

Binning Without a Model for Cone-beam CT

Tracy Petrie¹
scstp@leeds.ac.uk

Derek Magee¹
D.R.Magee@leeds.ac.uk

Jonathan Sykes²
Jonathan.Sykes@leedsth.nhs.uk

¹ School of Computing
University of Leeds
Leeds, LS2 9JT, UK

² Medical Physics and Engineering
Leeds Teaching Hospitals Trust
St James's Hospital
Leeds, LS9 7TF, UK

Abstract

We present the first results of a new technique to bin Cone Beam projections without imposing any motion model. Such a technique is required for studying motion in regions of the body, such as the pelvis, where motion exists and is unpredictable. All motion information is obtained directly from the projections and the binning is performed through a type of best first search through the graph of possible complete assignments. Simplifying assumptions coupled with loss-less dimensional reduction using Principal Component Analysis, make the method tractable.

1 Introduction

Cone-beam CT (CBCT) is frequently used in image-guided radiotherapy (IGRT) to verify patient position and the validity of the treatment plan with respect to the planning CT. Approximately 660 radiographs are typically acquired during a two minute scan and a standard filtered backprojection algorithm [1] is used to reconstruct a volume. This volume can then be rigidly registered with the planning CT to evaluate changes in the treatment area and potential misalignments. Such evaluations are needed to dynamically adjust the treatment to correctly irradiate the target tissue and avoid healthy tissue. A critical issue with this process is the introduction of blur due to the long acquisition time. This blur makes it difficult for clinicians to assess alignment and see changes in tissues over the course of the treatment plan. Several approaches to mitigating this problem have been proposed or implemented but most rely on some form of radiograph, or projection, binning. Usually, some form of model is used to guide this binning and most often this is a breathing model where the respiratory motion is divided into phases and projections are then binned by phase (see [2][3]). Projections are matched to phases in many different ways: measuring respiration directly, estimating it from diaphragm detection in the radiographs, estimating it from external camera views of the patient, etc. An inherent limitation is that without a model, these methods fail. Few approaches have attempted binning outside of the lung region.

Our approach is to attempt to bin projections without any model at all. We extend our previous work [4] which implements an exhaustive search of possible binning assignments. One of the requirements of that method is a two-pass protocol. In this paper we remove the two-pass protocol change and replace the exhaustive search with a stochastic search.

Removing the protocol change allows our ideas to potentially work with existing CBCT configurations and, indeed, with existing data for retroactive studies in the future. These studies could be used to better understand tumour motion in many parts of the body that are currently not studied. We show early results with only two bins but demonstrate on synthetic and phantom data that both periodic motion and non-periodic motion can be estimated with visible and measurable improvements




Figure 1: ROI images from synthetic model with motion in three dimensions. (a) Normal reconstruction (b) ideal bin 1 (c) ideal bin 2 (d) estimated bin 1 (e) estimated bin2.




Figure 2: ROI images from QUASAR phantom animated with simulated single shift motion. (a) Normal reconstruction (b) ideal bin 1 (c) ideal bin 2 (d) estimated bin 1 (e) estimated bin2.

2 Methods and Materials

Our method poses bin assignment as a search problem. It involves grouping projections under the assumption of binning similarity in temporally adjacent projections, constructing reprojections to fill in data gaps, recasting reconstruction as an averaging process of individual backprojections, reducing the backprojection size by clipping to a region of interest, and further reducing the size through the use of principal component analysis. Computing reconstructions and fitness metrics in eigenspace decreases computational cost, but limits us to metrics that have meaning in both eigenspace and the original space. We demonstrate promising results by combining two such metrics.

Bucketing: the likelihood of adjacent projections belonging to the same bin is high when binning into two bins. We group projections into one-second buckets which reduces our binning task to 120 items, each containing between five and six projections, instead of 660.

Scoring: To assess the fitness of a given assignment, two volumes are reconstructed based on a hypothesised binning and evaluated. At the voxel level, correctly binned reconstructions should, on average, be different from each other reflecting the fact that they represent the two states of the moving tissue. Globally, however, the reconstructions should be similar in terms of greylevel distribution (as they are of roughly the same material). The latter criteria is necessary as it is possible to reconstruct two highly different volumes due

to one containing mainly light voxels and the other containing dark voxels. This type of difference we avoid by minimising the global difference while maximising the voxel differences. The difference between the greyscale means ($d_{mean} = \bar{V}_1 - \bar{V}_2$) of the two volumes is assumed to be Gaussian with zero mean (ie. decreases with increasing d_{mean}). We estimate the variance of the Gaussian by sampling reconstructions from our search space. We multiply this probability by the sum of squared differences of the two binned volumes to obtain our fitness score ($score = G(d|0, \theta) \times SSD(V_1, V_2)$). We also, optionally, add a penalising term if too many buckets are assigned to the same bin. If we know that we are imaging a region of the body which is likely to have periodic motion, we can make a simple guess at the period and this is sufficient to help the search process without requiring a precise model of the pattern.

Searching: Even with the bucketing reduction, 2^{120} is still too large to search exhaustively. We implemented a Best- N -First method that takes N candidate assignments (of buckets to bins), and generates child assignments by flipping each of the individual bucket assignments and evaluating the result. The best N assignments which are better than the parent’s score are kept then the best N assignments out of all the children generated are kept. Parents are added to a retired list and on each iteration, new children are vetted for uniqueness among their peers and the retired list. The process terminates when fewer than N total new children are generated.

Regions of Interest (ROIs): Reconstruction volumes are large and mostly contain information we don’t need from a binning standpoint. What we’re really interested in is the motion that occurs in the region labelled in the planning CT as the "planning treatment volume". We simulate this ROI in our experiments with cylindrical regions but nothing prevents the use of oddly shaped planning volumes. We use these ROIs in two ways: to generate 2D masks of the projections by projecting the ROIs onto a virtual detector using the same geometry as the Synergy machine, and using them directly as 3D masks.

Reprojections: A significant problem in reconstruction is the introduction of artefacts. Two sources of artefacts that impact us directly are motion artefacts and missing projection artefacts. When we bin projections, reconstructing one bin means all the projections belonging to other bins are missing. This creates severe artefacts in the form of streaks and misshapen structures. Conversely, retaining all the projections introduces the same kinds of artefacts if there is motion, which the whole effort is predicated on. We mitigate this problem by constructing filler projections. By taking the volume reconstructed from the original projections (V_{all}), we can reproject new projections that hypothetically contain the blur. We apply a small median filter to V_{all} before reprojecting to attempt to remove any motion incorporated into this reconstruction. Furthermore, to eliminate the introduction of motion artefacts from outside the ROI, we synthesize merged projections consisting of the reprojection outside the edge-blurred 2D ROI mask and the original projection inside the 2D ROI.

Individual Backprojections: Having the set of merged and reprojected projections at matching acquisition angles, we reconstruct a given bin volume by choosing the merged projections that are assigned to the bin and using the reprojected projections to fill in the missing projections. Given the many thousands of reconstructions that are needed by the search process, we factor out the backprojection part of the process. Typical filtered backprojection algorithms [1] take each filtered projection and backproject them across a single volume. This is equivalent to backprojecting them into individual volumes and subsequently averaging these "backprojection volumes". This preprocessing allows us to reconstruct from multiple hypothesised binnings merely by averaging the relevant backprojection volumes

together.

Reconstruction in eigenspace: If we apply Principal Components Analysis (PCA) to the ROI-clipped backprojection volumes, we obtain a set of low-dimensional eigenspace projections of the backprojection volumes. Taking the mean of a set of vectors in this space is the same as taking the mean of the original vectors, after projecting back into the original space. Likewise, Euclidean distance between vectors is preserved. This means we can easily calculate the sum of squared differences between volumes in eigenspace (as SSD and Euclidean distance in the original space are equivalent). Finally, if we take the dot product of a vector of the means of the eigenspace basis vectors (which may be pre-calculated) and any vector in the eigenspace, we obtain the mean of that vector in our original space (i.e. the greylevel mean). This is a key property because it allows us to calculate the mean of a reconstruction without first projecting it back into the original space. Coupled with the SSD property, we are able to generate, in eigenspace, the fitness score. Using PCA, we can take vectors that contain tens of thousands of elements and truncate them – without loss – to vectors of length ≈ 1320 (the number of merged projections plus the number of reprojections). Besides making reconstruction fast enough for our purposes, PCA also allows us to evaluate our reconstructions in eigenspace. The two averaged vectors that constitute two binned reconstructions can be subtracted from each other and the resulting dot product gives us the sum of squared differences in the original space. Likewise, by taking the dot product of the vector with the means of the eigenspace column vectors, we obtain means of the vectors in the original space.

We show results from four experiments, two synthetic and two animated. In the synthetic case, we construct mathematical phantoms composed of ellipses, cubes, and cylinders and project them using the same geometry as the Synergy machine. In the first experiment, we simulate a sinusoidal motion along the z-axis, the axis of rotation. In the second experiment, we include motion in the x and y axis to create a curved motion.

We now introduce a new method for generating motion from a physical phantom. A QUASARTMrespiratory motion phantom with a wood cylindrical insert containing an acrylic sphere was scanned in 16 static positions with the insert being moved 1mm in the z-axis direction each time. A simple utility allows us to literally draw the z-axis motion we want over time. The correct projections from the 16 sets of projections are then copied to simulate the motion in a manner similar to “stop motion animation” techniques. In the first experiment, we simulate a fast breathing pattern with 22 breaths in the two minute duration. In the second experiment, we simulate a single shift in the sphere representing the type of motion that might occur in the prostate region.

3 Results

In each experiment, we construct the motion and so can establish the correct binning assignment and compare the results of our search estimation against the “ideal” binning volumes. To calculate the correct assignment, we take a representative point in motion, e.g. the centre of the moving sphere, and apply k-means clustering ($k=2$). Ideal binning volumes (V_{gt1}, V_{gt2}) are generated from this binning in the mannner described previously. In Figures 1 and 2, the success of our approach is clear. Table 1 quantifies this by comparing the mean sum of normalised absolute differences between matched volume pairs. Estimated volumes are matched so as to minimise this difference metric. The differences between the estimated volumes and ideal binning volumes is significantly smaller than between the estimated bins and the normal (blurry) reconstruction volume constructed by averaging all backprojection

Comparison	1	2	3	4
Ideal (V_{gt_1}, V_{gt_2}) vs. Normal (V_{all})	12.2	12.4	4.19	3.57
Estimated (V_{est_1}, V_{est_2}) vs. Normal (V_{all})	12.0	12.4	4.28	3.64
Estimated (V_{gt_1}, V_{gt_2}) vs. Ideal (V_{gt_1}, V_{gt_2})	2.87	.906	.662	2.82

Table 1: Comparison of reconstructions for four experiments. All comparisons are mean normalised sum of absolute differences (over matched volume pairs). Column 1 shows the synthetic z-axis only experiment; column 2 shows the synthetic phantom with three dimensions of motion; column 3 shows the two-state wood phantom results; column 4 shows the 22-breath phantom simulation results.

volumes. The differences between the ideal bins and the normal reconstruction is shown for comparison. Of interest is the fourth column; it reveals that our current algorithm still has room for improvement.

4 Discussion

We demonstrate initial feasibility of a new binning method requiring no motion model on synthetic and physical phantoms for the two-bin problem. We also illustrate a technique for simulating motion using a real phantom that doesn't require additional scans. The principle limitation of this new technique is that it is constrained to 1D motion. However, it is useful for exploring non-periodic types of motion in lower abdominal regions of the body. There are two potential approaches to extending the method to greater than two bins (and hence non-linear motion). Firstly, the number of bins in the search assignment could be increased. However, as the search space has size N^M (where N is the number of bins, and M is the number of buckets) this would greatly increase the size of the search space. The alternative approach is to perform multi-step binary binning; splitting each bin recursively. This has (approximately) $O(2)$ complexity in the number of bins and is thus more feasible. It also has the advantage of using a similarity based stopping criteria to determine the number of bins required. We are currently working on developing this approach. Subsequently we will attempt to register the reconstructed volumes to obtain a motion model. If successful, this technique has the potential to enable large retrospective motion studies on the set of CBCT projections already in existence.

References

- [1] Avinash C. Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*. Classics in applied mathematics ; 33. SIAM,, Philadelphia, 2001.
- [2] T. Petrie, D.R. Magee, and J.Sykes. Modeless motion binning in cone beam ct. *ISBI 2010. IEEE International Symposium on*, pages 1245–1249, April 2010.
- [3] David Sarrut, Bertrand Delhay, Pierre-Frédéric Villard, Vlad Boldea, Michael Beuve, and Patrick Clarysse. A comparison framework for breathing motion estimation methods from 4-d imaging. *IEEE Trans Med Imaging*, 26(12):1636–1648, Dec 2007.
- [4] J. J. Sonke, L. Zijp, P. Remeijer, and M. van Herk. Respiratory correlated cone beam ct. *Medical Physics*, 32(4):1176–86, 2005.