Optimizing 4DCBCT Projection Allocation to Respiratory Bins

Ricky T O'Brien[1], John Kipritidis[1], Chun-Chien Shieh[1,2] and Paul J Keall[1]

[1] Radiation Physics Laboratory, Sydney Medical School, The University of Sydney, NSW 2006, Australia. [2] Institute of Medical Physics, School of Physics, The University of Sydney, NSW 2006, Australia

E-mail: ricky.obrien@sydney.edu.au

Abstract. Four dimensional cone beam computed tomography (4DCBCT) is an emerging image guidance strategy used in radiotherapy where projections acquired during a scan are sorted into respiratory bins based on the respiratory phase or displacement. 4DCBCT reduces the motion blur caused by respiratory motion but increases streaking artefacts due to projection under-sampling as a result of the irregular nature of patient breathing and the binning algorithms used. For displacement binning the streak artefacts are so severe that displacement binning is rarely used clinically. The purpose of this study is to investigate if sharing projections between respiratory bins and adjusting the location of respiratory bins in an optimal manner can reduce or eliminate streak artefacts in 4DCBCT images. We introduce a mathematical optimization framework and a heuristic solution method, which we will call the optimized projection allocation algorithm, to determine where to position the respiratory bins and which projections to source from neighbouring respiratory bins. Five 4DCBCT datasets from three patients were used to reconstruct 4DCBCT images. Projections were sorted into respiratory bins using equispaced, equal density and optimized projection allocation. The standard deviation of the angular separation between projections was used to assess streaking and the consistency of the segmented volume of a fiducial gold marker was used to assess motion blur. The standard deviation of the angular separation between projections using displacement binning and optimized projection allocation was 30%-50% smaller than conventional phase based binning and 59%-76% smaller than conventional displacement binning indicating more uniformly spaced projections and fewer streaking artefacts. The standard deviation in the marker volume was 20%-90% smaller when using optimized projection allocation than using conventional phase based binning suggesting more uniform marker segmentation and less motion blur. Images reconstructed using displacement binning and the optimized projection allocation algorithm were clearer, contained visibly fewer streak artefacts and produced more consistent marker segmentation than those reconstructed with either equispaced or equal-density binning. The optimized projection allocation algorithm significantly improves image quality in 4DCBCT images and provides, for the first time, a method to consistently generate high quality displacement binned 4DCBCT images in clinical applications.

1. Introduction

Four dimensional cone beam computed tomography (4DCBCT) imaging is an emerging image guidance strategy used to position patients for treatment in radiotherapy. 4DCBCT was first published between 2003 and 2005 (Taguchi 2003) and (Sonke et al. 2005) and commercially released by Elekta (Stockholm, Sweden) in 2009 and Varian (Palo Alto, USA) in 2013. 4DCBCT was developed to overcome image blurring in (3D)CBCT which is caused by respiratory motion. On the day of treatment, 4DCBCT provides valuable information on the average tumour position, the magnitude of the tumour motion, validation of the treatment plan and the changing tumour size and shape.

CBCT (or 3D CBCT) imaging involves rotating the kilovoltage imager around a patient at a constant speed and acquiring 2D projections (kilovoltage images) with a constant time interval and constant gantry angle between projections. The series of 2D projections can be reconstructed into a CBCT image using the Feldkamp-Davis-Kress (FDK) algorithm (Feldkamp et al. 1984). An example of a CBCT image is given in Figure 1. It can take between 30 seconds and 4 minutes to rotate the gantry around the patient in which time the patient's anatomy moves due to respiratory motion.

4DCBCT imaging differs from CBCT imaging in that projections are allocated to different phases of the respiratory cycle and used to reconstruct an image for the corresponding respiratory phase. For example, projections acquired at peak exhale are allocated to a peak exhale respiratory bin and are used to reconstruct a CBCT image in the peak exhale respiratory bin, see Figure 2 for an example. Within each respiratory bin there is little anatomical motion and blurring artefacts are greatly reduced. The main problem with 4DCBCT imaging is streak artefacts. Streak artefacts occur because a constant gantry speed and constant projection pulse rate are used leading to large angular gaps between projections (Leng et al. 2008). Figure 2 shows the gantry angle for the projections acquired in two neighbouring displacement bins with large angular gaps, due to irregular breathing, in both respiratory bins. In this study we 'close' the gaps by sharing projections between neighbouring respiratory bins.

In recent years there have been a number of studies with the focus on reducing the streak artefacts in 4DCBCT images. There are two common approaches taken: (1) A reconstruction approach where prior images, compressed sensing, iterative reconstruction and deformable image registration are used (Mckinnon & Bates 1981), (Leng et al. 2008), (Sidky & Pan 2008), (Chen et al. 2008) and (2) A hardware approach where the gantry speed and projection acquisition are modulated in order to evenly distribute projections in each respiratory bin (O'Brien et al. 2013), (O'Brien et al. 2014), (Cooper et al. 2013) and (Fast et al. 2013).

Of the reconstruction approaches, there are two where projections, or pixels from projections, are sourced from outside the respiratory phase of interest. Bergner *et al.* introduce the auto adaptive phase correlated (AAPC) reconstruction for 4DCBCT (Bergner et al. 2009). In this approach pixels, or regions in each projection, with

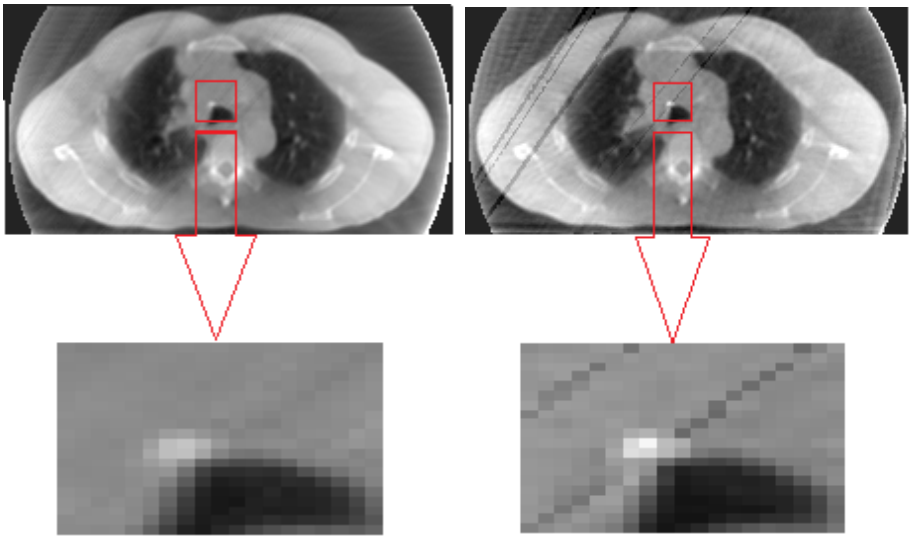

Figure 1. Left: A 3D CBCT image containing 2388 projections of a lung cancer patient with an expanded view of the region containing the marker. Right: A 4DCBCT displacement binned image in the peak exhale respiratory bin. The marker is blurred and less intense in the 3D CBCT image due to respiratory motion. However, the streaks are clearly evident in the 4DCBCT image with a light streak running through the marker.

little motion were identified. The pixels in regions at rest were weighted for use in the reconstruction. However, it was demonstrated that the AAPC method produced more motion blur than the leading 4DCBCT reconstruction algorithms because data was sourced from respiratory phases that were too far away from the phase of interest (Bergner et al. 2010). Projections from the neighbouring respiratory bin have been used to improve micro-CT of free breathing mice (Armitage et al. 2012). However, they exploited the regular, and predictable, nature of the cardiac cycle to determine which projections to source from the neighbouring respiratory bin so the algorithms are not applicable to respiratory induced 4DCBCT.

In this paper we make the best use of the projections acquired during 4DCBCT imaging by optimising the respiratory bin position, size and projection allocation. Although the projections in neighbouring bins are from a different part of the breathing phase, they contain valuable information about the patient's anatomy and can be used to reduce streak artefacts. The result of this optimisation are 4DCBCT images with fewer streak artefacts. We will also assess the amount of motion blur introduced using 4DCBCT scans of lung cancer patients with implanted fiducial gold markers.

2. Theory

The theoretical treatment begins by presenting the full set of equations that represent the optimal respiratory bin and projection allocations using Mixed Integer Programming techniques (MIP) (Nemhauser & Wolsey 1988). Even with the state of the art MIP

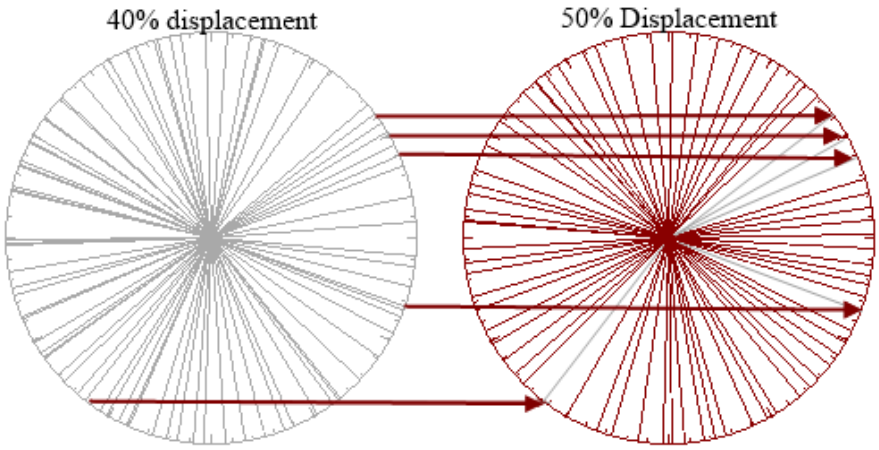

Figure 2. Projection clustering and optimized projection allocation schematic. The gantry angle for each projection in two neighbouring respiratory bins generated using 4DCBCT, the breathing trace from a lung cancer patient, and displacement binning. The arrows indicate projections that could be shared, or used in both respiratory bins, to try to close the angular gaps between projections. Peak inhale is at 100% displacement, mid inhale occurs at 50% displacement and peak exhale is at 0% displacement.

solvers the problem is *NP*-Hard and we are only able to solve problems to optimality with 300 projections or fewer on current computers; this is well below practical problem sizes used for 4DCBCT imaging in radiotherapy. We will then present a simple heuristic solution method that can obtain a near optimal, but not provably optimal, solution to the MIP model. We will show that the simple heuristic greatly improves image quality. The equations presented apply to both phase based binning and displacement binning.

2.1. Modelling the respiratory bin size and position

Let r_j and θ_j be the respiratory signal, either phase or displacement, and gantry angle respectively for projection j of P projections acquired across B respiratory bins. We assume that the projections are sorted in order from the smallest gantry angle $j = 1$ to the largest gantry angle $j = P$. To model the location of each respiratory bin we let R_b^u and R_b^l be the upper and lower respiratory signal respectively for respiratory bin b ($b = 1, 2, \ldots, B$). The values of R_b^u and R_b^l are to be determined as part of the optimisation but to make sure that the respiratory bins are contiguous we must have

$$R_b^u = R_{b+1}^l \text{ for } b = 1, 2, \ldots, B-1,$$

with $R_1^l \equiv r_{min} = \min_j\{r_j\}$ and $R_B^u \equiv r_{max} = \max_j\{r_j\}$.

To ensure that respiratory bins are not too large, and span a range of respiratory signals where significant anatomical motion takes place, we need to place restrictions on the maximum size of each respiratory bin. If respiratory bins were evenly spaced then we would expect the size of each respiratory bin, $(R_b^u - R_b^l)$, to be $\Delta = (r_{max} - r_{min})/B$.

We allow the respiratory bins to grow by a factor g or shrink by a factor s

$$\Delta(1-s) \leq R_b^u - R_b^l \leq \Delta(1+g) \text{ for } b = 1, 2, \ldots, B.$$

In our simulations we will allow the respiratory bins to be 50% larger ($g = 0.5$) or smaller ($s = 0.5$) than Δ.

2.2. Allocating projections to respiratory bins

To determine if a projection is allocated to a respiratory bin we introduce binary variables $\delta_{b,j}$ which take the value 1 if projection j is allocated to respiratory bin b and zero otherwise. The values of $\delta_{b,j}$ are to be determined as part of the optimisation. The proximity constraints used to determine if the projection belongs to a respiratory bin are

$$r_j \geq r_b^l \delta_{b,j} - \Delta_b^l \Delta, \tag{1}$$

$$r_j \leq r_b^u + \Delta_b^u \Delta + r_{max}(1 - \delta_{b,j}), \tag{2}$$

for $b = 1, 2, \ldots, B$ and $j = 1, 2, \ldots, P$. The values of Δ_b^l and Δ_b^u are zero if we do not allow sharing of projections, but can take a value if sharing of projection from neighbouring respiratory bins is allowed. The notation $\Delta_b^l \Delta$ allows us to refer to sharing of projections from the whole neighbouring respiratory bin using $\Delta_b^l = 1$ and $\Delta_b^u = 1$ or half the neighbouring respiratory bin using $\Delta_b^l = 0.5$ and $\Delta_b^u = 0.5$.

To make sure that every projection is allocated to a respiratory bin we must have

$$\sum_b \delta_{b,j} = 1 \text{ if projections cannot be shared between respiratory bins,}$$

$$\sum_b \delta_{b,j} \geq 1 \text{ if projections can be shared between respiratory bins,}$$

2.3. The objective function

An objective function that has been found to correlate well with image quality is the standard deviation, σ, of the angular separation between projections (Shieh et al. 2014), (O'Brien et al. 2013) and (O'Brien et al. 2014). If the standard deviation is zero then the projections in each respiratory bin are uniformly spaced and image quality is likely to be good. For a respiratory bin with 120 projections, and a standard deviation of 3°, the standard deviation is equal to the average separation between projections (i.e the average separation is $360°/120 = 3°$) and we might expect to see a cluster of two projections close together followed by a large gap before the next cluster of projections. Larger gaps between consecutive projections contribute to the standard deviation more than nearly uniformly spaced projections, so the optimisation algorithms are likely to reduce the standard deviation more by closing the large gaps.

For illustration purposes let $\theta_{b,k}$ be the k^{th} largest gantry angle for the projections in respiratory bin b ($k = 1, 2, \ldots, P_b$ where $P_b = \sum_{j=1}^{P} \delta_{b,j}$ are the number of projections

in respiratory bin b) then the standard deviation in respiratory bin b is

$$\sigma_b^2 = \left[\sum_{k=1}^{P_b-1} \left(\theta_{b,k+1} - \theta_{b,k} - \mu_b\right)^2 + \left(\theta_{b,1} + 2\pi - \theta_{b,P_b} - \mu_b\right)^2 \right]$$

where $\mu_b = 2\pi/P_b$ is the mean angular separation between projections. We average the standard deviation for each respiratory bin to give the average standard deviation of the angular separation between projections for the 4DCBCT scan

$$\sigma = \sum_{b=1}^{B} \sigma_b / B \tag{3}$$

We use equation 3 when we want to calculate the standard deviation from a given set of projections. We will make extensive use of this equation in our heuristic solution methods below. Unfortunately, we cannot minimise the standard deviation in an optimisation algorithm because equation 3 contains quadratic terms and the value of P_b is to be determined as part of the optimisation. However, equation 3 can be reduced to linear terms which are suitable for use in mixed integer programming (MIP) solvers. To avoid interrupting the flow of this paper we present the linear version of the standard deviation in Appendix A.

One problem with the standard deviation is that it can be made small by trading-off one respiratory bin for another. For example, one respiratory bin may contain two projections (at 0° and 180°) while the neighbouring respiratory bin may contain 500 evenly spaced projections. To stop this happening we add additional constraints to ensure that each respiratory bin contains at least 120 projections and the standard deviation in each respiratory bin is less than 3°.

3. Method

Even with the fastest commercial mixed integer programming solvers (XPRESS-MP, GUROBI and ILOG-CPLEX) it is not possible to solve problems with more than about 300 projections to optimality. For larger problems the computation time grows exponentially with problem size and in our experiments an optimal solution to a 1200 projection problem was not found within one month on a 16 Core 3.1GHz machine.

3.1. Heuristic solution methods

Although a provably optimal solution is difficult to obtain, we can make considerable progress using heuristic solution methods (Talbi 2009). Heuristic solution methods are used to obtain a good/near optimal solution in a small amount of time. There are a large range of heuristic solution methods available and each have their advantages and disadvantages (Talbi 2009). Below we have used a very simple heuristic solution method that obtains a solution in one or two seconds making it a viable candidate for clinical use.

3.1.1. The simple heuristic: There are two steps in the simple heuristic. The first step involves optimising the location of the respiratory bins and the second step involves determining if projections in neighbouring respiratory bins will improve the standard deviation.

The Respiratory Bin Position Algorithm (BPA): The respiratory bin position algorithm moves the position of each respiratory bin up or down a small increment, δ. The algorithm moves the position of the respiratory bin down N increments, for a total distance of δN, and then up N increments, for a total distance of $N\delta$. The position of the respiratory bin that produces the lowest standard deviation is accepted as the best solution.

Algorithm 1 The Respiratory Bin Position Algorithm (BPA).

Initialise variables used only in this algorithm:
$N = 20$
$\delta = (\Delta_b^l + \Delta_b^u)\Delta/(2N)$
$\sigma_{min} = 10000$
for $b = 1, 2, \ldots, B$ **do**
  Initialise temporary variables:
  $U_{init}^b = (b+1)\Delta$ or to the best value of R_b^u known.
  for *For* $i = -N, \ldots, N$ **do**
    Move bin boundaries up, or down, a small increment via
    $R_b^{l+1} = U_{init}^b - \Delta_b^l\Delta + i\delta$
    $R_b^u = R_b^{l+1}$
    for *For* $j = -N, \ldots, N$ **do**
      if *All constraints, e.g. bin sizes, are satisfied.* **then**
        if *Extensive Heuristic* **then**
          Calculate σ using the projection allocation algorithm (PAA).
        else
          Calculate σ from equation 3.
        end
        if $\sigma < \sigma_{min}$ **then**
          $\sigma_{min} = \sigma$
          Record R_b^u as the best value found.
        end
      end
    end
  end
end

Optimized Projection Allocation Algorithm (PAA): The projection allocation algorithm identifies projections in the neighbouring respiratory bins that can fill gaps. If theses projections reduce the standard deviation then they are accepted as the best solution.

Algorithm 2 Optimized Projection Allocation Algorithm (PAA). This algorithm can be run several times.

If not running the extensive heuristic then run the bin position algorithm (BPA).
$\sigma_{min} = 10000$
for $b = 1, 2, \ldots, B$ **do**
  For each respiratory bin find the candidate list, L_i for $i = 1, 2, \ldots, M$, of M projections that are in a neighbouring respiratory bin and satisfy the proximity constraints (equations 1 and 2).
  Sort L_i from the largest angular gap between projections, $i = 1$, to smallest angular gap between projections, $i = M$.
  for $i = 1, 2, \ldots, M$ **do**
    Add projection L_i to respiratory bin b.
    if *All constraints are satisfied* **then**
      Calculate σ from equation 3.
      if $\sigma < \sigma_{min}$ **then**
        $\sigma_{min} = \sigma$
        Permanently add the projection to the respiratory bin.
      end
    end
  end
end
return σ_{min}.

3.1.2. The extensive heuristic: Applying the simple heuristic produces a good solution within 1 or 2 seconds of computation time and is a very useful algorithm to apply in practice. To determine if the simple heuristic produces a solution that is close to optimality we apply a more extensive heuristic that takes approximately 12 hours to determine if better solutions are available. The extensive heuristic is a simple extension on the simple heuristic and applies the projection allocation algorithm at each window position in the respiratory bin position optimisation algorithm. That is we run the optimized projection allocation algorithm at each step of the respiratory bin position algorithm.

3.2. Patient image data

To test the optimized projection allocation algorithm, five 4DCBCT data sets from three patients in the study by (Roman et al. 2012) were used. The 4DCBCT datasets were selected because the patients had fiducial gold markers implanted from which we could extract a respiratory signal during the entire 4DCBCT scan. The marker around the thoracic cavity was used to extract the respiratory signal because it was the most easily segmented of the three markers.

The five scans consist of 2360 (patient 1 scan 1), 2406 (patient 1 scan 2), 2508 (patient 2 scan 1), 2435 (patient 2 scan 2) and 2390 (patient 3 scan 1) half fan projections

respectively with a resolution of 1024 by 768 pixels (0.388mm per pixel). For phase based binning, respiratory cycles were first identified by determining the peak inhale and exhale points.

3.3. Binning options

In all simulations a total of ten phase or displacement bins were used. The following binning options apply to both phase and displacement binning.

3.3.1. Equispaced binning: Respiratory cycles were divided into 10 equally spaced respiratory bins based on either phase or displacement. For displacement binning the peak inhale and exhale points were averaged to determine the location of the respiratory bins. Any projections falling above or below the average points were allocated to the peak inhale or peak exhale respiratory bins respectively.

3.3.2. Equal density binning: The recorded respiratory signals for each projection were sorted from lowest to highest in either phase or displacement. The lowest 10% were allocated to the first respiratory bin, the second 10% were allocated to the second respiratory bin and so on. This method ensures that every bin has exactly the same number of projections but the size and location of each respiratory bin cannot be controlled.

3.3.3. Optimized projection allocation (Optimized binning): Both the simple and extensive heuristic have been used to determine the location of respiratory bins and the allocation of projections to the respiratory bins. In the optimisation the respiratory bins were allowed to grow, or shrink, by 50% and projections could be sourced from either half way ($\Delta_b^l = \Delta_b^u = 0.5$) or from the whole neighbouring respiratory bin ($\Delta_b^l = \Delta_b^u = 1$). A minimum of 120 projections per respiratory bin were required to ensure that image quality is adequate. We refer to this binning method as optimized projection allocation or optimized binning. We will refer to the simple or extensive heuristic when using optimized projection allocation with the simple or extensive heuristic respectively.

3.4. Image reconstruction

Images were reconstructed using COBRA‡ to give 96 transverse slices of dimension 224×224 pixels. Each slice in the reconstructed image was 2mm and the voxel size was 2mm×2mm×2mm which is commonly used clinically.

3.5. Marker volume estimation

Because our analysis involves real patient data, we do not have a ground truth image to make comparisons. To determine if sourcing projections from neighbouring respiratory

‡ COBRA, Exxim Computing Corporation, 3825 Hopyard Road, Suite 220, Pleasanton, CA 94588.

bins has blurred the image we segment a fiducial gold marker and calculate the volume of the marker in each respiratory bin. If significant blurring of the marker has taken place then we expect the volume of the marker to be inconsistent across the respiratory bins. The fiducial gold markers were 0.35mm×10mm or 20mm in length (Visicoil, RadioMed Corp., Tynsboro, MA) (Roman et al. 2012). It should be noted that recovering the original volume of the marker is difficult because the diameter of the marker, 0.35mm, is less than the voxel size and the coils stretch and curl on insertion. However, apart from a small amount of deformation between respiratory bins we expect the volume of the marker to be reasonably consistent from one respiratory bin to the next. This gives us a method to analyse the patient data for motion blur with more conclusive metrics to be calculated from the XCAT phantom.

To calculate the size of the gold markers from the reconstructed 4DCBCT images, we select a region around the markers of 30mm×30mm×24mm in the lateral, anterior-posterior and superior-inferior directions respectively. We compute the mean and standard deviation of the voxels within the region. Voxels with an intensity of two standard deviations above the mean were selected as candidate voxels. From the candidate voxels the largest connected cluster of voxels were selected as the marker. We then calculate the centre-of-mass (COM) of the marker to represent the location of the marker.

Failed segmentations occur when large streaks run through the marker. For example, in Figure 4, respiratory bin 0 with equal density binning has streaks running through the marker which result in a failed segmentation. Failed segmentations were identified by comparing the position of the segmented marker with the position in the two neighbouring respiratory bins. If the distance is greater than 3mm then the respiratory bin is marker as a failed segmentation. If consecutive respiratory bins had failed segmentation then comparisons were made with the closest bin having a successful segmentation or the phase binned marker trajectories were used as a reference.

3.6. Digital XCAT phantom data

We do not have a ground truth for the patient data, so the digital XCAT phantom (Segars et al. 2010) was used to facilitate ground truth comparisons. To make the tumour easy to segment, a 19mm tumour was placed in the middle of the right lung. For each projection in the patient data, and corresponding displacement signal, the XCAT phantom was deformed according to the breathing signal and forward projected. These projections were generated with the same resolution as the patient data so that the only difference between the patient data and the XCAT data was the medium through which the projections were generated.

For the XCAT phantom the tumour size was known so the difference between the segmented tumour volume was compared to the known tumour volume. The absolute value of the marker volume error was calculated and then averaged across the 10 respiratory bins to give the mean marker volume error. The standard deviation in

the marker volume was also calculated in the same way as for the patient data.

3.7. Estimating streaks from the images

To quantify streaking, the streak ratio of (Leng et al. 2008) has been calculated for each reconstructed image. The streak ratio uses the total variation (the sum of the image gradient) as an estimate of the number of streaks in an image. The streak ratio is the ratio of the total variation of the reconstructed image divided by the total variation of a ground truth, or streak free, image. A value of 1 indicates that there are no streaks while values above 1 indicate more extensive streaking.

For each patient, we have calculated the streak ratio for each slice in the reconstructed image set by using the 3D (or motion blurred) image as the streak free baseline image. We calculate the streak ratio for each slice and each respiratory bin and then average the values across slices and respiratory bins to give the average streak ratio.

4. Results

There are three different parameters that we present in our analysis of the results: (1) reconstructed images (2) the standard deviation of the angular separation between projections and (3) an analysis of the consistency of the marker size segmented in the reconstructed images.

4.1. Reconstructed images

A transverse slice of patient 1 scan 1 was selected showing the location of the fiducial gold marker. Figure 3 contains reconstructed images for phase based binning using equispaced, equal density and optimized binning. Optimized binning produces better, or similar, quality images in all respiratory bins when compared to equispaced or equal density binning. For bin 0 with equispaced binning and bin 6 with equal density binning optimized binning was able to eliminate the dominant streak artefacts. For each respiratory bin, including the bins not shown, optimized binning was as good, or better than, the best image from either equispaced or equal density binning. There are no large streaks observed, as seen in bin 0 with equispaced binning, in the 10 respiratory bins for optimized binning.

Figure 4 gives reconstructed images for displacement binning using equispaced, equal density and optimized binning. For displacement binning the image quality is visibly improved using optimized binning with the most significant improvement in the peak inhale and peak exhale respiratory bins (bins 0 and 9). In the peak inhale and peak exhale respiratory bins, although the streak artefacts have been reduced with optimized binning, there are still some streak artefacts. The streaks can be reduced further by allowing projections to be sourced from further into the neighbouring respiratory bin,

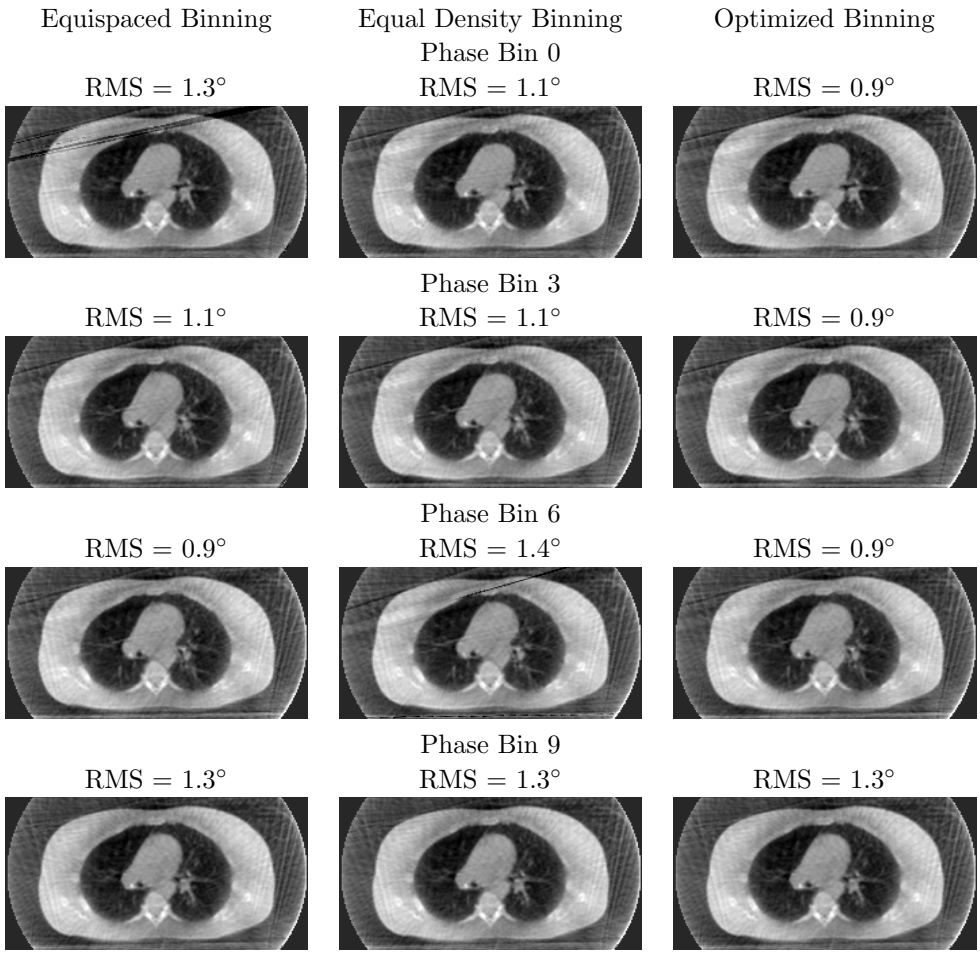

Figure 3. A transverse slice for patient 1 scan 1 using phase binning. Columns one, two and three use equispaced, equal density and optimized projection sorting respectively. Rows one, two, three and four are respiratory bins 0, 3, 6 and 9 respectively. Optimized binning allowed projections to be sourced from half way into the neighbouring respiratory bin.

or, allowing the bins to grow by more than 50%. Similar image quality is observed in the respiratory bins now shown.

4.2. The standard deviation between projections

As the objective of the optimisation is to reduce the standard deviation of the angular separation between projections it is important to establish how much we can reduce the standard deviation. In Table 1 we present the standard deviation for the five 4DCBCT datasets using equispaced, equal density and optimized binning (the simple heuristic). In all cases the standard deviation is reduced using optimized binning when compared to either equispaced or equal density binning. The improvement in the standard deviation is more significant for displacement binning. For displacement binning the standard deviation is reduced by at least 50% over both equal density and equispaced binning for all parameter settings and patient scans. For phase based binning the improvement

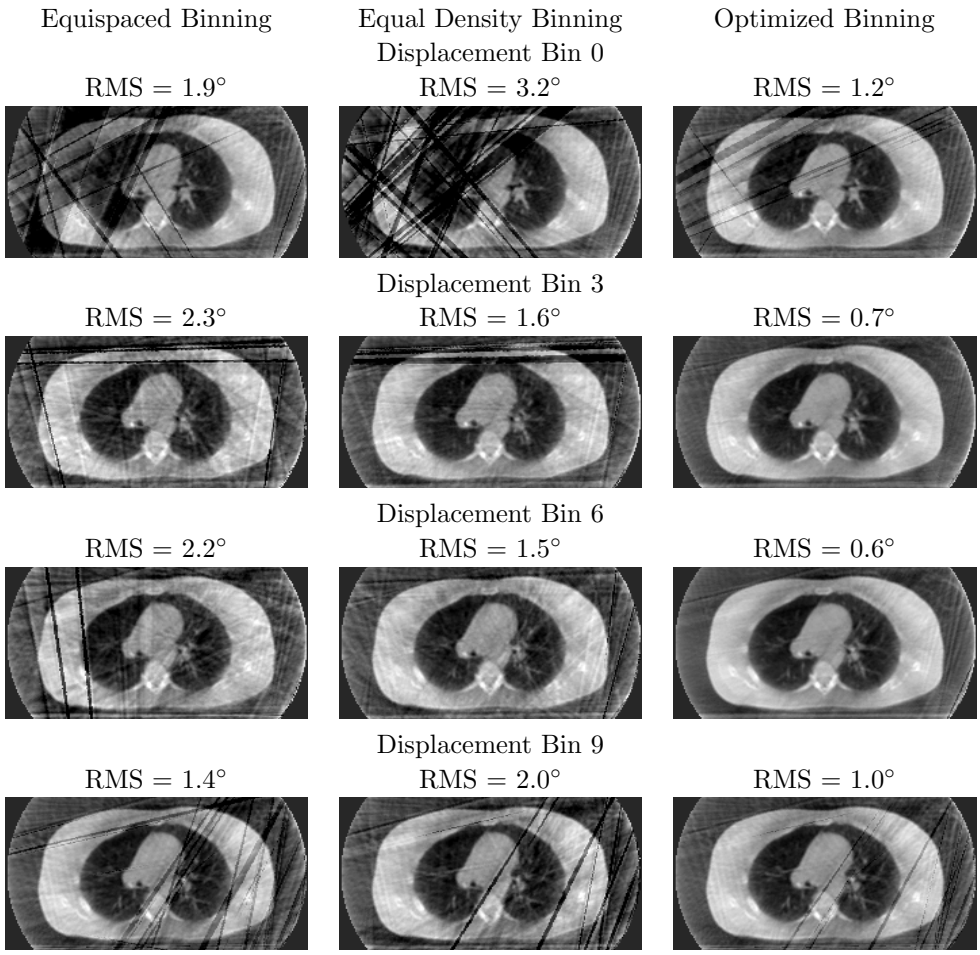

Figure 4. A transverse slice for patient 1 scan 1 using displacement binning. Columns one, two and three use equispaced, equal density and optimized projection sorting respectively. Rows one, two, three and four are respiratory bins 0, 3, 6 and 9 respectively. Optimized binning allowed projections to be sourced from half way into the neighbouring respiratory bin. With optimized binning the streak artefacts are significantly reduced and image quality is significantly better.

in the standard deviation is usually in the range of 10-30% over both equal density and equispaced binning.

The simple heuristic includes two components, the bin position algorithm (BPA) and the projection allocation algorithm (PAA). In Table 1 we have used only one algorithm at a time to examine if one algorithm dominates. Comparing the BPA, PAA columns to the simple heuristic (SH), we can see that for phase binning the BPA produces a smaller standard deviation than the PAA algorithm for 4 of 5 cases when sharing is from half way into the neighbouring bin. However, the PAA performs better when sourcing projections from the entire neighbouring bin. The final bin positions, are highly dependent on the patients unique breathing pattern so no systematic behaviour in the bin position adjustment was observed. For displacement binning the PAA performs better than the BPA and produces the same standard deviation as the SH in 9 of 10

Table 1. The standard deviation, equation 3, of the angular separation between projections. The sharing distance parameters (Δ_b^u and Δ_b^l) can be either 0.5 or 1.0 allowing projections to be sourced from half way and the full neighbouring respiratory bin respectively. BPA applies only the bin position algorithm. PAA applies only the projection allocation algorithm. The simple heuristic (SH) applies both the BPA and PAA. SH allows bins to shrink by 50% and grow by 50%. SH100 allows bins to shrink by 50% and grow by 100%. The numbers in brackets are the average streak ratio for the patient dataset.

				$\Delta_b^l = \Delta_b^u = 0.5$			$\Delta_b^l = \Delta_b^u = 1.0$		
Patient -Scan	Equi-spaced	Equal-density	BPA	PAA	SH	SH100	PAA	SH	SH100
				Phase Based Binning					
1-1	1.2° (1.4)	1.3° (1.4)	1.0°	1.2°	0.9°	0.9°	1.0°	0.9° (1.3)	0.9°
1-2	1.6° (1.5)	1.3° (1.4)	1.0°	1.1°	0.9°	0.9°	0.9°	0.8° (1.3)	0.8°
2-1	1.6° (1.4)	1.5° (1.4)	1.4°	1.3°	1.2°	1.3°	1.2°	1.1° (1.3)	1.2°
2-2	1.5° (1.5)	1.5° (1.5)	1.3°	1.4°	1.2°	1.2°	1.2°	1.1° (1.5)	1.1°
3-1	1.8° (1.6)	1.8° (1.5)	1.7°	1.6°	1.6°	1.6°	1.5°	1.5° (1.5)	1.3°
				Displacement Binning					
1-1	1.9° (1.9)	1.9° (1.8)	1.8°	0.8°	0.8°	0.8°	0.6°	0.6° (1.2)	0.6°
1-2	2.6° (3.2)	3.9° (3.6)	2.6°	1.1°	1.1°	1.1°	0.7°	0.7° (1.7)	0.7°
2-1	2.5° (2.3)	2.4° (2.5)	2.4°	1.0°	1.0°	1.0°	0.7°	0.6° (1.3)	0.6°
2-2	2.2° (2.1)	4.3° (2.5)	2.1°	0.8°	0.8°	0.8°	0.6°	0.6° (1.2)	0.6°
3-1	2.0° (1.8)	3.2° (2.1)	2.0°	0.8°	0.8°	0.8°	0.7°	0.7° (1.3)	0.7°

cases. The BPA could be considered optional for displacement binning.

In Table 1 the difference between SH and SH100 is that bins are allowed to grow by 100% (i.e. double in size) in the SH100 simulation. For displacement binning SH and SH100 produce the same standard deviation. For phase binning, only a small improvement is obtained when using SH100 compared to SH. These results are not surprising because if one bin increases in size then the other bins must shrink to account for the increased respiratory bin size; which degreases the image quality in these bins as a consequence. For the remainder of this study we will allow bins to shrink or grow by 50% and vary the distance from which projections can be shared.

The streak ratio, the numbers in brackets in Table 1, follow a similar pattern to the standard deviation of the angular separation between projections. This indicates that the standard deviation of the angular separation between projections is a good metric for estimating streaks in images.

Table 2 lists the percentage of projections sourced from neighbouring respiratory bins. With phase based binning, the number of projections sourced from neighbouring respiratory bins ranges from 22% to 32%. For displacement binning more projections are sourced from neighbouring respiratory bins with the number ranging from 44% to 60% of the total projections used to reconstruct each image. The number of projections sourced from the neighbouring respiratory bin can be reduced by reducing the parameters Δ_b^l

Table 2. The simple heuristic versus the extensive heuristic with projection sharing from the neighbouring respiratory bin ($\Delta_b^u = \Delta_b^l = 1.0$). Foreign projections are the percentage of projections that have been sourced from the neighbouring respiratory bin.

	Simple Heuristic			Extensive Heuristic		
Patient -Scan	σ °	Total Proj's	Foreign Proj's	σ °	Total Proj's	Foreign Proj's
Phase Based Binning						
1-1	0.9°	2995	21%	0.8°	2873	23%
1-2	0.8°	3236	26%	0.8°	3258	25%
2-1	1.1°	3412	26%	1.2°	3224	22%
2-2	1.1°	3451	29%	1.1°	3581	32%
3-1	1.5°	3097	23%	1.4°	3248	26%
Displacement Binning						
1-1	0.6°	4732	50%	0.6°	3746	49%
1-2	0.7°	5150	53%	0.7°	5361	55%
2-1	0.6°	6317	60%	0.6°	6206	60%
2-2	0.6°	4755	49%	0.5°	4955	51%
3-1	0.7°	4535	47%	0.7°	4292	44%

and Δ_b^u. Comparing the simple and extensive heuristic in Table 2 gives an indication on the quality of the simple heuristic. For both phase and displacement binning the extensive heuristic was able to improve the standard deviation over the simple heuristic. However, the improvements are small and they do not justify the increase in computation time (up to 12 hours). The implication is that the simple heuristic will produce a good solution for most problems of practical interest within the clinically viable computation time of 1-2 seconds.

For phase binning, projections are only shared in the irregular breathing cycles with more sharing occurring for the more irregular breathing cycles. For displacement binning, in addition to the irregular breathing cycles, more projection sharing occurs in the peak inhale and peak exhale respiratory bins.

4.3. Quantifying marker blur

Table 3 lists the segmented marker volume in voxels for the five 4DCBCT datasets. For patient 1, scans 1 and 2, we have segmented different markers but both markers are of the same size (10mm in length). To give further variety to our results, for patient 2, we have segmented two markers of different length which are a marker of 10mm length for scan 1 and 20mm length for scan 2.

We are interested in the standard deviation in the marker volume across the 10 respiratory bins because if the standard deviation is higher then the implication is that the marker is more blurry, and has a less consistent volume, across the 10 respiratory bins. Phase based binning indicates similar standard deviations in the marker volume

between the three binning methods for all five patient scans which indicates that marker blur is primarily caused by using phase as the respiratory signal. For phase based binning, patient 2 scan 1 is the only scan where there was a failure to segment the marker in one or more respiratory bins with optimized binning failing in fewer cases than both equispaced and equal density binning.

Displacement binning indicates that optimized binning is able to segment the marker in all cases while equispaced and equal-density binning have a large number of failed segmentations (usually at inhale and exhale limits). Failed segmentations occur more often with equal-density binning than equispaced binning during mid inhale. The standard deviation for optimized binning is always lower, and often by a significant amount, than both equispaced and equal-density binning.

An important result is that displacement binning, with the optimized projection allocation algorithm, produces a lower standard deviation in marker volumes than phase based binning for all datasets except for patient 1 scan 2. This indicates that on average the marker can be more consistently segmented using displacement binning than phase binning. These results are not surprising because phase binning sources projections from a much larger range in displacement than optimized projection allocation with displacement binning. To illustrate this further, in Figure 5, we have displayed the first 60 seconds of the breathing trace for patient 2 scan 1. We can see that optimized projection allocation with displacement binning sources projections from a narrow band in displacement (usually between 9.9mm and 10.7mm) with a few projections sourced from the neighbouring respiratory bins outside this band (between 9.4mm and 11mm in this case). However, equispaced phase binning, which is currently used clinically, sources projections from a much wider displacement band (usually 9.4 to 11.1mm) with some projections sourced from 7.9mm to 12.4mm.

Figure 5 also highlights the difficulties encountered calculating phase. In Figure 5 we have calculated phase as rising linearly from 0 to 100 from peak inhale to peak inhale. A different result will be obtained if phase rises from 0 to 50 from inhale to exhale then 50 to 100 from exhale to inhale as a different amount of time is spent in inhale versus exhale. It is difficult to accurately identify the peak inhale and exhale points in many of the breathing cycles because images are only acquired at a rate of 10hz. There are several small short breathing cycles that can be identified as either short breathing cycles or irregularities within a cycle.

In Table 3 the results in brackets are from the digital XCAT phantom. The values in the mean column represent the mean error in the marker volume when compared to the known marker volume. A similar pattern to the patient data emerges , i.e. projection sharing with displacement binning performs consistently well in comparison to all other methods. We also note, that when the mean marker volume error is low the standard deviation in the marker volume is also low; this gives us confidence that the standard deviation in the marker volume is a useful metric.

Table 3. Segmented marker volumes. The mean and standard deviation measured in voxels across the 10 respiratory bins for the cases where segmentation was successful for the patient data. Failed is the number of respiratory bins for which the segmented location had an error greater than 3mm. The simple heuristic allowed respiratory bins to shrink or grow by 50% and projections to be source from the entire neighbouring respiratory bin. The values in brackets are the mean marker volume error from the digital XCAT phantom and marker volume standard deviation.

	Phase Binning			Displacement Binning		
Binning Method	Failed	Mean (Error)	Standard Deviation	Failed	Mean (Error)	Standard Deviation
Patient 1 scan 1						
Equispaced	0	20 (49)	2.1 (45)	2	22 (48)	8.4 (42)
Equal-Density	0	21 (47)	1.8 (75)	2	22 (67)	7.0 (106)
Simple heuristic	0	20 (29)	2.6 (40)	0	19 (25)	1.7 (32)
Patient 1 scan 2						
Equispaced	0	25 (31)	7.4 (19)	2	46 (84)	30.0 (84)
Equal-Density	0	22 (28)	3.0 (17)	3	45 (81)	27.8 (104)
Simple heuristic	0	20 (24)	3.6 (19)	0	22 (26)	5.2 (33)
Patient 2 scan 1						
Equispaced	3	35 (12)	9.3 (16)	3	39 (127)	13.9 (77)
Equal-Density	4	40 (14)	8.1 (17)	6	32 (100)	20.1 (91)
Simple heuristic	2	39 (28)	7.0 (30)	0	41 (19)	7.0 (12)
Patient 2 scan 2						
Equispaced	0	15 (38)	2.6 (38)	1	15 (76)	3.6 (86)
Equal-Density	0	15 (50)	1.9 (70)	1	14 (160)	4.1 (312)
Simple heuristic	0	15 (61)	2.2 (77)	0	16 (12)	1.3 (19)
Patient 3 scan 1						
Equispaced	0	16 (52)	10.5 (59)	1	11 (32)	2.1 (45)
Equal-Density	0	13 (38)	4.3 (46)	1	11 (55)	2.4 (90)
Simple heuristic	0	13 (82)	6.1 (112)	0	12 (25)	1.2 (30)

5. Discussion

Significant improvements in image quality have been observed using displacement binning with optimized projection allocation. In clinical practice phase based binning is usually preferred because of data sufficiency problems and streak artefacts present in the displacement binned images. For 4DCT it has been demonstrated that displacement binning is more accurate, contains fewer motion artefacts and recovers the tumour size and shape better than phase binning (Abdelnour et al. 2007), (Fitzpatrick et al. 2006) and (Li et al. 2012). Our algorithms confirm this result for 4DCBCT imaging and allows the user to establish a trade-off between reducing streak artefacts and increasing motion blur. Our algorithms also presents a pathway to reliably generate low streak displacement binned 4DCBCT images.

Streak artefacts significantly degrade the quality of deformable image registration. Applications using deformable image registration, such as lung ventilation studies, are likely to benefit using the optimized projection allocation algorithm. Additionally, several CBCT iterative reconstruction techniques utilise deformable image registration

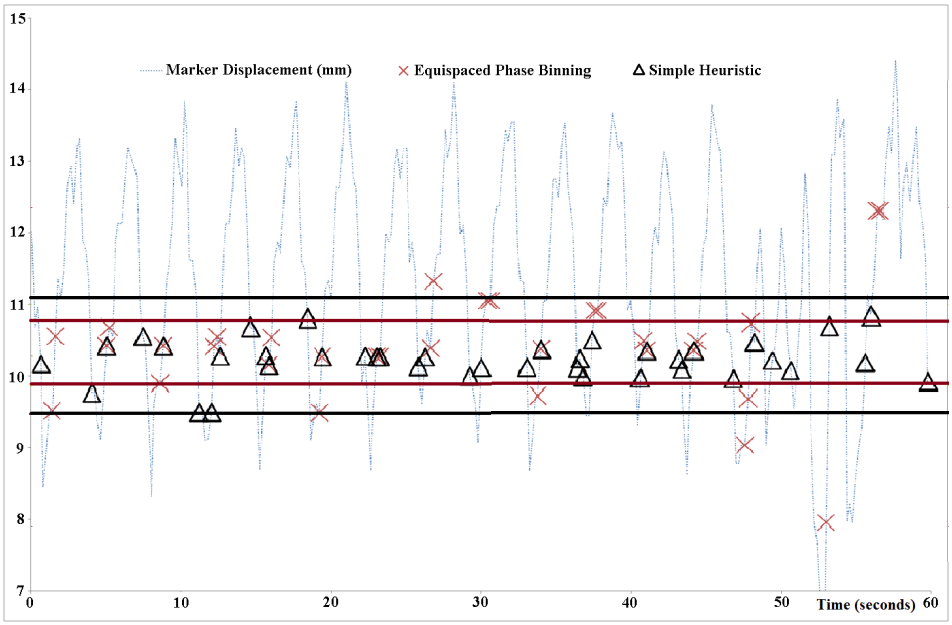

Figure 5. Phase binning versus the corresponding displacement bin using optimized projection allocation for patient 2 scan 1. The first 60 seconds of the patients respiratory signal is given together with the projections taken using phase binning (red crosses) and optimized projection allocation (black triangles) for the corresponding respiratory bin. Although optimized projection allocation binning sources some projections from the neighbouring respiratory bin, there is a smaller range in displacement than for phase binning.

either between prior images or between respiratory bins; these techniques will also benefit from the optimized projection allocation algorithm.

Further progress needs to be made in several areas. From an optimisation point of view solving the equations to optimality would give a ground truth to benchmark the faster, but not guaranteed optimal, heuristic solution methods. Further development on fast heuristic methods can be made, in particular, a heuristic solution method that can run during projection acquisition, that updates as projections are acquired, would be a useful development so that 4DCBCT images could be available as soon as the scan finishes. Further 4DCBCT images need to be examined to determine how far into neighbouring respiratory bins we source projections.

Sourcing projections from a large distance into the neighbouring bin is likely to produce images with fewer streaks but the trade-off is that there will be more motion blur in the images. In this study we have chosen to control streaking and motion blur by limiting the amount that bins can grow/shrink and limiting the distance that projections can be sourced from the neighbouring bins. An optimal trade-off between image quality and motion blur should be established with motion blur taken into account in the objective function. This could be achieved by calculating an estimate of motion blur from the projections themselves, or respiratory signal, and including this measure in the objective function with an appropriate weight. Placing additional constraints on

the motion blur and streak width (i.e. limiting both values to 3mm) would enable the user to control the accuracy of the reconstructed images.

For some clinical applications there may be a requirement that respiratory bins are located at specific locations rather than the optimized locations presented in this work. Deformable image registration may be necessary to recover images at the specified locations. With better image quality there is the potential for reducing both imaging time and imaging dose (fewer projections) and further studies need to be performed to determine how much the imaging time and imaging dose can be reduced.

6. Conclusions

We have developed an optimized 4DCBCT projection allocation algorithm that leads to improvements in 4DCBCT images. The improvement in image quality using our optimized projection allocation approach is more significant when using displacement binning than phase binning. Our results suggest that optimized projection allocation using displacement binning has the potential to produce better 4DCBCT images than phase based binning, which opens the door for the clinical use of displacement binning to improve image quality and patient positioning in radiotherapy.

7. Acknowledgements

The authors would like to acknowledge the support of a National Health and Medical Research Council (NHMRC) Australia Fellowship. This project was supported in part by NHMRC project grant 1034060. The authors would like to thank Fair Isaacs Corporation (FICO), IBM-ILOG and GUROBI optimization for academic licenses to XPRESS-MP, CPLEX and GUROBI respectively.

Appendix A. Linearizing the standard deviation

The equations presented in this section can be used to linearise the standard deviation for use with standard Mixed Integer Programming solvers. A solution to these equations represents the best possible value of the standard deviation while the heuristics presented in the main text are fast but do not produce a provably optimal solution. The standard deviation in bin b (σ_b) can be expressed in linear form as

$$\sigma_b = -4\pi^2 - 2\theta_{bsl} + 4\pi(\theta_{bs} - \theta_{bl}) - 2\pi\delta_{\theta,b} + \sum_j \left(2\delta_{b,j}\theta_j - 2\theta_{b,j}\right),$$

with

$$\begin{aligned}
\theta_{b,k} &\leq M\delta_{b,k}, \\
\theta_{b,k} &\leq M\sum_{j>k}\delta_{b,j}, \\
\theta_{b,k} &\leq M(2 - \delta_{b,k} - \delta_{b,j}) + \theta_k\theta_{k'},
\end{aligned}$$

$$
\begin{aligned}
\theta_{bl} &\leq \theta_k + M\sum_{j>k}\delta_{b,j},\\
\theta_{bl} &\geq \theta_k\delta_{b,k},\\
\theta_{bs} &\leq \theta_k + M(1-\delta_{b,k}),\\
\theta_{bs} &\geq \theta_k - M\sum_{j<k}\delta_{b,j},\\
\theta_{bsl} &\geq \theta_{bl}\theta_k + M(1-\delta_{b,k}),\\
\theta_{bsl} &\geq \theta_k - M\sum_{j<k}\delta_{b,j},\\
\delta_{\theta,b} &= 2\pi + \sum_j C_{b,j},\\
C_{b,k} &\leq \delta_{\theta,b},\\
C_{b,k} &\leq 2\pi\delta_{b,k},\\
C_{b,k} &\geq \delta_{\theta,b} + M(\delta_{b,k}-1),
\end{aligned}
$$

for all k and k' and M is a suitably large number which in this case can be $4\pi^2$. By programming these equations into a Mixed Integer Programming solver the optimal solution for the standard deviation can be found.

References

Abdelnour A F, Nehmeh S A, Pan T, Humm J L, Vernon P, Schoder H, Rosenzweig K E, Mageras G S, Yorke E, Larson S M & Erdi Y E 2007 Phase and amplitude binning for 4D-CT imaging *Phys Med Biol* **52**(12), 3515–29.

Armitage S E, Pollmann S I, Detombe S A & Drangova M 2012 Least-error projection sorting to optimize retrospectively gated cardiac micro-CT of free-breathing mice *Med Phys* **39**(3), 1452–61.

Bergner F, Berkus T, Oelhafen M, Kunz P, Pa T, Grimmer R, Ritschl L & Kachelriess M 2010 An investigation of 4D cone-beam CT algorithms for slowly rotating scanners *Med Phys* **37**(9), 5044–53.

Bergner F, Berkus T, Oelhafen M, Kunz P, Pan T & Kachelriess M 2009 Autoadaptive phase-correlated (AAPC) reconstruction for 4D CBCT *Med Phys* **36**(12), 5695–706.

Chen G H, Tang J & Leng S 2008 Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets *Med Phys* **35**(2), 660–3.

Cooper B J, O'Brien R T, Balik S, Hugo G D & Keall P J 2013 Respiratory triggered 4D cone-beam computed tomography: a novel method to reduce imaging dose *Med Phys* **40**(4), 041901.

Fast M, Wisotzky E, Oelfke U & Nill S 2013 Actively Triggered 4d Cone-Beam CT Acquisition *Medical Physics* **40**(9), 0.

Feldkamp L A, Davis L C & Kress J W 1984 Practical Cone-Beam Algorithm *Journal of the Optical Society of America a-Optics Image Science and Vision* **1**(6), 612–619.

Fitzpatrick M J, Starkschall G, Antolak J A, Fu J, Shukla H, Keall P J, Klahr P & Mohan R 2006 Displacement-based binning of time-dependent computed tomography image data sets *Med Phys* **33**(1), 235–46.

Leng S, Zambelli J, Tolakanahalli R, Nett B, Munro P, Star-Lack J, Paliwal B & Chen G H 2008 Streaking artifacts reduction in four-dimensional cone-beam computed tomography *Med Phys* **35**(10), 4649–59.

Li H, Noel C, Garcia-Ramirez J, Low D, Bradley J, Robinson C, Mutic S & Parikh P 2012 Clinical

evaluations of an amplitude-based binning algorithm for 4DCT reconstruction in radiation therapy *Med Phys* **39**(2), 922–32.

Mckinnon G C & Bates R H T 1981 Towards Imaging the Beating Heart Usefully with a Conventional Ct Scanner *Ieee Transactions on Biomedical Engineering* **28**(2), 123–127.

Nemhauser G & Wolsey L 1988 *Integer and Combinatorial Optimization* John Wiley and Sons.

O'Brien R T, Cooper B J & Keall P J 2013 Optimizing 4D cone beam computed tomography acquisition by varying the gantry velocity and projection time interval *Phys Med Biol* **58**(6), 1705–23.

O'Brien R T, Cooper B J, Kipritidis J, Shieh C & Keall P J 2014 Respiratory motion guided four dimensional cone beam computed tomography: Encompassing irregular breathing *Phys Med Biol* **59**(3), 579–95.

Roman N O, Shepherd W, Mukhopadhyay N, Hugo G D & Weiss E 2012 Interfractional positional variability of fiducial markers and primary tumors in locally advanced non-small-cell lung cancer during audiovisual biofeedback radiotherapy *Int J Radiat Oncol Biol Phys* **83**(5), 1566–72.

Segars W P, Sturgeon G, Mendonca S, Grimes J & B.M.W T 2010 4D XCAT phantom for multimodality imaging research *Medical Physics* **37**(9), 4902–4915.

Shieh C, Kipritidis J, O'Brien R T, Cooper B J, Kuncic Z & Keall P J 2014 Image quality in thoracic 4D cone-beam CT: A sensitivity analysis of respiratory signal, binning method, reconstruction algorithm, and projection angular spacing *Medical Physics* **41**(4), 041912.

Sidky E Y & Pan X 2008 Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization *Phys Med Biol* **53**(17), 4777–807.

Sonke J J, Zijp L, Remeijer P & van Herk M 2005 Respiratory correlated cone beam CT *Medical Physics* **32**(4), 1176–1186.

Taguchi K 2003 Temporal resolution and the evaluation of candidate algorithms for four-dimensional CT *Medical Physics* **30**(4), 640–650.

Talbi E G 2009 *Metaheuristics: From Design to Implementation (Wiley Series on Parallel and Distributed Computing)* John Wiley and Sons. Hoboken, New Jersey.