objects. As a result, it can smooth away the small moving structures and the high-intensity structures of one frame can partially spread out to the previous and next frames. The diaphragm, small structures in the lungs and lung tumors seem to gradually fade from one position to another, while they should have sharp edges and a distinct position on every frame. This effect can partly be explained by noting that, while TV is often said to favor piecewise-constant functions, the 1D TV of a monotonic function is exactly the same as the 1D TV of a step function with the same lower and upper bounds. 1D TV therefore favors piecewise-monotonic functions, not only piecewise-constant functions. When used to regularize along time, it can cause the above-described blurring effect between consecutive frames. To tackle this issue, one could try to find a better regularizer along time than 1D TV, e.g. based on the $\ell 0$ norm of the temporal gradient. In the MA-ROOSTER algorithm, we have chosen to circumvent the problem by bending the regularization trajectories, so as to follow the moving structures.

As explained in the introduction, lung cancer patients undergo a 4D CT at the beginning of the radiotherapy, on which their treatment is planned. From the 4D planning CT, a 4D DVF can be estimated, which maps the end-exhale frame to each other frame. The method by which the DVF estimation is performed has little importance for MA-ROOSTER, only the DVF itself matters. Since this *a priori* motion estimation is available, we propose improving ROOSTER by performing temporal regularization along curved trajectories following the motion. This is illustrated in figure 1. Note that it is mathematically equivalent, but much easier to implement, to first warp all frames onto the end-exhale frame by 'backward mapping' (see chapter 10.2 of Moeslund (2012)), and then to apply straight regularization along time on the warped sequence. The regularized frames must then be inverse-warped to their original phase (see section 2.6). Note that trilinear interpolation-based image warping is a linear process, and can therefore be described by a matrix. To describe the MA-ROOSTER temporal regularization in a formal way, assuming that we reconstruct a sequence f of ten frames, let

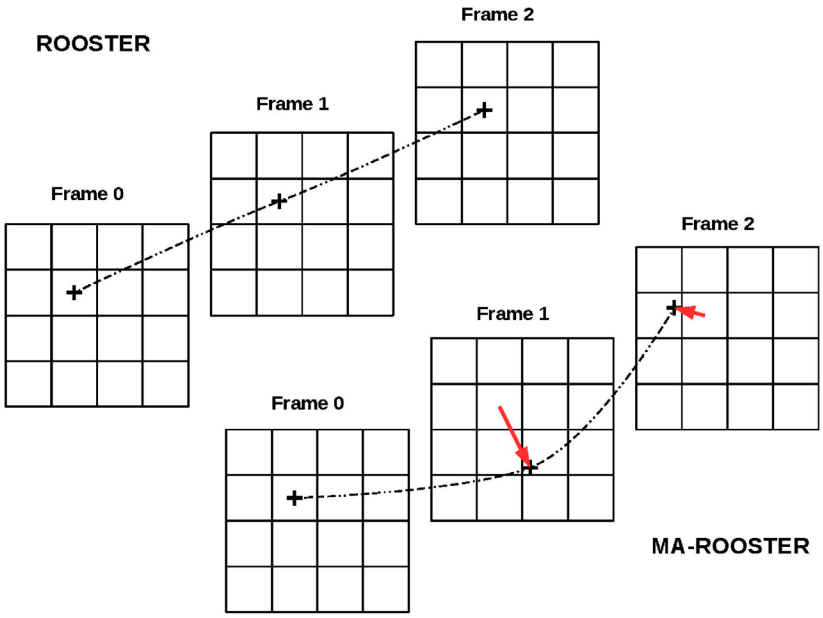

Figure 1. Illustration of the voxel selection for temporal regularization. Top: the behavior of ROOSTER, select in each frame the voxel that is located at a fixed spatial position. Bottom: behavior of MA-ROOSTER, follow the trajectory of the voxel, given by the input DVF. The red arrows are the vectors of the input DVF.

- N be the number of voxels in a 3D frame.
- $f_{\text{noisy}} = \begin{pmatrix} f_0 \\ \vdots \\ f_9 \end{pmatrix}$ be the sequence to regularize, with $f_j \in \mathbb{R}^N, j \in \{0..9\}$ the individual frames
- $W = \begin{pmatrix} W_0 & & 0 \\ & \ddots & \\ 0 & & W_9 \end{pmatrix}$ be the 4D warping operator, with $W_j : \mathbb{R}^N \to \mathbb{R}^N, j \in \{0..9\}$ the 3D warping operators for each frame.
- $Wf_{\text{noisy}} = \begin{pmatrix} W_0 f_0 \\ \vdots \\ W_9 f_9 \end{pmatrix}$ be the warped sequence (each 3D frame in f is warped independently).
- D_{time} be the TV denoising along time operator (note that D_{time} is not linear).
- $D_{\text{time}}(Wf_{\text{noisy}})$ be the denoised warped sequence.

Then $f_{\text{denoised}} = W^{-1} D_{\text{time}}(Wf_{\text{noisy}})$ is the denoised and inverse-warped sequence, which will be used as input for the next main loop iteration of MA-ROOSTER.

An inverse DVF can only be obtained when the DVF is diffeomorphic (Arsigny *et al* 2006). And even when a diffeomorphic DVF and its inverse are available, the associated warping operators are not the exact inverse of one another, because of interpolation errors. Since inverting W in reasonable time is impossible, we will always use the warping operator associated with the inverse DVF, which from a strict linear algebra standpoint is an approximation of W^{-1}. Nevertheless, by an abuse of notation, we will write W^{-1} to denote it. Details on how to mitigate the effects of using an imperfect W^{-1} are given in section 2.6. Note that a null DVF

results in $W = W^{-1} = I$, i.e. in temporal regularization along straight lines as in ROOSTER. ROOSTER can therefore be seen as a specific case of MA-ROOSTER with a null DVF. Intuitively, this means that MA-ROOSTER should yield better results than ROOSTER as soon as the DVF is better than null, even when it does not perfectly describe the real displacement.

The temporal regularization enforces a trade-off between the motion present in the projection data and the one described by the input DVF. It discourages deviating too far from the input motion, but does not prevent it. MA-ROOSTER is thus expected to show some robustness to motion estimation inaccuracies. The temporal regularization also discourages variations in lung tissue attenuation during the breathing cycle, but does not prevent them. The attenuation of lung tissue in MA-ROOSTER reconstructions is therefore expected to vary slowly throughout the breathing cycle, as a result of the air and blood flows in the lungs.

2.6. Inverse warping

Some motion estimation methods, including the one we used, provide two so-called 'inverse-consistent' DVFs (Christensen and Johnson 2001, Janssens *et al* 2011, Wang and Gu 2013b). Let W and W^{-1} be their associated warping operators. They are approximately the inverse of each other. In practice, $\hat{f} = W^{-1}Wf$ is a blurry approximation of f, since it has undergone two trilinear interpolations, one contained in W, the other one in W^{-1}. With a small trick, however, the blurring can be limited. We compute f_{denoised} as follows:

$$\begin{aligned} f_{\text{denoised}} &\approx W^{-1}D_{\text{time}}(Wf_{\text{noisy}}) \\ &\approx W^{-1}(D_{\text{time}}(Wf_{\text{noisy}}) - Wf_{\text{noisy}} + Wf_{\text{noisy}}) \\ &= W^{-1}(D_{\text{time}}(Wf_{\text{noisy}}) - Wf_{\text{noisy}}) + f_{\text{noisy}} \end{aligned} \tag{3}$$

In equation (3), $D_{\text{time}}(Wf_{\text{noisy}}) - Wf_{\text{noisy}}$ is the correction brought by regularization along time to Wf_{noisy}. Restricting the approximate inverse warping by W^{-1} to that correction allows one to avoid blurring f_{noisy}.

MA-ROOSTER can also handle the case where only a single DVF is available, i.e. the motion estimation method used does not provide the inverse DVF. The inverse warping in that case relies on convex optimization, and has been described in Mory and Rit (2015).

2.7. Metrics

Since the main purpose of MA-ROOSTER is to provide a 4D reconstruction on which tumor motion can be measured, we have designed a quantitative metric to measure tumor motion. A secondary objective is to yield a reconstruction that is close, in terms of 'image quality', to what doctors are used to in 3D (i.e. FDK), and which contains as few streak artifacts as possible, so as not to hamper visual interpretation. This secondary objective being both quite vague and extremely difficult to quantify with metrics, in particular on patient data (McCollough *et al* 2012), we have chosen to leave it to visual evaluation.

The motion of tumors throughout the breathing cycle was measured as follows:

- A small region of interest (ROI) containing the full trajectory of the tumor was delineated manually. The ROI is static, and in any frame the whole tumor is inside the ROI.
- Each frame was cropped to keep only the ROI. A translation-only motion was estimated between each cropped frame and the end-exhale cropped frame, using the open-source software Elastix, with the mean squared difference as a similarity measure, and a pyramid of four resolutions.

- The Euclidian norm of the obtained 3D translation vector was computed and is reported in graphs (see figures 4, 6, 8, 10 and 12).

2.8. Table subtraction

FDK can be used to reconstruct any voxel contained in the field of view, independent of the others. The reconstructed volume can therefore be a small part of the attenuating object. Optimization-based tomography methods, on the other hand, require the reconstructed volume to fully contain the object (Ziegler *et al* 2008). In other words, the reconstructed volume should be large enough so that any object appearing in the measured projections lies inside it. If this requirement is not met, overshoot appears on the borders of the reconstructed volume, which then causes streak artifacts once back-projected, and the reconstruction quickly diverges. In some radiation therapy centers, patients lie in a stereotactic body frame (SBF), which is typically much larger than them. Setting the reconstructed volume to be large enough to contain the SBF could increase the computation time and memory footprint by a factor of 4 or more, which is unacceptable for MA-ROOSTER. The SBF therefore has to be removed from the projections before reconstruction. To this end, we performed a static 3D reconstruction of the full volume (patient + SBF + table), masked out the patient, forward projected through the residual 3D volume (SBF + table) and subtracted the result from the original projections. The corrected projections were then used for reconstruction. This simple pre-processing, very close to the one proposed in Ziegler *et al* (2008), proved sufficient to reconstruct a volume as small as possible, containing only the patient, without generating overshoot and streak artifacts.

2.9. Physical phantom data

Since the patient's breathing motion of the day can differ from the one estimated on the 4D planning CT, it is important to evaluate whether MA-ROOSTER can handle inexact input motion information, and how errors in its input DVF, extracted from the 4D planning CT, propagate to the 4D CBCT reconstruction. MA-ROOSTER should theoretically show some robustness to such errors. To evaluate this robustness, we performed reconstructions with both underestimated and overestimated DVFs.

Acquisitions were performed on a 4D thorax phantom built at the Université Catholique de Louvain (UCL), and shown on figure 2. The phantom's diaphragm has a controllable motion period and amplitude, and can be stopped at specific positions. An insert was added close to the diaphragm, simulating a small tumor. The figures focus on slices extracted from the end-exhale phase, but each MA-ROOSTER reconstruction is a 4D sequence of eight volumes of size $244 \times 284 \times 308$, with isotropic 1 mm voxel size, representing the whole breathing cycle. We performed two dynamic acquisitions with either an 18 mm or a 9 mm amplitude in diaphragm motion on the cranio–caudal axis, and ten fixed acquisitions, with the 'diaphragm' at the following positions : 0, 2, 4, 6, 8, 10, 12, 14, 16 and 18 mm. The fixed acquisitions were reconstructed with the FDK algorithm (Feldkamp *et al* 1984), then arranged in two sequences to simulate 4D planning CTs. The first sequence was $\{0, 2, 4, 6, 8, 10, 8, 6, 4, 2\}$ and the second $\{0, 2, 4, 6, 8, 10, 12, 14, 16, 18, 16, 14, 12, 10, 8, 6, 4, 2\}$. From each of these fake 4D CT sequences, a pair of inverse-consistent DVFs was estimated using the method described in Janssens *et al* (2011). MA-ROOSTER's implementation can handle DVFs of an arbitrary number of frames (not necessarily the same as the number of reconstructed frames in f), by interpolating along time if required. The DVF-pair extracted from the first sequence describes a 10 mm amplitude motion, whereas the one extracted from the second sequence describes an 18 mm amplitude motion.

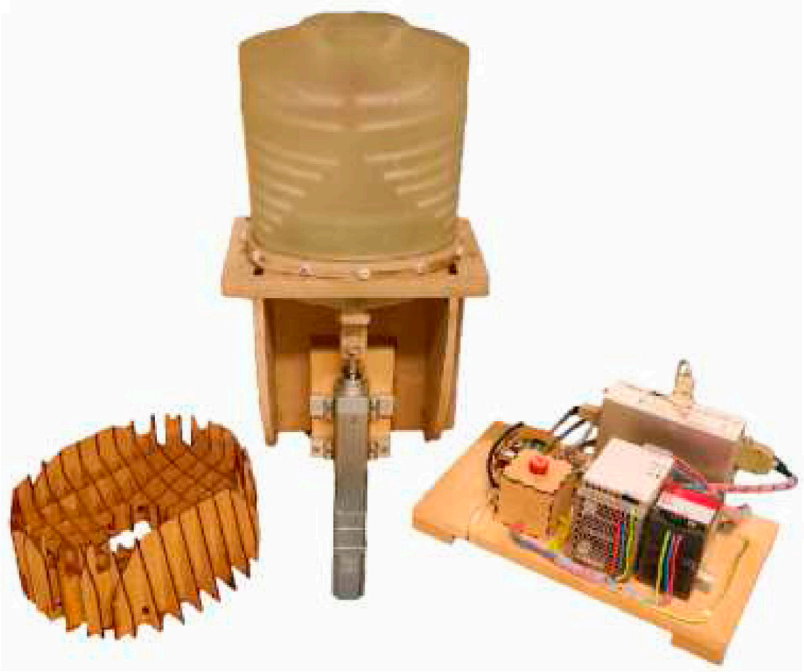

Figure 2. Photograph of the 4D thorax phantom. From left to right: support grid, actual thorax phantom with its hydraulic cylinder and electronic control board.

Note that these are the amplitudes for the motion of the diaphragm. The motion of the insert we placed on top of the diaphragm is of slightly lower amplitude, and the motion of the structures near the neck is of much lower amplitude, since the lungs are made of compressible foam.

Three experiments were conducted:

- Reconstructing the acquisition with 18 mm amplitude motion using the 18 mm amplitude DVF (correct motion).
- Reconstructing the acquisition with 18 mm amplitude motion using the 10 mm amplitude DVF (underestimated motion).
- Reconstructing the acquisition with 9 mm amplitude motion using the 18 mm amplitude DVF (overestimated motion).

In both ROOSTER and MA-ROOSTER, the γ_{time} parameter controls the trade-off between the attachment to the projection data and the attachment to the input DVF: $\gamma_{\text{time}} = 0$ means no regularization along time, i.e. the motion in the reconstruction is only the result of the motion in the projection data (desirable), but it also means a lot of streak artifacts (not desirable). On the other hand, $\gamma_{\text{time}} = +\infty$ means no streaks (desirable), but a strict attachment to the input DVF (not desirable). In the case of ROOSTER, the input DVF is null, so $\gamma_{\text{time}} = +\infty$ means no motion. Setting γ_{time} to obtain robustness to errors in the input DVF and high image quality, i.e. good contrast, sharp structures and few streaks, therefore requires a few trials. In radiotherapy, the primary goal of a 4D CBCT is to determine the real motion of the tumor, so for this application γ_{time} should be set to a small value, even if it means leaving some streaks in the image. In addition to evaluating the robustness of MA-ROOSTER to erroneous input motion information, the results presented in this section suggest a way to determine a suitable γ_{time} experimentally: acquire or generate projections of a moving phantom with known motion, perform several reconstructions with incorrect input motion and different values of γ_{time}, and

choose a value sufficiently low to yield robustness to erroneous input motion. In each experiment, four different values of γ_{time} were tested.

2.10. Clinical data

2.10.1. 4D planning CT. For both patients, the planning CT, from which the 4D DVF was extracted, has been acquired on a brilliance big bore 16-slice 4D CT scanner (Philips Medical Systems, Cleveland, OH). Each frame of this 4D CT is a $512 \times 512 \times 170$ voxel volume, with a voxel size of $0.88 \times 0.88 \times 2$ mm (the last dimension is the cranio–caudal axis). As with the phantom, for each patient a pair of inverse-consistent DVFs was estimated on the 4D planning CT using the method described in Janssens *et al* (2011).

2.10.2. CBCT data. Two CBCT acquisitions performed on different patients were reconstructed. Each one contains approximately 635 projections, each made of 512×512 pixels of size 0.8×0.8 mm. They have been acquired on an Elekta Synergy CBCT, along a 360 degrees trajectory, at 5.48 frames per second, using an off-center detector (Cho *et al* 1996) to enlarge the field of view. All 4D methods were set to reconstruct the respiratory cycle as a sequence of ten volumes. The size of the reconstructed volumes depends on the patient's size. It was $220 \times 280 \times 370$ voxels for patient 1, and $285 \times 270 \times 307$ voxels for patient 2, both with isotropic voxels of $1\,\text{mm}^3$.

2.10.3. Parameters. In both ROOSTER and MA-ROOSTER, γ_{space} was set to the same value as for the phantom study, i.e. $\gamma_{\text{space}} = 0.000\,05$. An animated GIF sequence available in the supplementary material shows that $\gamma_{\text{space}} = 0.0001$ leads to over-regularization (stacks.iop.org/PMB/61/6856/mmedia). As a rule of thumb, a regularization parameter should be set to the highest value that does not cause adverse effects, so we have set $\gamma_{\text{space}} = 0.000\,05$, which is close to the limit. γ_{time} was set to the value taken from the phantom study, to obtain robustness to erroneous input motion and limited streak artifacts removal, i.e. $\gamma_{\text{time}} = 0.0002$. The supplementary material also contains an animated GIF sequence showing the results obtained with a large number of different γ_{time} parameters (patient2_gamma_time.gif). Two other animated GIF sequences, with the same layout as figures 9 and 11, show the reconstructions of patients 1 and 2 through a full breathing cycle (patient1.gif and patient2.gif, respectively).

2.10.4. Reference. On clinical data, no ground truth of the DVF is available. The DVF estimated on the 4D planning CT may not be an accurate estimate of the patient's breathing motion on treatment days, especially if morphological changes, e.g. tumor shrinkage or atelectasis evolution, have occurred in the meantime. As a workaround, we propose using two respiration-correlated reconstructions, 4D FDK and 4D CG, as references for the motion of high-contrast structures (the low-contrast ones are either blurred or hidden by under-sampling artifacts). Since we extracted similar motion information from both of these unregularized reconstructions, we have assumed that this motion information was cross-validated and trustworthy.

3. Results

3.1. On a physical phantom

3.1.1. Correct motion. Figure 3 shows ROOSTER, MA-ROOSTER and MC-FDK reconstructions of the phantom with four different values of γ_{time}. The acquisition was performed with 18 mm amplitude motion, and the DVF used for the MA-ROOSTER and MC-FDK

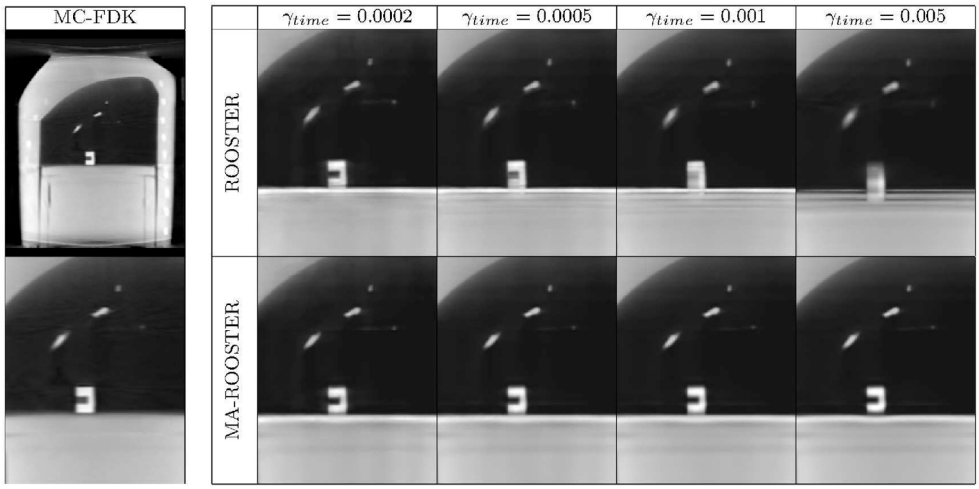

Figure 3. MC-FDK (full volume and zoom in), ROOSTER and MA-ROOSTER (zoom in) reconstructions of the UCL 4D phantom, with several values of the temporal regularization parameter. The real motion amplitude is 18 mm, and so is the motion amplitude in the input DVF.

reconstructions represents that motion. Note that ROOSTER does not make use of the DVF. When γ_{time} is high, MA-ROOSTER yields sharper results that ROOSTER, which is not surprising given the additional information it has made use of. Also, the MA-ROOOSTER and MC-FDK reconstructions have comparable image quality: when the input DVF is accurate, both MA-ROOSTER and MC-FDK perform well.

Figure 4 contains two graphs which show the amplitude of the insert's motion measured in the MA-ROOSTER reconstructions (on the left) and in the ROOSTER reconstructions (on the right), as well as on the 4D-FDK reconstruction (on both). The ideal profile, i.e. the one that shows perfect robustness to erroneous input motion information, is the dotted black line of the 4D-FDK. In this case, with various values of γ_{time}, the insert's position in the MA-ROOSTER reconstructions does not differ from the reference position by more than 1 mm. On the other hand, as γ_{time} increases, the motion amplitude in the ROOSTER reconstruction decreases: a higher γ_{time} means a stricter attachment to the input DVF, which for ROOSTER consists in null motion.

3.1.2. Underestimated motion. Figure 5 shows ROOSTER, MA-ROOSTER and MC-FDK reconstructions of the phantom with four different values of γ_{time}. The acquisition was performed with 18 mm amplitude motion, and the DVF used for the MA-ROOSTER and MC-FDK reconstructions only represents 10 mm amplitude motion. Since ROOSTER does not make use of the DVF, the ROOSTER results are the same as in figure 3. When γ_{time} is high, MA-ROOSTER still yields sharper results than ROOSTER, which confirms that MA-ROOSTER should yield better results than ROOSTER as soon as the DVF is closer to the real motion than a null DVF (see section 2.5). With a small value of γ_{time}, MA-ROOOSTER yields a much sharper reconstruction than MC-FDK: correctly tuned, MA-ROOSTER shows some robustness to errors in the DVF, while MC-FDK does not.

Figure 6 shows the amplitude of the insert's motion measured in the MA-ROOSTER reconstructions and in a 4D-FDK reconstruction. With a small value of γ_{time}, both ROOSTER and MA-ROOSTER yield an accurate reconstruction of the motion pattern. With higher values of γ_{time}, both methods tend to a stricter attachment to their input DVF. Since the one

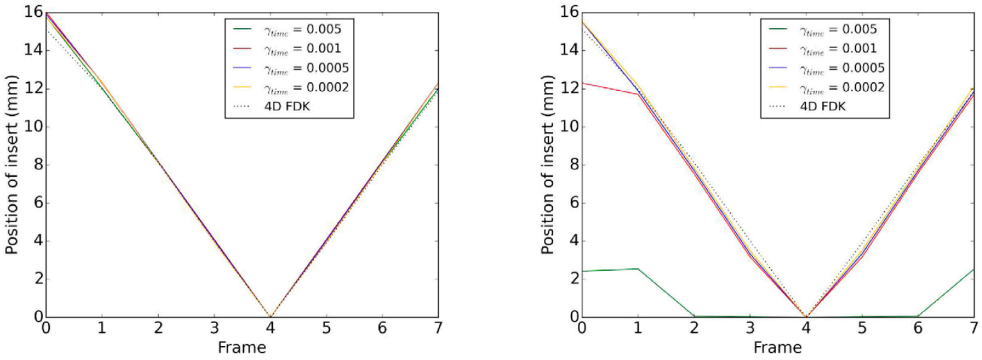

Figure 4. Detected position of the moving insert in the MA-ROOSTER reconstructions (on the left) and in the ROOSTER reconstructions (on the right) with correct input motion. The 4D FDK is used as a reference.

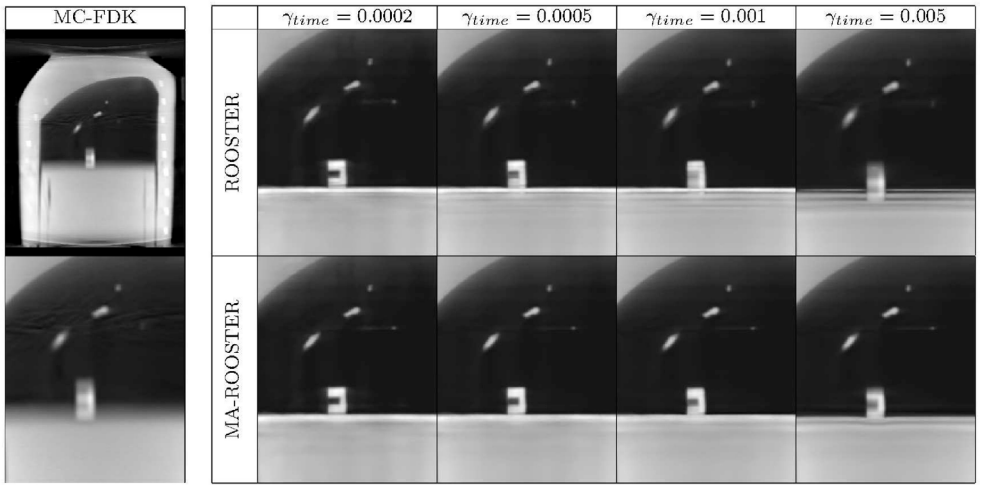

Figure 5. ROOSTER, MA-ROOSTER and MC-FDK reconstructions of the UCL 4D phantom, with several values of the temporal regularization parameter. The real motion amplitude is 18 mm, but the input DVF only models 10 mm amplitude motion.

used in MA-ROOSTER is closer to reality than the null motion assumed by ROOSTER, MA-ROOSTER performs better in that case.

3.1.3. Overestimated motion. Figure 7 shows ROOSTER, MA-ROOSTER and MC-FDK reconstructions of the phantom with four different values of γ_{time}. The acquisition was performed with 9 mm amplitude motion, and the DVF used for the MA-ROOSTER and MC-FDK reconstructions represents 18 mm amplitude motion. Again, ROOSTER does not make use of the DVF, but the ROOSTER results are different from those in figures 3 and 5 since the projection data have changed. This time, even with γ_{time} high, there is no clear ranking between MA-ROOSTER, which assumes a motion of amplitude 18 mm, and ROOSTER, which assumes a motion of amplitude 0 mm (the real motion amplitude is 9 mm). This observation is consistent with the statement in section 2.5: if the MA-ROOSTER's input DVF is not a better estimate of the real motion than a null DVF, there is

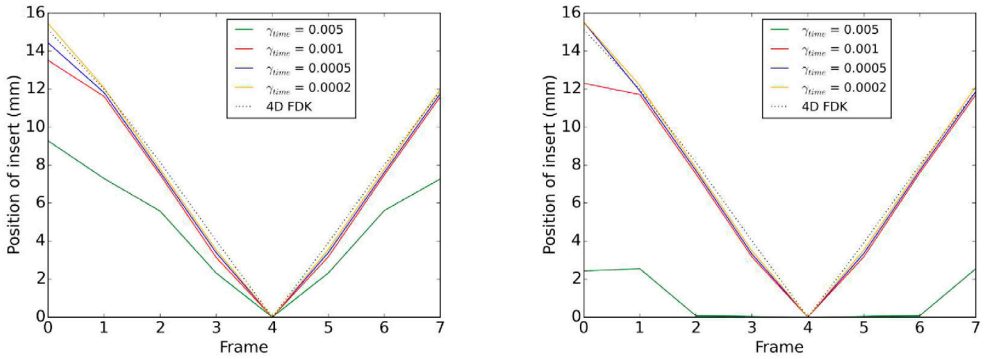

Figure 6. The detected position of the moving insert in the MA-ROOSTER reconstructions (on the left) and in the ROOSTER reconstructions (on the right) with underestimated input motion. The 4D FDK is used as a reference.

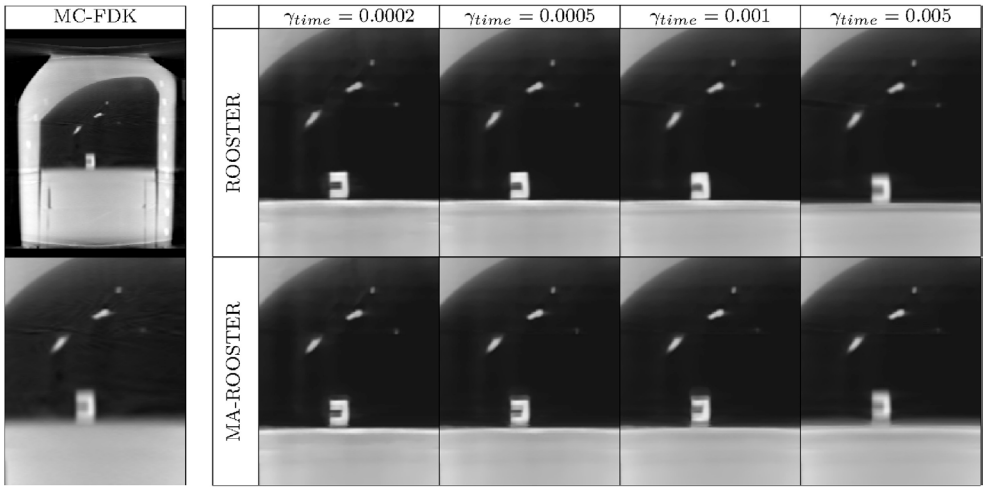

Figure 7. ROOSTER, MA-ROOSTER and MC-FDK reconstructions of the UCL 4D phantom, with several values of the temporal regularization parameter. The real motion amplitude is 9 mm, but the input DVF models 18 mm amplitude motion.

no reason why MA-ROOSTER should yield better results than ROOSTER. Similarly, as in figure 5, with a small value of γ_{time} MA-ROOOSTER yields a much sharper reconstruction than MC-FDK.

Figure 8 shows the amplitude of the insert's motion measured in the MA-ROOSTER reconstructions and in a 4D-FDK reconstruction. The ROOSTER results allow one to locate the tumor more precisely than MA-ROOSTER ones. In this case, assuming an overestimated motion (MA-ROOSTER) is worse than assuming null motion (ROOSTER). Such a situation, however, is very unlikely to occur in a clinical context.

3.2. On patients

As explained in section 2.5, the temporal regularization enforces a trade-off between the motion present in the projection data and the one described by the input DVF. Since in

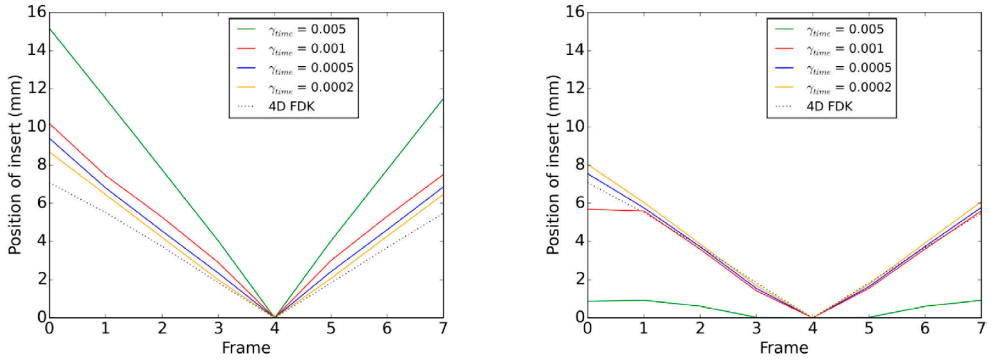

Figure 8. Detected position of the moving insert in the MA-ROOSTER reconstructions (on the left) and in the ROOSTER reconstructions (on the right) with overestimated input motion. The 4D FDK is used as a reference.

radiotherapy one of the the primary goals of 4D CBCT is to compare the breathing motion of the treatment day with that of the planning CT, we recommend seting γ_{time} to stick to the motion of the projection data. From the phantom study, $\gamma_{\text{time}} = 0.0002$ seems the safest choice, and $\gamma_{\text{time}} = 0.0005$ would be 'strong but acceptable' regularization. We have chosen $\gamma_{\text{time}} = 0.0002$. In other applications where motion evaluation is less critical, e.g. cardiac road-mapping (Knecht *et al* 2008), γ_{time} could be set to a higher value.

3.2.1. Patient 1. Figure 9 shows sagittal and coronal slices of the reconstructions obtained using 4D FDK, 4D CG, ROOSTER and MA-ROOSTER on patient 2, at the end-inhale (rows 1 and 3) and end-exhale (rows 2 and 4) phases. MC-FDK should show the end-exhale since the reference image of the DVF is the end-exhale. The small moving structures in the lungs are better contrasted in the MA-ROOSTER reconstruction than in the ROOSTER one, in particular on the end-inhale frame.

In terms of the contrast and sharpness of small structures, MA-ROOSTER does much better than 4D CG and 4D FDK, but still has some streaks that are absent from the MC-FDK (e.g. around the ribs). We recall that the MC-FDK is used only as a reference for image quality, since it cannot provide information on the motion on the treatment day.

Figure 10 shows the tumor motion throughout the breathing cycle, measured as described in section 2.7. Mostly, it is cranio–caudal motion, but in some patients (not here) the tumor may also undergo high-amplitude antero–posterior motion. The reference motion amplitude, i.e. the difference between tumor position at end-inhale and end-exhale, measured on both unregularized 4D reconstructions, differs from the one estimated on the planning CT by 2 to 3 mm. In both the ROOSTER and MA-ROOSTER reconstructions the tumor position is at maximum 1 mm away from the references.

Artifacts can be observed on the diaphragm and liver of patient 1 at end-exhale, in MA-ROOSTER reconstructions. In the corresponding area on the MC-FDK, the region is darker than clinically expected. These artifacts have the same profound cause, namely that in some phases the motion compensation leads outside the field of view. The difference between the nature of the artifacts arise from the different ways motion compensation is performed in the two methods. In MA-ROOSTER, all frames are warped to a single one and regularization along time is performed on the warped frames. For some frames, this warping implies forward or backward mapping outside the field of view, right in the cone-beam artifacts, which 'brings in' cone-beam artifacts. These are then spread out to the other frames through regularization

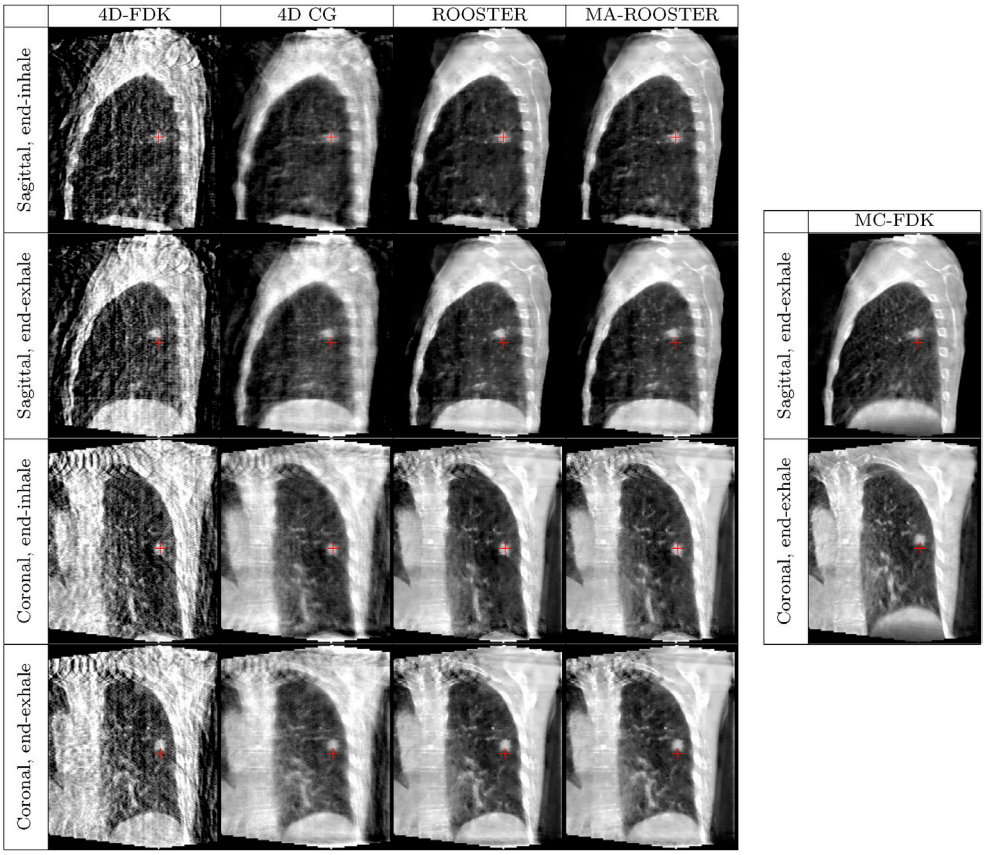

Figure 9. Slices through reconstructions of patient 1. Left panels (from left to right): 4D-FDK, 4D CG, ROOSTER and MA-ROOSTER. Right panels: MC-FDK of the end-exhale phase. The red cross marks a fixed spot, approximately at the center of the tumor in the end-inhale position, to ease visual evaluation of the motion's amplitude.

along time. In MC-FDK, when the DVF points outside the field of view, the projection data is assumed to be null. Fewer back-projected rays reach the inside of the FOV, which results in an attenuation that is lower than expected. Patient 1 also has metal artifacts, visible mostly in the coronal view, caused by a tracheotomy device. Note that cone-beam artifacts can also be observed at the top of the reconstruction of both patients, although these do not interfere with motion compensation.

3.2.2. Patient 2. Figure 11 shows sagittal and coronal slices of the reconstructions obtained using 4D FDK, 4D CG, ROOSTER and MA-ROOSTER on patient 2, in end-inhale (rows 1 and 3) and end-exhale (rows 2 and 4) phases. MC-FDK shows the end-exhale.

The comparison with ROOSTER yields the same results as for patient 1. On this patient, however, MA-ROOSTER achieves a slightly higher contrast than MC-FDK on small structures, especially below the tumor in the sagittal view, but has a slightly lower contrast on the tumor itself. Even with the red cross as a reference spot, it is difficult to notice a change in motion amplitude between the various methods. However, this time, the motion estimation results are clearly in favor of MA-ROOSTER, as shown in figure 12, which is similar to figure 10 but for patient 2. The reference motion amplitude, measured on both

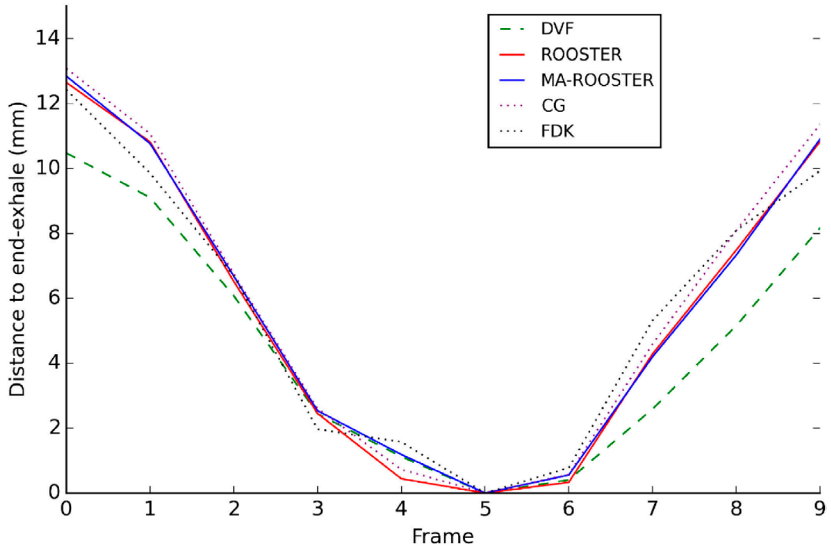

Figure 10. Distance to end-exhale position (in mm) estimated around the tumor on several 4D reconstructions of patient 1, and computed from the DVF.

unregularized 4D reconstructions, is 4 mm larger than the one estimated on the planning CT. The discrepancy between tumor position in the ROOSTER reconstruction and in the references reaches 3 mm at end-exhale, while it remains within 1 mm throughout the cycle in the MA-ROOSTER result. Again, this tends to prove that MA-ROOSTER is robust to inaccuracies of the input DVF.

3.2.3. Linear attenuation coefficients of lung tissue. An interesting feature of MA-ROOSTER is that it allows a variation of the linear attenuation coefficients of lung tissue throughout the respiratory cycle. To measure this variation, we warped all frames onto frame 50% and summed the attenuation of all voxels contained in the mask shown on figure 13. This was performed on the planning CT images and the MA-ROOSTER reconstructions. Since the planning CT images are in Hounsfield units (HU), a small calculation is necessary to make sure that the HUs and linear attenuation coefficients are supposed to follow the same variations. For a position $\mathbf{x}$ and a phase p,

$$\begin{aligned} \mathrm{HU}(\mathbf{x},p) &= \frac{\mu(\mathbf{x},p)-\mu_{\mathrm{air}}}{\mu_{\mathrm{water}}-\mu_{\mathrm{air}}}\times 1000 - 1000 \\ \mathrm{HU}(\mathbf{x},p)+1000 &\approx \frac{\mu(\mathbf{x},p)}{\mu_{\mathrm{water}}}\times 1000, \quad \text{since } \mu_{\mathrm{air}}\approx 0. \end{aligned} \tag{4}$$

The ratio between attenuation in frame p and in frame 50% is

$$\frac{\mu(\mathbf{x},p)}{\mu(\mathbf{x},50\%)} = \frac{\mathrm{HU}(\mathbf{x},p)+1000}{\mathrm{HU}(\mathbf{x},50\%)+1000}. \tag{5}$$

The ratios between the sum of voxel values in frame p and in frame 50% should therefore be the same in the planning CT expressed in 'HU + 1000' and in the MA-ROOSTER reconstruction. Figure 14 shows these ratios throughout the breathing cycle on patient 1. The mean attenuation in lung tissue in MA-ROOSTER reconstructions follows the same trend as in the 4D CT reconstruction.

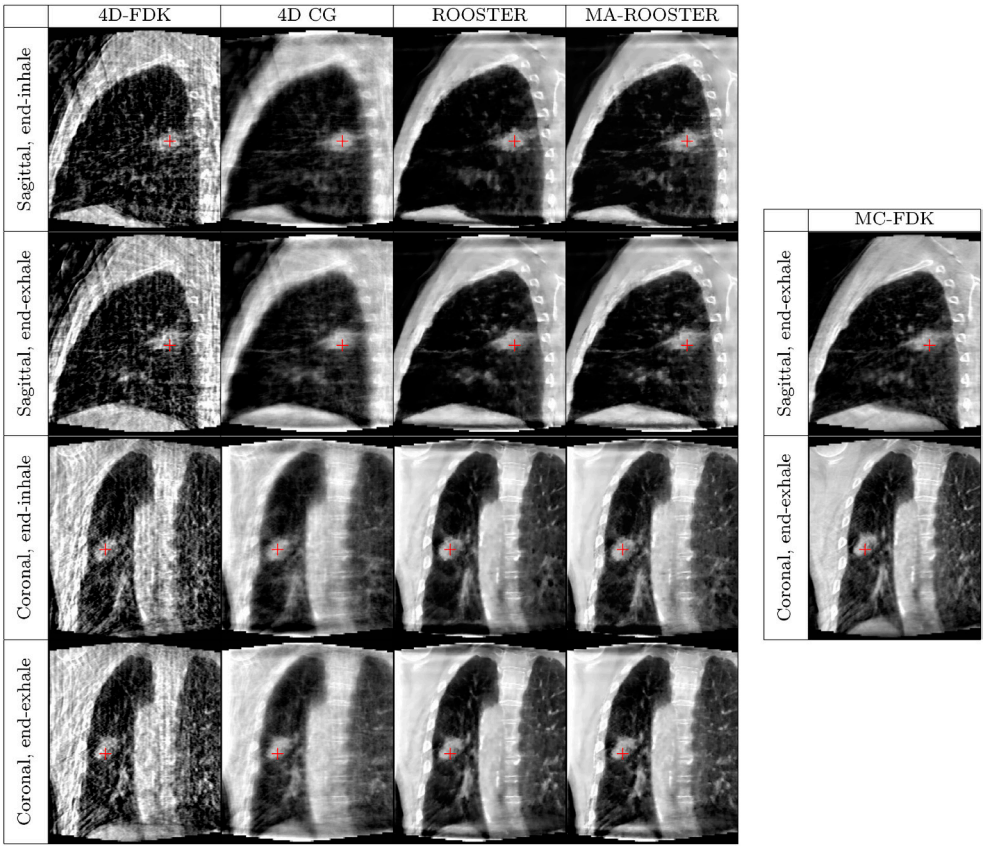

Figure 11. Slices through reconstructions of patient 2. Left panels (from left to right): 4D-FDK, 4D CG, ROOSTER and MA-ROOSTER. Right panels: MC-FDK of the end-exhale phase. The red cross marks a fixed spot, approximately at the center of the tumor in end-inhale position, to ease visual evaluation of the motion's amplitude.

4. Discussion

4.1. Comparison to ROOSTER and MC-FDK

MA-ROOSTER outperforms ROOSTER when the DVF used is closer to the real motion than a null DVF. MA-ROOSTER also outperforms MC-FDK when the DVF does not exactly represent the motion of the day, since it corrects for some of the DVF's inaccuracies. For real image-guided radiation therapy (IGRT) cases, any sensible motion estimation on the planning CT, even a rough one, will usually be a better estimate of the motion of the day than nothing. On the other hand, even if the motion estimation on the planning CT is perfect, it is impossible to know *a priori* whether or not it represents accurately the motion of the day. Therefore, for real IGRT cases, MA-ROOSTER is likely to provide more reliable reconstructions than both ROOSTER and MC-FDK.

With respect to the classical approach used in 'motion-compensated' tomography, which consists in bending the forward and back-projection trajectories, motion-aware regularization is a new way to make use of an existing motion estimation. The additional parameter to tune, γ_{time}, is to be seen more as an additional degree of freedom than as an additional burden, since

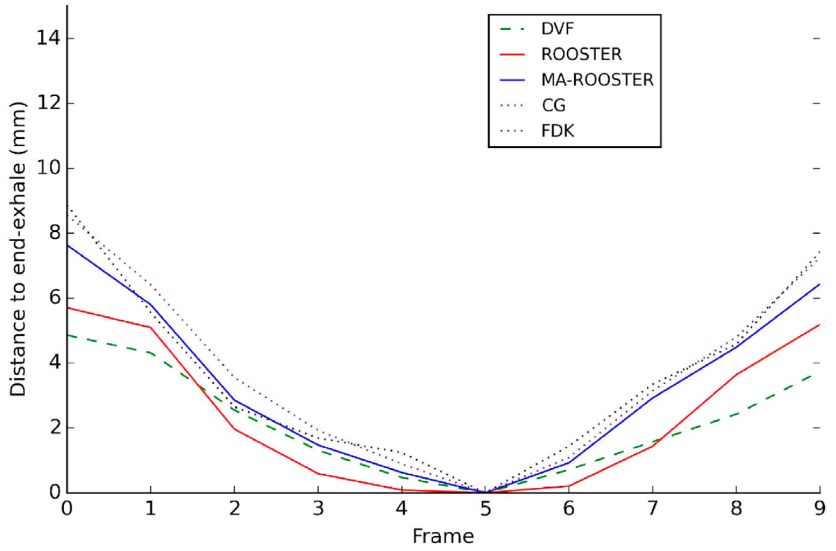

Figure 12. Distance to end-exhale position (in mm) estimated around the tumor on several 4D reconstructions of patient 2, and computed from the DVF.

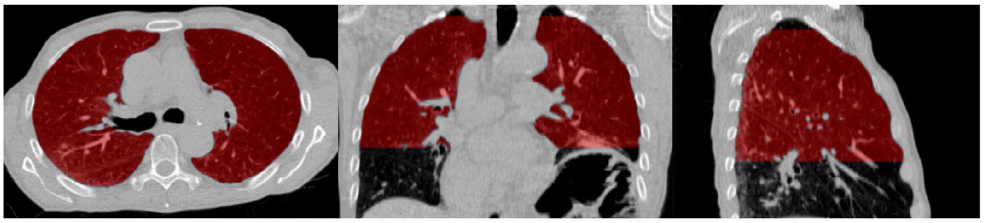

Figure 13. Mask used to measure attenuation variation in lung tissue, shown on frame 50% of the planning CT.

setting $\gamma_{\text{time}} = +\infty$ yields results very similar (although not theoretically identical) to motion-compensated ones.

Other denoising methods exist which could be described as 'motion-aware', e.g. temporal non-local means (TNLM) (Tian *et al* 2011), since in regularizing between consecutive frames they assume that the underlying structures may have moved. TNLM does not require an input DVF, but is computationally more demanding than MA-ROOSTER. In addition, the implicit motion it uses is likely to be very irregular, and therefore not a proper description of the real motion.

4.2. Phantom studies

It is very unusual, in a clinical context, to have such a large discrepancy in the motion amplitude between the planning CT and the CBCT as the ones we used in the phantom study: when anatomical changes occur that are likely to alter the patient's breathing amplitude, doctors usually order a re-planning of the treatment on a new 4D planning CT. Our phantom experiments are therefore quite extreme cases (Seppenwoolde *et al* 2002, Rit *et al* 2012). On real data, we expect the planning CT's motion to be closer to the motion of the day.

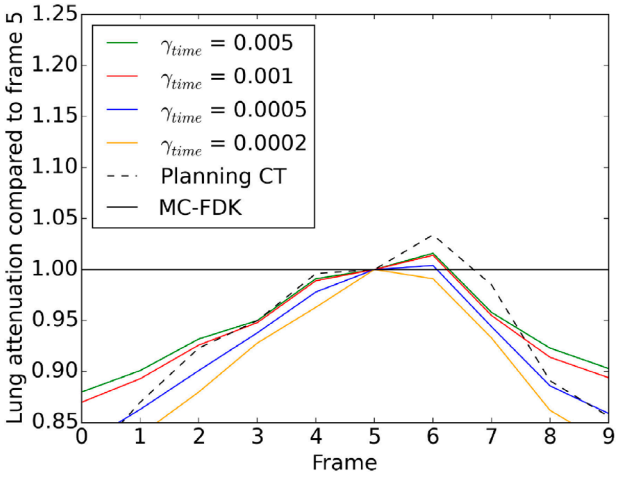

Figure 14. Ratio between the mean attenuation in lungs in a given frame and in frame 50%, in the planning CT and MA-ROOSTER reconstructions of the CBCT data. MC-FDK is shown to recall that it cannot describe such variations.

4.3. Variation of lung tissue attenuation during the respiratory cycle

While it is clear that lung tissue should not have a constant attenuation over time (Guerrero *et al* 2006), it is hard to say how this attenuation should vary. We restricted ourselves to pointing out that the variation of attenuation observed in the MA-ROOSTER reconstructions is consistent with that of the planning CT, as shown on figure 14. This feature of MA-ROOSTER could prove important in applications that require quantitative CBCT, e.g. (Bernchou *et al* 2015).

4.4. Other applications

In interventional cardiology, an accurate 4D reconstruction of the patient's beating heart would allow functional analysis such as left ventricle ejection fraction measurement, detection of hypo- or a-kynesia of some myocardium segments, and road-mapping for numerous interventions like electrophysiology or aortic valve replacement. But many interventional cardiology procedures do not require a 4D CT scan, therefore no DVF is available. For patients who did undergo a 4D CT before their 4D CBCT, MA-ROOSTER could be tested on 4D cardiac CBCT data.

Since MA-ROOSTER has been proved to correct for some of the input DVF's inaccuracies, further work could involve estimating a new DVF from the MA-ROOSTER reconstruction, and performing a second MA-ROOSTER with this new DVF as the input.

4.5. Convergence

Neither ROOSTER nor MA-ROOSTER come with a convergence proof. Furthermore, as each iteration of the main loop is rather long, the stopping criterion is not based on some convergence measurement, but on the number of iterations. Although ROOSTER and MA-ROOSTER behave well in practice for the cases we have studied, there is no theoretical guarantee that they converge, nor that the solution they yield at convergence is more desirable than their output after ten iterations. An animated GIF sequence has been added

to the supplementary material which follows the reconstructed image through 100 iterations (patient2_iterations.gif). Ten iterations seems to be a good choice, since when the number of iterations rises no major image quality improvement is observable, and some dark and bright dots appear. These dots are presumably caused by imperfect inverse warping, which causes some voxels to be modified by temporal regularization when they should not.

Observations on the convergence of ROOSTER, involving the theory of non-expansive mappings, can be found in Mory *et al* (2014). As each step of ROOSTER can be expressed as a proximal mapping, an algorithm similar to ROOSTER can be obtained by minimizing a carefully chosen cost function with the Chambolle–Pock method (Chambolle and Pock 2011). Such an attempt can be found in Mory and Jacques (2014), but the resulting algorithm proved impractical because of its slow initial convergence. Future work includes transforming ROOSTER and MA-ROOSTER into efficient proximal algorithms.

4.6. Regularization

Spatial TV regularization has been shown to be better suited to phantom images than to real clinical data (Mory *et al* 2012), as it favors piecewise-constant images. The spatial TV denoising step could be replaced with some wavelet-based denoising. Finding better-suited regularizers and implementing them efficiently is also part of the future work on MA-ROOSTER.

4.7. Computational cost

The reconstructions were performed on an Intel Xeon E5-2620 CPU with 12 cores, equipped with an nVidia GTX780 GPU, running OpenSuse 13.1. All three methods were implemented using the Reconstruction ToolKit (RTK) (Rit *et al* 2013), an open source C++ software based on the Insight ToolKit' (ITK). With this set-up, the total reconstruction time with MA-ROOSTER for patient 1 is 21 min, divided as follows: 4D CG optimization took 17 min, spatial TV denoising 45 s, warping 80 s, inverse warping 100 s and the other operations can be neglected. When the motion estimation method used on the planning CT only yields a single 4D DVF (instead of two inverse-consistent 4D DVFs), MA-ROOSTER uses an iterative procedure to perform the inverse warping, which increases the duration of that step to 19 min, while all the other execution times remain the same.

4.8. Implementation

All the reconstruction methods used in this paper (4D-FDK, MC-FDK, 4D CG, ROOSTER and MA-ROOSTER) have been implemented in the RTK library (http://openrtk.org/), an open-source software based on ITK. RTK is available to anyone, documented, and we provide help on how to compile, use and modify the code through a mailing list open to everyone, as well as via a wiki (http://wiki.openrtk.org/). A page on the RTK wiki specifically describes how to use ROOSTER and MA-ROOSTER (http://wiki.openrtk.org/index.php/RTK/Examples/4DROOSTERReconstruction) and contains links to the data of both patients (4D planning CT, projections, geometry, respiratory signal and DVFs), as well as the command lines to reproduce the results on patients. The phantom data (projections, geometry, simulated respiratory signal and DVFs) can be made available on request. The implementation of the method we used to compute the 4D DVFs from the 4D planning CT (Janssens *et al* 2011), on the other hand, is not open source. Note that MA-ROOSTER can use DVFs generated by any motion estimation method, and that how the motion estimation is performed is beyond the scope of the present paper.

5. Conclusion

The proposed method, MA-ROOSTER, yields a reconstruction that is visually close to MC-FDK, while being robust to motion estimation inaccuracies. MA-ROOSTER appears as an excellent trade-off between the MC-FDK technique (which provides high-quality 3D images, but from which one cannot retrieve the motion of the day nor the attenuation variation due to the flows of air and blood in the lungs) and unregularized 4D reconstruction techniques (from which one can estimate the motion of the day of large structures, but which have poor image quality).

Acknowledgments

The authors would like to thank Sébastien Brousmiche and Guillaume Janssens for their help in performing the acquisitions on the phantom. This research work has been partially funded by the iMagX project. iMagX is a public partnership between UCL and IBA funded by the Walloon region under convention number 1017266 and 1217662, and by grant ANR-12-BS01-0018 (DROITE project).

References

Andersen A H and Kak A C 1984 Simultaneous algebraic reconstruction technique (SART): a superior implementation of the art algorithm *Ultrason. Imaging* **6** 81–94

Arsigny V *et al* 2006 A log-Euclidean framework for statistics on diffeomorphisms *Medical Image Computing and Computer-Assisted Intervention* (*Lecture Notes in Computer Science* vol 4190) ed R Larsen *et al* (Berlin: Springer) pp 924–31

Bergner F *et al* 2010 An investigation of 4D cone-beam CT algorithms for slowly rotating scanners *Med. Phys.* **37** 5044

Bernchou U *et al* 2015 Prediction of lung density changes after radiotherapy by cone beam computed tomography response markers and pre-treatment factors for non-small cell lung cancer patients *Radiother. Oncol.* **117** 17–22

Boyd S and Ye Y 2011 *Foundations and Trends in Optimization* (New York: Dover)

Brehm M *et al* 2012 Self-adapting cyclic registration for motion-compensated cone-beam CT in image-guided radiation therapy *Med. Phys.* **39** 7603–18

Chambolle A and Pock T 2011 A first-order primal-dual algorithm for convex problems with applications to imaging *J. Math. Imaging Vis.* **40** 120–45

Cho P S, Rudd A D and Johnson R H 1996 Cone-beam CT from width-truncated projections *Comput. Med. Imaging Graph.* **20** 49–57

Christensen G E and Johnson H J 2001 Consistent image registration *IEEE Trans. Med. Imaging* **20** 568–82

Feldkamp L A, Davis L C and Kress J W 1984 Practical cone-beam algorithm *J. Opt. Soc. Am.* A **1** 612–9

Guerrero T *et al* 2006 Dynamic ventilation imaging from four-dimensional computed tomography *Phys. Med. Biol.* **51** 777–91

Jacques L, Hammond D K and Fadili J 2010 Dequantizing compressed sensing: when oversampling and non-Gaussian constraints combine (arXiv:0902.2367v4)

Jaffray D A *et al* 2002 Flat-panel cone-beam computed tomography for image-guided radiation therapy *Int. J. Radiat. Oncol. Biol. Phys.* **53** 1337–49

Janssens G *et al* 2011 Diffeomorphic registration of images with variable contrast enhancement *Int. J. Biomed. Imaging* **2011** 891585

Jia X *et al* 2010 4D computed tomography reconstruction from few-projection data via temporal non-local regularization *Medical Image Computing and Computer-Assisted Intervention* (*Lecture Notes in Computer Science* vol 6361) ed T Jiang *et al* (Berlin: Springer) pp 143–50

Knecht S *et al* 2008 Computed tomography-fluoroscopy overlay evaluation during catheter ablation of left atrial arrhythmia *Europace* **10** 931–8

Leng S *et al* 2008 High temporal resolution and streak-free four-dimensional cone-beam computed tomography *Phys. Med. Biol.* **53** 5653–73
Li T, Koong A and Xing L 2007 Enhanced 4D cone-beam CT with inter-phase motion model *Med. Phys.* **34** 3688
Liu J *et al* 2015 5D respiratory motion model based image reconstruction algorithm for 4D cone-beam computed tomography *Inverse Problems* **31** 115007
Lu J *et al* 2007 Four-dimensional cone beam CT with adaptive gantry rotation and adaptive data sampling *Med. Phys.* **34** 3520
McCollough C H *et al* 2012 Achieving routine submillisievert CT scanning: report from the summit on management of radiation dose in CT *Radiology* **264** 567–80
Moeslund T B 2012 *Introduction to Video and Image Processing: Building Real Systems and Applications* (Berlin: Springer)
Mory C and Jacques L 2014 A modified 4D ROOSTER method using the Chambolle–Pock algorithm *Proc. 3rd Int. Conf. on Image Formation in X-ray Computed Tomography (Salt Lake City, USA) pp* 191–3 (http://ucair.med.utah.edu/CTmeeting/ProceedingsCTMeeting2014.pdf)
Mory C and Rit S 2015 Improving iterative 4D CBCT through the use of motion information *Proc. of Fully 3D 2015 (Newport, Rhode Island, USA)*
Mory C *et al* 2012 ECG-gated C-arm computed tomography using L1 regularization *Proc. 20th European Signal Processing Conf.* pp 2728–32
Mory C *et al* 2014 Cardiac C-arm computed tomography using a 3D+ time ROI reconstruction method with spatial and temporal regularization *Med. Phys.* **41** 021903
Rit S *et al* 2009a On-the-fly motion-compensated cone-beam CT using an *a priori* model of the respiratory motion *Med. Phys.* **36** 2283–96
Rit S, Sarrut D and Desbat L 2009b Comparison of analytic and algebraic methods for motion-compensated cone-beam CT reconstruction of the thorax *IEEE Trans. Med. Imaging* **28** 1513–25
Rit S *et al* 2011 Comparative study of respiratory motion correction techniques in cone-beam computed tomography *Radiother. Oncol.* **100** 356–9
Rit S, van Herk M, Zijp L and Sonke J J 2012 Quantification of the variability of diaphragm motion and implications for treatment margin construction *Int. J. Radiat. Oncol. Biol. Phys.* **82** e399–407
Rit S, Vila Oliva M, Brousmiche S, Labarbe R, Sarrut D and Sharp G C 2013 The Reconstruction RoolKit (RTK), an open-source cone-beam CT reconstruction toolkit based on the Insight ToolKit (ITK). *Proc. Int. Conf. on the Use of Computers in Radiation Therapy*
Ritschl L *et al* 2012 Iterative 4D cardiac micro-CT image reconstruction using an adaptive spatio-temporal sparsity prior *Phys. Med. Biol.* **57** 1517–25
Seppenwoolde Y *et al* 2002 Precise and real-time measurement of 3D tumor motion in lung due to breathing and heartbeat, measured during radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **53** 822–34
Sidky E Y and Pan X 2008 Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization *Phys. Med. Biol.* **53** 4777–807
Sonke J J *et al* 2005 Respiratory correlated cone beam CT *Med. Phys.* **32** 1176
Tian Z *et al* 2011 Low-dose 4DCT reconstruction via temporal nonlocal means *Med. Phys.* **38** 1359–65
Vandemeulebroucke J *et al* 2012 Automated segmentation of a motion mask to preserve sliding motion in deformable registration of thoracic CT *Med. Phys.* **39** 1006–15
Wang J and Gu X 2013a High-quality four-dimensional cone-beam CT by deforming prior images *Phys. Med. Biol.* **58** 231
Wang J and Gu X 2013b Simultaneous motion estimation and image reconstruction for 4D cone-beam CT *Med. Phys.* **40** 101912
Wolthaus J W *et al* 2008 Comparison of different strategies to use four-dimensional computed tomography in treatment planning for lung cancer patients *Int. J. Radiat. Oncol. Biol. Phys.* **70** 1229–38
Wu H *et al* 2012 Spatial-temporal total variation regularization for 4D-CT reconstruction *Proc. SPIE* **8313** 83133J
Ziegler A, Nielsen T and Grass M 2008 Iterative reconstruction of a region of interest for transmission tomography *Med. Phys.* **35** 1317