EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling

Parametric Models for Motion Correction in Medical Imaging

Joseph Field

Contents

1. Introduction

Medical imaging techniques, such as radiography and computed tomography (CT), are key diagnostic tools in many areas of healthcare, and are paramount for disease detection and for patient monitoring during ongoing care. In the U.S.A., for example, 70% of the population are expected to undergo a medical or dental X-ray procedure each year. Due to advancements in digital imaging, it is now possible to produce images with fine anatomical detail, allowing disease detection to be made much earlier than before.

A crucial development in X-ray imaging is digital tomosynthesis (DTS), in which a 3D image is reconstructed from a set of separately-captured 2D measurements. Over the last decade, a lot of effort has been directed at improving the accuracy and computational efficiency of the algorithms used to reconstruct DTS measurements. Standard DTS acquisition methods (see Figure 1) involve a single X-ray tube and a paired flat-panel detector, which are computer-controlled to move along predefined paths. The X-rays are emitted and detected, before reconstruction algorithms are employed to create three-dimensional section images. The X-ray intensity measurements obtained by the detector panel are used to rebuild the geometry of the body being scanned, as well as its density.

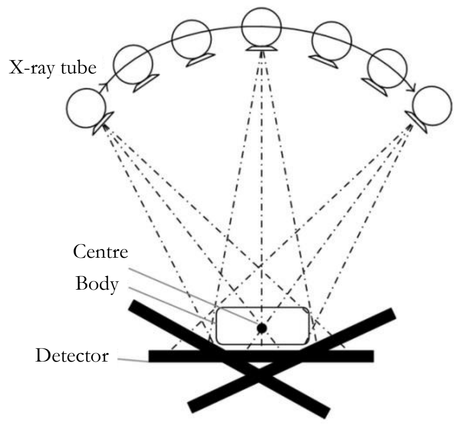

Figure 1: In a standard DTS acquisition, a set of two-dimensional images is produced by moving an X-ray emitter and tilting the detector around the body. Three-dimensional reconstructions can then be obtained by analysing the set of images together.

Adaptix have developed a portable, low-cost X-ray source consisting of multiple low-power emitters arranged in an array, to be used in conjunction with a flat-panel detector. With this device it is possible to fire a sequence of low-dose X-rays, and use the resulting intensity measurements to produce 3D image reconstructions. By using an array of smaller emitters, it is possible to reconstruct 3D images with a radiation dose similar to a single radiography procedure.

Patient movement during measurement acquisition may create blurring in the reconstructed 3D images. To ensure that image accuracy is maximized, all motion should be corrected for in the reconstruction.

During a complete firing sequence of the X-ray emitters there may be small body movements, either from small fluctuations in a patient's resting position, or from internal changes, e.g. if a patient is unable to hold their breath during the capture process. The 3D reconstructions may therefore be blurred. With this in mind, Adaptix wish to investigate the reconstruction of images which are affected by motion. It would first appear that reconstructions from individual emitters can be superimposed to obtain a complete image, but this requires a perfectly-trackable reference point within the body, which is not often achievable in medical imaging.

Our approach is to implement and review two possible methods for motion correction. In one method, we use information captured over time to build a single 3D image. In the other method, we treat temporal information in such a way that we reconstruct a video of a 3D image moving over time.

2. Formulating the Problem

The physical device that Adaptix have developed consists of a pair of emitter and detector panels that are positioned opposite one another. Each emitter flashes a conical beam of X-ray radiation, with all detectors within the base of the cone recording a particular measured X-ray intensity. The number of emitters is far fewer than the number of detectors, and the body of interest is positioned closer to the detector panel to increase the amount of cone coverage – see Figure 2.

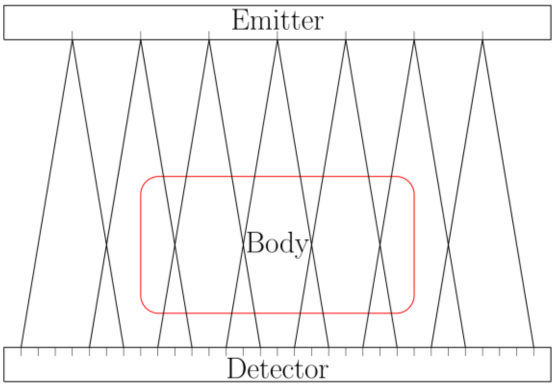

Figure 2: Schematic showing a configuration for seven emitters and 32 detectors, with the body placed closer to the detector for increased cone coverage.

Glossary of terms

- Sparsity: The proportion of nonzero entries in a matrix is known as its sparsity, such that if a matrix is mostly comprised of zeros then it is said to be sparse.
- Discretisation: The process of converting continuous information into discrete counterparts is known as discretisation. For example, we divide time into discrete parts (seconds, hours, days) even though time is continually progressing.
- Compressed Sensing: Rather than capturing a huge amount of data and then compressing it afterwards, compressed sensing combines these two ideas by sampling a small amount of information from a signal while still being able to reconstruct the original signal with great accuracy. Compressed sensing relies on the assumption that the vector of signal information is sparse.
- Wavelets: Wavelets are wave-like functions that can be used to represent signals using limited information. Wavelets often have specific properties that make them useful for signal processing.
- Voxels: A voxel is a volume element in a 3D image, much like how a pixel is a picture element in a 2D image.
- Phantom: A phantom is an object that is used in the field of medical imaging to evaluate and analyse the performance of an imaging device.

Mathematical model

We assume that, during measurement acquisition, no overlap occurs between radiation cones from different emitters. This is akin to assuming that all emitters flash at separate times, or that emitters which flash at the same time are positioned such that their radiation cones do not cross.

In our model, the total volume to be measured, referred to as the *scan domain*, is taken to be a unit cube. We discretise the scan domain into cubic voxels, and we assume that each voxel has a constant density value. In particular, we model a cubic phantom as a body with a density of 1, with the body being positioned symmetrically in the scan domain, while the remaining voxels have a density of 0.

We model emitters and detectors as points, and we model the cone of radiation produced by an emitter as a set of rays starting from the point emitter which then pass through the point detectors contained within the base of the cone – see Figure 3. The cones of radiation are assumed to be identical for each emitter, and are entirely defined by their *collimation angle*, i.e. the complete open angle from which rays can be emitter. To agree with the design of the physical device, the collimation angle of each emitter is chosen to be 20°.

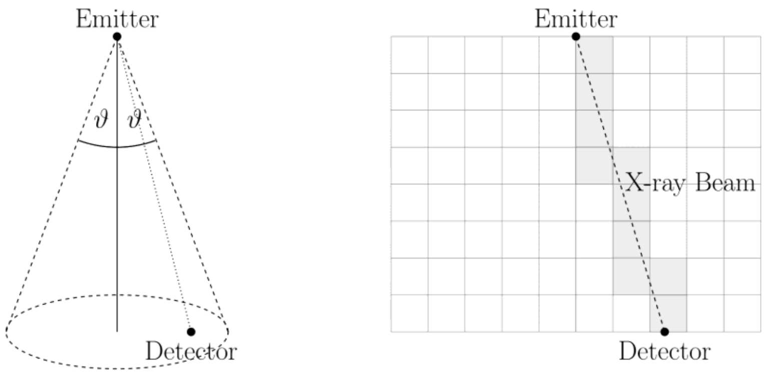

Figure 3: (Left) Schematic of a cone with a collimation angle of 2ϑ. (Right) A single ray is traced through the voxels that it crosses, represented here by two-dimensional pixels.

A single measurement is obtained by recording the X-ray intensity at detector that lies within the base of the radiation cone. We assume that all emitters have the same intensity, and that the detected intensity is less than or equal to this due to interactions such as scattering and body absorption. We do not model the effect of scattering, and so we assume that all reductions in intensity are due to absorption. If the ray passes only through material of density zero, then the ray intensity does not decrease.

We use the *YALL1* algorithm to find the reconstruction which best satisfies the system $\boldsymbol{Ax} = \boldsymbol{b}$, while satisfying the sparsity assumption

Ray attenuation is governed by two factors - the attenuation coefficient of the body that the ray passes through, which is related to the density of the material, and the length of the path along which the ray travels. DTS reconstructions are made from measurements following the Beer-Lambert Law for X-ray absorption, with each ray requiring a separate equation. All rays in the model are therefore represented by a system of linear equations, $\boldsymbol{Ax} = \boldsymbol{b}$. The number of columns in $\boldsymbol{A}$ reflects the number of voxels in the model, while the number of rows represents how many measurements are made. From the standpoint of compressed sensing, we wish to reconstruct the information about all voxels without the need for a huge number of measurements, and so we assume that $\boldsymbol{A}$ has fewer rows than columns. We also assume that the solution to reconstruct is sparse, i.e. the vector $\boldsymbol{x}$ does not have many nonzero entries. The system $\boldsymbol{Ax} = \boldsymbol{b}$ may not have a unique solution, nor will it necessarily have a solution at all, and so we use the $YALL1$ compressed sensing algorithm to find the best fitting solution (under certain additional constraints).

When investigating the effect of movement we assume that the body undergoes simple linear motion at a constant velocity, in the direction of a single coordinate axis.

3. Results

We first investigate the recoverability of static bodies to provide a baseline for how accurately the reconstructions can be resolved. We then move onto reconstructing measurements of moving bodies, using two different approaches, to see whether we can obtain clear reconstruction using either approach.

Static body results

We first run multiple simulations of a static cubic body, to make sure that the simpler, zero-motion model can be reconstructed with reasonable accuracy.

In Figure 4, we show the setup cubic body with density equal to 1, alongside a numerical reconstruction of the same body found using the $YALL1$ algorithm. We find that a cuboid-like body is clearly resolved, but that many voxels have density values far higher than the density of the original setup. The reconstructed solution not only has densities that are over ten times greater than the original setup, but we also observe artefacts at the top and bottom of the scan domain. However, in a medical environment, results are often interpreted by visual judgement, such that we must ask ourselves the question: *Can a body be identified in the reconstruction, even if the values are not entirely correct?*

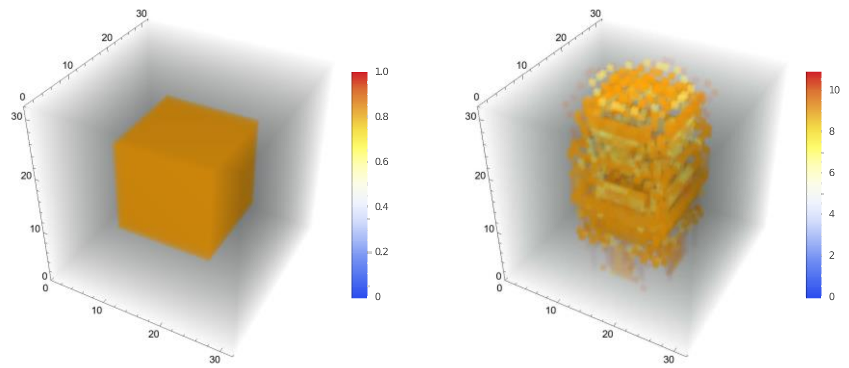

Figure 4: (Left) Original setup of the constant density, cubic body. (Right) Reconstruction obtained from the *YALL1* algorithm, using an emitter/detector ratio of 100/121. The scale in each figure gives the densities of the voxels.

By increasing the number of emitters and detectors in the model, we obtain reconstructions with greater detail. However, we also increase the amount of information being measured, and thus we make the reconstruction problem more computationally expensive.

The number of emitters (and detectors) is taken to be a model parameter, and so we can increase this number if we want to reconstruct over more measurements. The reconstruction in Figure 4 used an emitter/detector ratio of 100/121, while in Figure 5 (Left) we show the results of using a ratio of 576/625. When we use more measurements, we find that the resolution of the cube is increased and the density values better represent the actual solution.

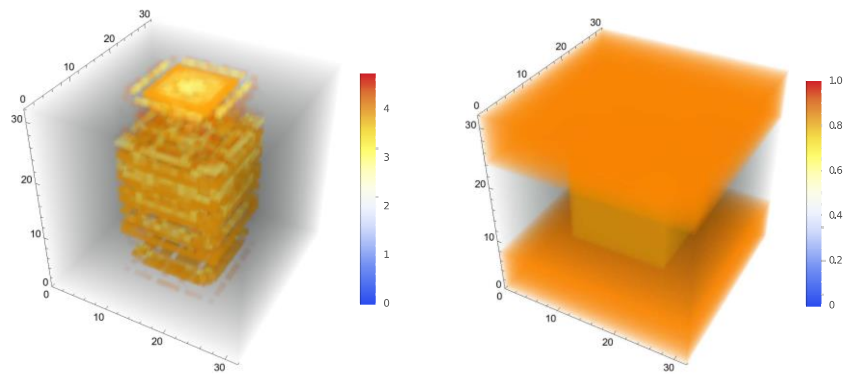

Figure 5: (Left) Reconstruction obtained from the *YALL1* algorithm, using an emitter/detector ratio of 576/625. (Right) Reconstruction obtained from the *YALL1* algorithm, using a wavelet transformation and an emitter/detector ratio of 576/625.

By using wavelets, we find reconstructions with much sharper edge resolution. However, bodies with curved faces are not reconstructed well, and so different wavelet types should be investigated further.

To resolve the edges of the cubic body more clearly, we transform the problem to one in which the body information is represented by wavelets. In this way, we infer density differences between neighbouring voxels, rather than solving for voxel information directly. In this approach, the $YALL1$ solver is able to find the exact boundary of the cubic body – see Figure 5 (Right). Reconstruction artefacts that occur in the top and bottom of the scan domain now appear to be spread out evenly, rather than being concentrated in certain areas, as is the case in Figure 5 (Left). We find that the same wavelets are unable to resolve bodies with curvature, such as a cylinder, although we do not present results here. Thus, the specific choice of wavelets we use may not have been optimal, and in future work we will look at different wavelet types to determine their utility.

We now consider the case in which the body is moved at a given speed through the scan domain, from the 'back' of the domain to the 'front'.

Dynamic body results: Method I

To reconstruct measurements from taken over a finite time interval, we must necessarily divide the measured information over a certain number of frames, e.g. if we look at 10 frames then each frame only contains 10% of the original information. In Figure 6 we show part of a 4D reconstruction without any temporal correlation between frames. We see that, in each frame, there is not enough information to gain any idea of what the phantom looks like.

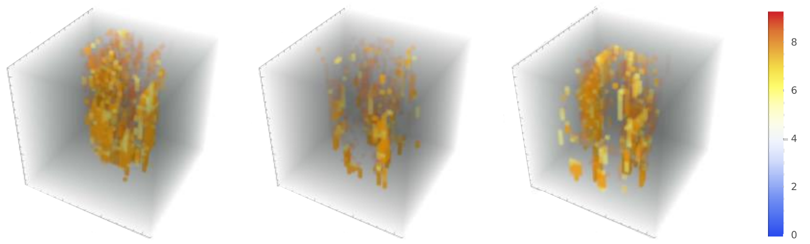

Figure 6: Frames 6, 10, and 14 (of 16). In each frame only a limited amount of information is obtained, and so it is not obvious that the moving body is cubic.

Using a temporal transform, voxel densities can be inferred during frames in which the voxels are not explicitly measured

We exploit the assumption that voxel information between neighbouring frames is correlated, i.e. we expect measurements from two or more sequential frames to share information. Therefore, we transform the problem to solve for voxel densities not only spatially, but also as a function of time. We use two different transforms to connect temporal information of a body undergoing constant linear motion. The first transformation is formed from smooth periodic functions, and the second is formed from 'square-shaped' wavelets that have abrupt jumps in their value. In Figure 7 we show the impact of using a periodic function transformation to obtain additional information.

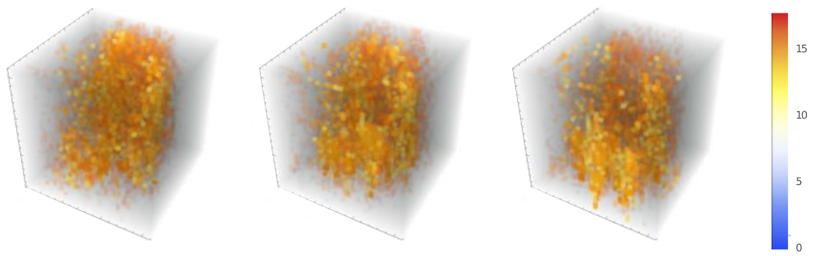

Figure 7: Frames 6, 10, and 14 (of 16). In each frame we are able to reconstruct more of the body than is actually measured, due to the assumption of correlated temporal information.

Dynamic body results: Method II

The second approach we use focuses on correcting for linear motion while also attempting to parametrise the motion of the body, i.e. we wish to determine the velocity at which it moves across the scan domain. We achieve this by translating the voxels of each frame in the sequence so that the voxels of different frames overlap, in the hope of then reconstructing the body more accurately.

We assume that, for maximal reconstruction accuracy, the exact body motion is known, so that the voxels can be perfectly translated – this is a strong assumption to make in a real life situation. Therefore, the numerical algorithm must also make estimates for the body velocity, so that we expect to find the most accurate reconstruction when the exact velocity is determined. Initial numerical tests find that the algorithm that we use does not always converge to the correct velocity, suggesting that the solution to the problem is non-unique.

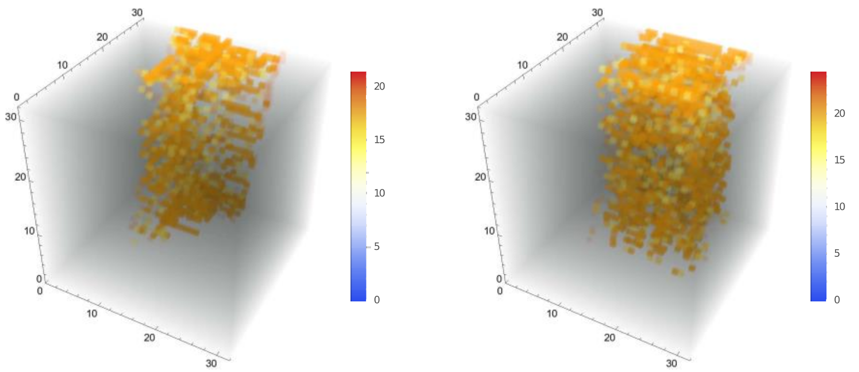

Figure 8: (Left) Reconstruction from 8 frames, with the emitter firing sequence running orthogonally to the body motion. (Right) Reconstruction from 8 frames, with the emitters fired in a random sequence.

Results from the second approach are highly dependent on the order in which the measurements are obtained (see Figure 8) and so more work must be done to identify the optimal emitter sequence to recover maximal information.

4. Discussion, Conclusions, & Recommendations

We have designed and implemented two models for tackling the problem of motion correction during digital tomosynthesis (DTS) measurement acquisition. We used the first model to identify methods for exploiting voxel information that is correlated over time, in order to infer information during frames that certain voxels are not measured. It will be useful to implement a variety of temporal transforms, to determine what the strengths and limitations are of each, and to then maximise the accuracy of the four dimensional reconstruction. We used the second model to not only detect the motion and correct the DTS reconstruction, but also to determine the most likely motion that the body could have undergone. The choice of numerical solver is very important in this approach, and more work must be done to identify the mathematical structure of the problem, to determine if there are solvers that will be particularly useful.

In the immediate future, it will be useful to analyse the recoverability of a body undergoing motion that is not simply linear. We are interested in periodic motions, such as a heart beating or a breathing cycle, and so the first method may work a lot better with a smooth, periodic temporal transform. The second approach will need to be greatly adapted, as the space of possible motions will increase dramatically and so the numerical solver should be chosen accordingly. We have focused on cubic bodies, which are unrealistic as proxies for biological structures. It will be of great benefit to reconstruct bodies that better represent the types of bodies that we hope to detect in real medical situations.

5. Potential Impact

Our results will help to make current reconstruction algorithms more robust to patient motion and body dynamics, in a clinical (as opposed to laboratory) setting. In the long term, this work will enable "functional" information to be obtained during imaging, allowing for the quantification of lung function and perhaps cardiac monitoring.

Gil Travish, Chief Science Officer for Adaptix, remarked that, *"Adaptix are bringing 3D imaging to all the places that currently use 2D, including at the bedside and in the clinic. We have worked for the past three years with Oxford University to create the image reconstruction software. Now, through InFoMM, we are working on motion correction. Patient movement and anatomical processes – breathing, blood flow, heart beats – all can potentially degrade image quality. We are excited by the initial results and these new ideas should allow us to compensate for these motion effects. We hope our future collaboration will extend the utility of our ground breaking X-ray sources and enable new clinical applications. We are excited to see the fruits of this academic work brought to the clinic."*