Fast binary CT using Fourier Null Space Regularization (FNSR)

G. A. Jones[1]‡ and P. Huthwaite

Department of Mechanical Engineering, Imperial College London, South Kensington Campus, London. SW7 2A

E-mail: glennarthurjones@gmail.com

Abstract. X-ray CT is increasingly being adopted in manufacturing as a non destructive inspection tool. Traditionally, industrial workflows follow a two step procedure of reconstruction followed by segmentation. Such workflows suffer from two main problems: (1) The reconstruction typically requires thousands of projections leading to increased data acquisition times. (2) The application of the segmentation process a posteriori is dependent on the quality of the original reconstruction and often does not preserve data fidelity. We present a fast iterative X-ray CT method which simultaneously reconstructs and segments an image from a limited number of projections called Fourier Null Space Regularization (FNSR). The novelty of the approach is in the explicit updating of the image null space with values derived from a regularized image from the previous iteration, thus compensating for any missing projections and effectively regularizing the reconstruction. The speed of the method is achieved by directly applying the Fourier Slice Theorem where the Non-Uniform Fast Fourier Transform (NUFFT) is used to compute the frequency spectrum of the projections at their positions in the image k-space. At each iteration a segmented image is computed which is used to populate the null values of the image k-space effectively steering the reconstruction towards a binary solution. The effectiveness of the method to generate accurate reconstructions is demonstrated and benchmarked against other iterative reconstruction techniques using a series of numerical examples. Finally, FNSR is validated using industrial X-ray CT data where accurate reconstructions were achieved with 18 or more projections, a significant reduction from the 5000 needed by filtered back projection.

Submitted to: *Inverse Problems*

‡ Present address: College of Science, Swansea University, Swansea, SA2 8PP

1. Introduction

X-ray CT is becoming an increasingly common method of non destructively assessing and quantifying material and dimensional properties of industrial components e.g., [1, 2]. Traditionally, industrial X-ray CT workflow consists of a reconstruction followed by image segmentation [1]. Once a segmented image is generated it can be used to assess the specimen for flaws, material properties, quality control and dimensional analysis. However, these workflows typically suffer from time consuming acquisition and a final segmented image which is dependent on the quality of the original reconstruction.

Industrial X-ray CT reconstruction is typically computed using filtered back projection (FBP) or equivalent [3, 4]. The main advantage of FBP methods are its computational speed and low memory requirements. However, FBP methods require hundreds or thousands of projections, each evenly distributed for an accurate reconstruction [3, 4]. These large datasets lead to increased acquisition times particularly in aerospace applications where superalloys are routinely used and long exposure times, on the order of seconds, are necessary to achieve adequate signal-to-noise ratios (SNR). Inevitably these long acquisition times leads to manufacturing bottlenecks which require additional X-ray CT capacity to be purchased or a reduction in the exposure times of the individual X-rays projections, degrading image quality. Additionally, the workflow suffers from the a posteriori application of the segmentation process which is dependent on the quality of the reconstruction and often does not preserve X-ray projection fidelity in the final image e.g.,[5].

An alternative to this two stage workflow is the simultaneous reconstruction and segmentation (SRS) of the image for which a number of methods have been developed. One such approach is to use prior information as a regularization term in the reconstruction [6]. Such priors include sparsity in the gradient of the image which can be exploited by Total Variation (TV) regularization [7, 8], or by assignment of pixels to a given class, using knowledge about the number of materials present [9, 10]. Alternatively, level-set segmentation methods can be incorporated into the reconstruction as shape based priors [5, 11, 12]. Typically the SRS methods are applied to datasets acquired using a full set of projections, which would not lead to reduced data acquisition times.

Binary or discrete tomography deals with image reconstruction from a small number of projections where the pixels are known to be limited to one or a number of distinct values [13, 14]. Unlike conventional tomography where the projections have continuous amplitudes that are related to the distance a raypath travels through an object, in binary tomography the projections are made up of discrete integer sums where only the number of pixels within each material traversed by a raypath is needed. Discrete tomographic problems have been shown to be highly ill posed, are highly unstable in the presence of noise [15] and difficult and complex to solve exactly [13, 14, 16]. The computational difficulty and in particular the highly unstable nature of the problem in the presence of noise means discrete tomography methods cannot be applied successfully to experimental data. However, a number of discrete tomographic methods have been specifically developed to either provide approximate solutions or to handle experimental projection data; examples of such methods include greedy algorithms [17], Monte Carlo based optimization methods [18, 19, 20] network flow methods [21] and iterative hybrid approaches that alternate between a continuous reconstruction using algebraic reconstruction and a discrete reconstruction [22, 23, 24, 25].

Despite significant improvements in computational power and the increasing application of graphics cards to tomographic problems [26], improvements to the resolution of the X-ray detectors and subsequent image resolution continue to make SRS and discrete tomographic methods computationally expensive. The computational expense of these methods renders FBP methods popular due to their speed and small computational footprint.

We introduce a new fast iterative Fourier based CT reconstruction method called Fourier Null Space Regularization (FNSR). The novelty of the approach is in its explicit inclusion of regularization into the null space to compensate for any missing projections. The method achieves its speed by directly applying the Fourier Slice Theorem to compute values for the image k-space corresponding to the projection data. A threshold image is computed at each iteration which is used to populate the reconstruction k-space where no values are available. The population of this null-space is used to regularize and steer the reconstruction towards a binary solution. Numerical experiments are conducted using different binary phantoms to benchmark the computational speed and reconstruction quality of FNSR against other algebraic reconstruction methods as a function of the number of projections. Finally, the method is validated using experimental X-ray data from a turbine blade.

2. Fourier Null Space Regularization (FNSR)

In this section we describe the FNSR method which builds upon the Fourier slice theorem and uses null space regularization to steer the solution towards a desired reconstruction.

2.1. Fourier slice theorem and back projection

The Radon transform relates the X-ray projection data to a line integral through an object $f(x, y)$, at an angle θ to the x axis with an offset from the centre of rotation t [3, 4]. The Radon transform is defined as

$$\begin{aligned}\mathcal{R}f \equiv p(\theta, t) &= \int_{(\theta,t)} f(x, y)\mathrm{d}S \\ &= \int\int f(x, y)\delta(x\cos\theta + y\sin\theta - t)\mathrm{d}x\mathrm{d}y, \qquad (1)\end{aligned}$$

where $t = x\cos\theta + y\sin\theta$ and δ is the Dirac delta function. The projection or sinogram dataset are subsequently a collection of line integrals taken at different offsets and angular rotations [3, 4].

The main objective of CT is to recover the image $f(x, y)$ from the sinogram data by computing the inverse of the Radon transform $(\mathcal{R}^{-1}p)$. A cornerstone in the calculation of $\mathcal{R}^{-1}p$ is the Fourier slice theorem (FST) which relates the 1D Fourier transform of the projection data at an angle θ to the corresponding slice through the 2D Fourier transform of the image of interest [3, 4]. Therefore, the recovery of the object is obtained by performing a Fourier transform of the sinogram slices which are used to populate the corresponding 2D spatial frequency positions. This is then followed by an inverse 2D Fourier transform to give the final image in cartesian coordinates.

The implementation of the FST requires the frequency domain of the projection data to be interpolated from a polar to a cartesian grid in order to compute the inverse 2D transform. Generally the direct implementation of the FST is avoided due to the significant errors introduced during the polar to cartesian grid interpolation in the spatial frequency domain [3]. An alternative formulation of the CT problem, and by far the most widely used is filtered back projection where the sinograms are filtered, then back projected (smeared) through the image. This filtering acts to remove the amplification of particular spatial frequency components introduced by the back projection stage. Provided there are sufficient projections with an even angular sampling the filtered back projection method yields an accurate image of the object [3, 4]. Recent development of the non-uniform sampled FFT (NUFFT) [27] [28] has allowed for the accurate application of the FST with similar computational overheads to the standard 1D and 2D FFT, with results equal to the back projection methods [29] [30].

2.2. Null space regularization

Consider the linear operation $\mathcal{L}$ which acts on an image x to generate some data d, which may be written as $\mathcal{L}x = d$. In X-ray CT, the linear operator is $\mathcal{L}$ is the Radon transform. The goal of an inverse problem is to recover the image thus effectively undoing the operator $\mathcal{L}$. In general, $\mathcal{L}$ is non invertible, so an iterative approach is required to obtain a solution [31] [32]. Additionally, $\mathcal{L}$ is typically ill conditioned which means that an infinite number of null space images $\mathcal{L}x_n = 0$ exist to which $\mathcal{L}$ is effectively insensitive to. The insensitivity of the operator $\mathcal{L}$ to the null space means that any recovered image which matches the data will consist of a combination of the true and null space images. Therefore, any solution to the inverse problem is non-unique where multiple images exist which match the measured data.

It is initially assumed that the values of the null space are zero. The addition of regularization terms allow for the null space to be updated towards more representative values, thus reducing the effects of non-uniqueness on the solution [31] [32]. The type and amount of regularization allows for certain preferential solutions to selected based on some prior assumptions e.g., smooth or edge preserving solutions [6], [7],[31], [32]. The generation of the null space values by regularization may be achieved either implicitly or explicitly. In the former both the data and the regularization terms are fitted and no consideration is given to the null space terms. In the explicit case the null space and measurable components are separated via some appropriate transform with the measurable terms matched with the data whilst the null space components are updated based on a regularized reconstruction.

Deal and Nolet [33] used a 'null space shuttle' to seismic tomographic problems where a non-linear filter is estimated and applied as a post processing step and only modifies components of the null space. Huthwaite *et al.,* [34] and Shi and Huthwaite [35] applied an iterative null space regularization method to limited view ultrasonic imaging, applying a threshold to the image from the previous iteration. This threshold image is used to generate synthetic data for the missing viewing angles which are incorporated with the measured data to perform the next iteration of the procedure. In this sense the threshold image is focusing the data onto the higher contrast areas of the image, attempting to minimise the lower contrast smearing artefacts which appear both in subsampled and limited angle of view problems.

Here, we apply a similar strategy to [34] and [35], but specifically tailored it to limited data X-ray CT. Below we review the theory used by [34] and [35] for ultrasonic imaging. As stated previously the main objective of any imaging approach is to reconstruct an image x, from a set of measurements d by undoing the linear operator $\mathcal{L}$. However, $\mathcal{L}$ is not invertible in general, since there are multiple sets of x which can produce the same data d. We define a corresponding imaging operator $\mathcal{I}$ which can therefore only generate x', an approximation of x,

$$x' = \mathcal{I}\left(d\right). \tag{2}$$

We define an operator $\mathcal{R}$ which maps an image into a regularised image

$$x_{reg} = \mathcal{R}\left(x\right). \tag{3}$$

We do not, at this point, make any assumptions other than that $\mathcal{I}$ and $\mathcal{L}$ are linear. The focus of this paper is to generate images when data is limited, which means that d is not known in its entirety. We therefore split this into two data sets:

$$d = \left\{ \begin{array}{c} d_m \\ d_u \end{array} \right\}, \tag{4}$$

where d_m is the known, measured, data, and d_u is a hypothetical, unmeasurable data set which completes the ideal, full data set. The value of splitting the data in this way is that we can target a solution which matches d_m while using regularisation to estimate d_u. This is achieved through iteration.

An initial image is generated by setting the unknown components d_u to zero, and combining these with d_m as in (4) to obtain the full-view data set at the first iteration

$$d^{(1)} = \left\{ \begin{array}{c} d_m \\ 0 \end{array} \right\}. \tag{5}$$

where we use the superscript (k) to indicate the value of a variable at iteration k. The image is then approximated from this data set as

$$x^{(1)} = \mathcal{I}\left(d^{(1)}\right). \tag{6}$$

The point of regularisation is to steer an image towards the true solution, based on some prior knowledge; the operator can be applied such that $x^{(n)}_{reg}$ should be a better representation of the true object than $x^{(n)}$

$$x^{(1)}_{reg} = \mathcal{R}\left(x^{(1)}\right). \tag{7}$$

This regularisation improvement, encoded within $x^{(1)}_{reg}$, needs to be combined with our measured data set d_m to provide an improved estimate for the unknown values at the next iteration. To do this, we use the forward model to generate a complete data set from $x^{(1)}_{reg}$

$$d^{(1)}_{reg} = \mathcal{L}\left(x^{(1)}_{reg}\right). \tag{8}$$

As before, the components of $d^{(1)}_{reg}$ can be subdivided into two sets depending on whether or not they can be measured from the original limited view array

$$d^{(1)}_{reg} = \left\{ \begin{array}{c} d^{(1)}_{reg,m} \\ d^{(1)}_{reg,u} \end{array} \right\}. \tag{9}$$

This subdivision enables us to identify two components. In general, $d^{(1)}_{reg,m} \neq d_m$, i.e. the regularisation will have moved the measured components away from the true values. However, we seek a solution where these components do match. The unknown components $d^{(1)}_{reg,u}$ have also been adjusted from the zero values they were set to before, which is caused by the regularisation. We wish to maintain this behaviour. Therefore, we seek a solution where the data for the next iteration is

$$d^{(2)} = \left\{ \begin{array}{c} d_m \\ d^{(1)}_{reg,u} \end{array} \right\}. \tag{10}$$

We have outlined the process for the first iteration; these steps can be repeatedly applied until convergence is achieved. As highlighted, the steps up to this point are as given in [34].

In X-ray CT reconstruction, the FST, introduced in Section 2.1, provides the means to transform between the image space (in which regularisation can be applied) and the data space (in which data can be combined and replaced). It is powerful in that it is fast, which is critically important when iterating, and it is invertible, relying only on the 2D FFT and its inverse, avoiding the solution drifting from the true data. This is the key contribution of the FNSR, and it should be noted that it is general for any regularisation approach allowing both binary and non-binary x-ray images to be recovered.

FNSR is applied as follows and summarised in Figure 1: (1) take the NUFFT of the projection data and populate the corresponding 2D k-space locations. The initial image is estimated by taking the inverse 2D FFT. (2) regularisation is applied, (3) the 2D FFT takes the image into k-space, (4) the measured components are corrected to their true components and (5) the latest image is generated by the inverse 2D FFT. This is then iterated from point (2).

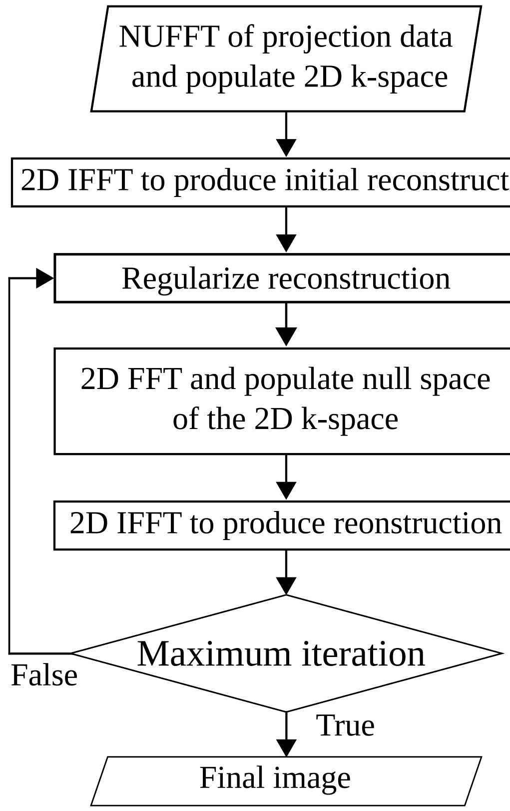

Figure 1. Summary workflow of the FNSR algorithm.

2.3. Data k-space stencil

Fundamental to the FNSR algorithm is the population of the data and null components of the image k-space. The FST is used to populate the 2D image k-space using the FFT of the projection data [3, 4]. Typically, the direct calculation of the FST is avoided due to significant interpolation errors associated with converting the projection data from a polar to cartesian grid. However, these interpolation errors are mitigated by sampling the 2D image k-space using a series of concentric squares [36], [37] and using the NUFFT to obtain exact horizontal or vertical grid point values (Figure. 2). This sampling scheme reduces the interpolation to a single dimension with results comparable to FBP [36].

The concentric sampling and data k-space estimation are computed as follows:

(i) Data are divided in two; sub-horizontal projection $\theta \in [-45^{\circ}, 45^{\circ})$ and sub-vertical projection $\theta \in [45^{\circ}, 135^{\circ})$ where θ is the angle of projection taken counter clockwise from the horizontal.

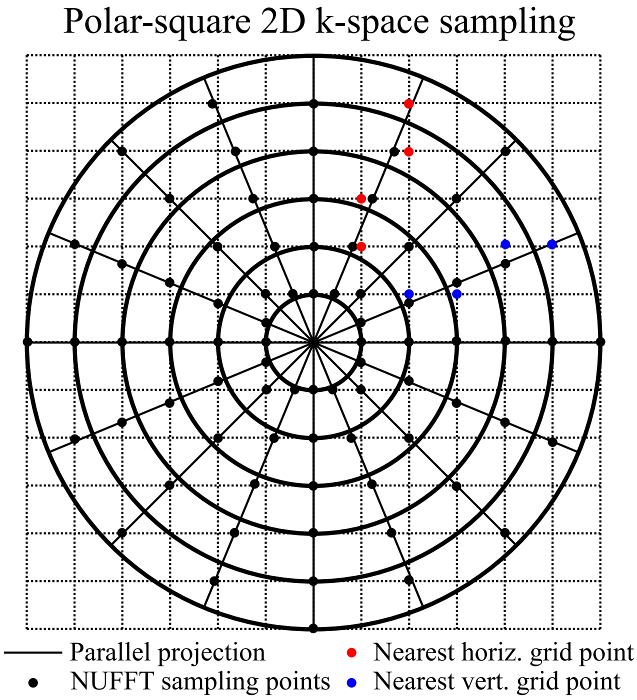

Figure 2. Illustration of the k-space sampling using a polar and concentric square grid. The the radial lines are the parallel projections the circles show the sampling positions obtained by Fourier Transforming the projection data. The black dots are the sampling positions we wish to obtain using the NUFFT. The red and blue dots show the location of the nearest vertical or horizontal grid point to each NUFFT sample.

(ii) Compute the 1D NUFFT of the projection data using the method of Greengard and Lee [28] with a desired output sampling interval of $\Delta k_p = \frac{\Delta k_x}{\cos\theta}$ for sub-horizontal and $\Delta k_p = \frac{\Delta k_y}{\sin\theta}$ for sub-vertical projections where Δk_x and Δk_y are the horizontal and vertical sampling rate of the 2D k-space. These points correspond to the black dots along the radial lines in Figure2.

(iii) 1D linear interpolation of the computed k-space values along vertical lines for sub-horizontal projections ($\theta \in [-45°, 45°)$, red dots Figure 2) and along horizontal lines for the sub-vertical projections ($\theta \in [45°, 135°)$ blue dots Figure 2).

This k-space, where the data have been interpolated onto the nearest horizontal or vertical grid space (red and blue dots in Figure 2) is called the data k-space X_{data} and corresponds to d_m in (4). Following, that we define a binary k-space stencil $X_{stencil}$ which is used to define the null space of the current system where the j^{th} pixel is defined as:

$$X^j_{stencil} = \begin{cases} 1 & \mid X^j_{data} \mid = 0 \\ 0 & \text{otherwise} \end{cases} . \tag{11}$$

By performing a point-by-point multiplication of this stencil with the FFT of the regularized image X_{reg} we eliminate the data component d_m in the k-space leaving only components belonging to $d^{(k)}_{reg,m}$. A combined dataset (10) is estimated by adding the X_{data} to the newly formed regularised k-space.

3. Binary FNSR algorithm

In this section we describe regularization scheme when combined with FNSR generates binary reconstructions of X-ray CT data. We begin by describing the binary steering approach of Censor [22] followed by a brief description on the implementation of the algorithm. A summary of the algorithm is given by the pseudo-code in Algorithm 1.

3.1. Binary regularization

The binary regularization scheme used is a binary steering approach which incrementally steers each iteration towards a discrete image [22]. Any iterative reconstruction may be written as

$$x^{k+1} = x^k + c(x^k, d), \tag{12}$$

where x is the image vector, d is the data, c is the correction function which updates the image, and k is the iteration number. The binary steering step is applied to the k^{th} iterate of the image vector x_k and is used as the input to the correction function. The method relies on three sequences of real numbers, $\alpha = \{\alpha_k\}_{k\geq 0}$ which defines the lower segmentation value, $\beta = \{\beta_k\}_{k\geq 0}$ which defines the upper segmentation value, and $\tau = \{\tau_k\}_{k\geq 0}$ which defines the final image threshold. The three sequences of numbers satisfy the conditions $0 \leq \alpha_k < \tau_k < \beta_k \leq 1$, $\alpha_k < \alpha_{k+1}$ and $\beta_{k+1} < \beta_k$. Typically τ is usually fixed at 0.5 with α and β defined as $\alpha_k = \frac{k}{K}\tau$ and $\beta_k = 1 - \alpha_k$, where K is the total number of iterations. At the k^{th} iterate α and β are applied to the j^{th} pixel of the image x^k to define the binarised image $\tilde{x}^k$ as [22],

$$\tilde{x}^k_j = \mathcal{B}\left(x^k\right) = \begin{cases} 0 & \text{if } x^k_j \leq \alpha_k \\ 1 & \text{if } x^k_j \geq \beta_k \\ x^k_j & \text{otherwise} \end{cases}, \tag{13}$$

However, practically the image pixels values are equal to the radiographic attenuation coefficient of the imaged object their values may not binary (i.e, 0 or 1). In order to apply the binarisation step the image is normalized between 0 and 1 using linear normalization which for the j^{th} pixel is given by:

$$\hat{x}^k_j = x^k_j \frac{1}{\max\left(x^k\right)}, \tag{14}$$

where $\max\left(x^k\right)$ is the maximum image pixel amplitude of the k^{th} iteration. The normalized image is subsequently binarised using (13).

Next an intermediary image $y^{k+1} = x^k + c(\tilde{x}^k, d)$ is computed by applying the correction function c to the binarised image $\tilde{x}^k$. The intermediate image y^{k+1} may result in the amplitudes of previously binarised pixels x^k significantly changing resulting in pixel 'conflict' [22] which can lead to the unstable binarisation and reconstructions. A pixel is defined as being in conflict when the intermediate image y^{k+1} crosses the image threshold τ i.e., if $x^k \leq \alpha_k$ and $y^{k+1} \geq \tau$ or if $x^k \geq \beta_k$ and $y^{k+1} \leq \tau$. To resolve pixel conflict and allow for the smooth and stable binarisation of the image a conflict resolution function $\mathcal{C}$ is defined as:

$$x_j^{k+1} = \mathcal{C}\left(y^{k+1}, x^k\right) = \begin{cases} \tau - \epsilon & \text{if } x_j^k \leq \alpha_k \text{ and } y_j^{k+1} \geq \tau \\ \tau + \epsilon & \text{if } x_j^k \geq \beta_k \text{ and } y_j^{k+1} \leq \tau \\ y_j^{k+1} & \text{otherwise} \end{cases} \tag{15}$$

where ϵ is a small constant greater than 0 which ensures that each pixel falls the correct side of the threshold value τ.

3.1.1. Implementation First, a binary image is generated from the previous iteration using (13) and (14). Next, the 2D FFT of the binarised image is computed and used to regularize the reconstruction. Differences in the amplitudes of the binarised image $\tilde{x}^k$ and the data derived image x must be accounted for prior to regularization. This is achieved by computing the modified inverse of 14 where $\max(x^k)$ is replaced by the median of all pixels which were identified with values greater than the threshold β (median $\left(\{x_j^k \mid \tilde{x}_j^N = 1\}\right)$). The median was used instead of the maximum to reduce the effects of any spuriously large pixel amplitude values on the regularization. Following the rescaling the 2D FFT of the image is computed to give the the binarised k-space X_b.

Next, the k-space of the binarised image, X_b is projected on to the null space of the current reconstruction by performing element by element multiplication with the stencil $X_{stencil}$. Following the calculation of the regularized reconstruction y^k any conflicts associated with changes in pixel assignment are resolved using (15) where the terms α, β, τ and ϵ are scaled according to the inverse of (13).

By binarising the image, information about sharp contrast objects will be introduced into the image null space via the FFT. The effect of taking the 2D FFT of an image with sharp edges in the image domain is a sinc function centred at 0 frequency in the k-space. The effect of any spurious ringing on the reconstruction will be limited for two main reasons (1) The majority of the information relating to an image is contained at low spatial frequencies in the k-space, which is adequately sampled even with a few projections due to the polar nature of the sampling geometry – considering two points at the same radius separated by a particular angle, the cartesian separation distance will be small when the radius is small. Therefore, at low spatial frequency the 2D k-space is dominated by the data. (2) The Fourier transform is linear and as such its inverse will reproduce the original sharp edges of the image.

Of greater concern to the quality of the reconstruction is the potential presence of high frequency noise in the image i.e., single pixel holes in an object. To alleviate this the final stage of the algorithm is a smoothing operation which is used to enforce coherency between adjacent pixels and further mitigate the mislabeling of pixels from the binarisation phase. From a practical perspective the smoothing is justified by the fact that the imaged objects are homogeneous with any flaws or voids expected to be made up of more than tens of pixels.

Algorithm 1 FNSR reconstruction

Input: Acquisition $p(\theta, t)$, iterations K, filter type and size $F(x)$, binarisation parameters (τ, S)
Output: binarised image reconstruction x

1D NUFFT $p(\theta, t)$ and populate 2D k-space X_{data}
$x^0 \leftarrow \mathcal{F}^{-1}(X_{data})$
$X_{stencil} \leftarrow X_{data}$ ▷ Binary k-space stencil (11)
for k = 1 to K **do**
  $\alpha_k \leftarrow \frac{k}{K}\tau$, $\beta_k \leftarrow (1-\alpha_k)$
  binarise normalized image by (13) and compute FFT:
  $\tilde{x} \leftarrow \mathcal{B}\left(x^{k-1}\right)$, $X_b \leftarrow \mathcal{F}(\tilde{x})$
  Compute regularized k-space and IFFT:
  $X_{reg} \leftarrow X_{data} + (X_b \circ X_{stencil})$, $y^k \leftarrow \mathcal{F}^{-1}(X_{reg})$
  Conflict resolution (15)
  $x_c \leftarrow \mathcal{C}\left(y^k\right)$
  Update reconstruction by filtering
  $x^k \leftarrow F(x_c)$
$x \leftarrow x^k$

4. Numerical Examples

A series of numerical experiments were conducted to assess the capability of the proposed algorithm to accurately reconstruct five different phantoms. The tests assess the FNSR sensitivity to iteration number, filter size and the convergence behaviour of the method. Finally a comparison is made between FNSR and the algebraic reconstruction method's capability to reconstruct an object using a limited number of projections. The projections are computed for a parallel beam geometry using (1) with each projection spaced evenly between 0 and 180°.

The quality of the reconstruction is assessed by comparing the percentage of mislabeled pixels and the root mean squared (RMS) of the difference between the phantom and reconstruction. The mislabeling of the j^{th} pixel is defined as

$$x_m^j = \begin{cases} 1 & \mid x_t^j - x_r^j \mid > 0.5, \\ 0 & \text{otherwise} \end{cases}. \tag{16}$$

where x_t is the true image and x_r is the reconstruction. Throughout the experiments the final segmentation threshold $\tau = 0.5$ and the conflict resolution parameter $\epsilon = 10^{-4}$

(a) (b) (c) (d) (e)

Figure 3. Phantoms used for the numerical simulations. (a) Phantom 1. (b) Phantom 2. (c) Phantom 3. (d) Phantom 4. (e) Phantom 5.

4.1. Phantoms

The numerical experiments are based on five binary representative phantoms (Figure 3), where Phantoms 2 and 4 were taken from [25] whilst phantom 5 was generated using the XDesign software package [38]. The trivial process of using the same forward model in the inversion as used to generate the simulated data is known as an 'inverse crime' and should be avoided [32]. Within CT imaging, inverse crimes are typically avoided by using continuous or high resolution phantom data to generate the projections that can then be used for the desired reconstruction at a lower resolution. Here, the projection data were computed using phantoms with pixel resolutions of 1024×1024 for phantom 1, 4096×4096 for phantoms 2-4 and 2048×2048 for phantom 5 and a desired image reconstruction of 512×512. The projection data were interpolated so that each detector pixel corresponded to the length of the pixels in the output image. For comparisons with the reconstructed data the phantoms were downsampled to 512 × 512 pixels.

4.2. Effect of number of iterations

A key factor of the performance of the FNSR algorithm is the total number of iterations, since this controls the rate of binarisation and the null space thresholding. The effect of the number of iterations on the quality of reconstruction as a function of the number of projections is tested here. For this test no filtering was applied to the results at the end of each iteration.

Figure 4 shows the RMS and the percentage number of mislabeled pixels for Phantoms 1 and 3 for different number of iterations. In both examples it can be seen that fewer iterations, generally 50 or below, produce results with the lowest number of misidentified pixels and RMS error. However, the two phantoms have distinctly different error trajectories. In the case of Phantom 1 the optimum number of iterations are dependent on the number of available projections; for example at 180 and 36 projections 10 iterations marginally produce the most accurate reconstructions but for fewer than 36 projections the image quality rapidly decreases. On average, 50 iterations produces the reconstructions with the lowest number of mislabeled pixels and RMS error. However, for Phantom 3 the error trajectories are distinct with 10 iterations producing the best reconstruction in all cases in Figure 4.

The variability in the reconstruction quality as a function of iteration is seen in Figure 5. Where an increasing number of iterations are used, leading to a finer set of binarisation thresholds, the images exhibit 'noise' in the form of mislabeled pixels. These mislabeled pixels and associated image noise are particularly prevalent at low numbers of projections. The mislabeling of pixels tends to occur during the final set of iterations where pixels with amplitudes close to the final segmentation threshold τ are set incorrectly.

Two main factors contribute to the mislabeling of pixels at higher iterations. The first is the influence of the null space regularization on the FNSR reconstruction which is dependent on the both the shape of the imaged object and relative size of the data and null spaces. This explains the variability in error trajectories between both phantoms in Fig 4. In general, as the size of the null space increases it exerts a greater influence on the solution and, as such, a pixel which has been labeled by the binarisation process

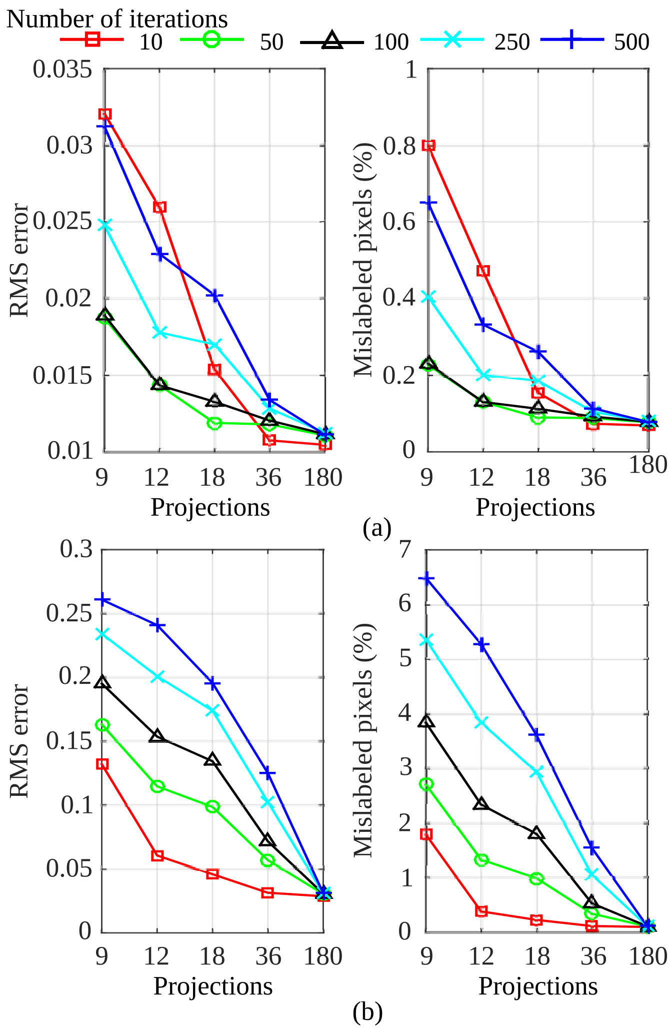

Figure 4. Comparison of the effect of different number of iterations on the final RMS error and misidentified pixels of (a) Phantom 1 and (b) Phantom 3 using a limited number of projections.

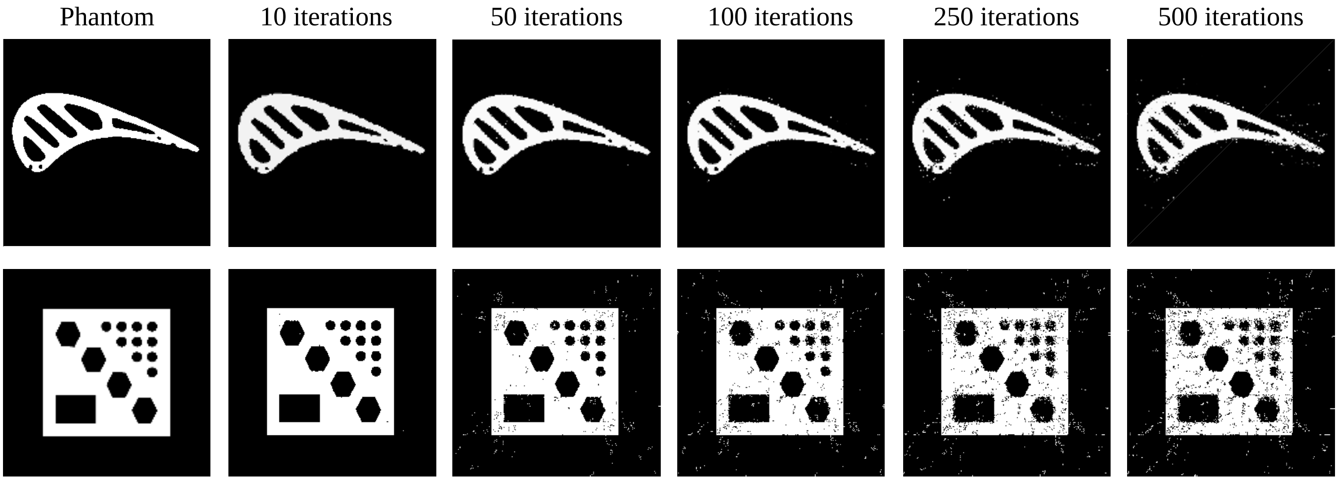

Figure 5. Comparison of the effect of different number of iterations on RMS error and misidentified pixels of (a) Phantom 1 and (b) Phantom 3 using 18 projections.

tends to remain with this designation in subsequent iterations.

The second influence on the solution is the conflict resolution phase (15) which stabilizes reconstruction by inhibiting significant change in pixel labels particularly close to τ. In addition, other non binarised pixels will now exhibit large changes in amplitude to compensate for the mislabeling of other pixels thus enhancing the effect. To combat these effects the number of iterations should be as low as possible to produce an accurate reconstruction whilst also reflecting the expected pixel amplitude variability for each iteration. In addition a filtering operation is applied at each iteration to regularize the solution, which is discussed next.

4.3. Filter regularization

In the previous section it was seen that some regularization was necessary in order to prevent the mislabeling of pixels and the generation of noise in the image (Figure 5). We select a 2D median filter as a method of regularizing the image and test how the window size affects the quality of the reconstruction. The median filter was selected because of its ability to minimise 'salt and pepper' type noise, which is how these errors typically manifest themselves, while avoiding blurring edges. All experiments were conducted using 50 iterations.

Figure 6 shows the RMS error and the percentage number of mislabeled pixels of Phantom 5 with varying filter size. From the two graphs it can clearly be seen that the addition of a small filter improves the reconstruction for a decreasing number of projections whilst larger filters degrade the image. The effect of the different filters for 18 projections is seen in Figure 7. We can see that the 3×3 filter reduces the number mislabeled pixels whilst preserving the edges of the internal features of the object. The larger filter windows 7×7 and 9×9 remove all mislabeled pixels at the expense of rounding the edges and removal of the internal features of the object. From both the

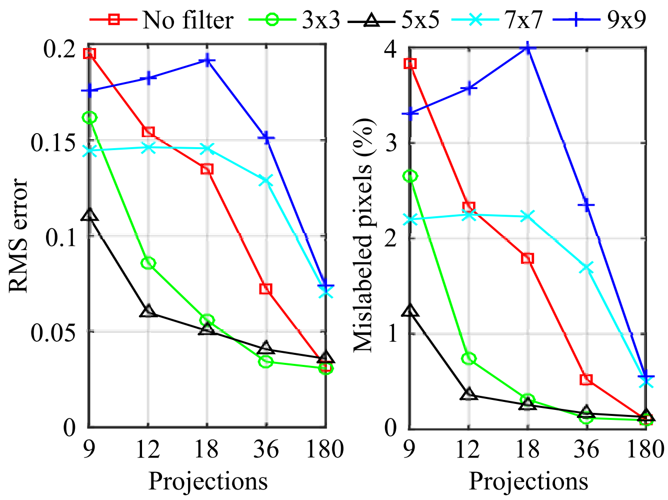

Figure 6. Comparison of the effects of different median filter lengths on the RMS error and misidentified pixels of Phantom 5 using a limited number of projections and FNSR.

graph (Figure 6) and image comparison (Figure 7) the 5×5 median filter produces the best result by removing the noise associated with the mislabeled pixels whilst preserving the edges of the internal features of the object.

The optimal filter size is dependent on the effect of the number of iterations on the reconstruction, which itself is a function of the number of projections (Figure 4). In general the best results are obtained when the lowest number of iterations and smallest filter size are used.

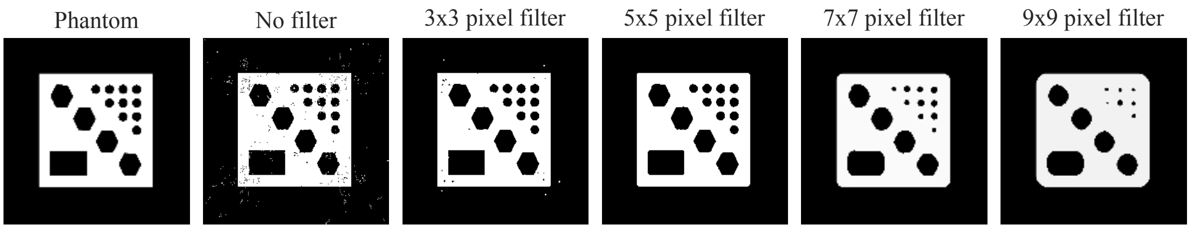

Figure 7. Comparison of the effects of different median filter lengths on the reconstruction of Phantom 5 both with 18 projections.

4.4. Evaluation of reconstruction algorithms

A number of iterative algorithms have been developed over the years to solve the inverse problem $d = Ax$, where d are the data, A is the model or projection matrix which links the data to the image x e.g., [3, 4, 31, 32]. Algebraic reconstruction techniques (ART) are a family of reconstruction methods which have been routinely applied to X-ray CT and efficiently solve the inverse problem in a least squares sense using a series of forward and back projections [3, 4, 31, 32]. Here, we test the quality and speed of the proposed FNSR algorithm against a number of ART algorithms using a limited number of projections.

All tests were performed on a desktop workstation (32 Intel Xeon E502690 processor, 2.90GHz CPU, 8 cores and 256 GB RAM) using MATLAB 15.

The details of the ART algorithms tested are given below:

4.4.1. Algebraic Reconstruction Technique (ART) ART is an iterative reconstruction algorithm of form (12) where the data correction function is calculated by sweeping through each row of the matrix A and projecting the solution onto orthogonal hyperplanes e.g., [3, 4, 31]. A single iteration is completed when all rows of A have been swept through [3, 4, 31].

For the numerical experiments we set the total number of iterations to 100 which is sufficiently large to ensure that convergence has been achieved. Additionally, pixel amplitude positivity following each row iteration of the reconstruction is enforced. The inclusion of the positivity constraint yields better results than those without, particularly for limited data case e.g., [39].

4.4.2. ART-TV The TV reconstruction of Sidky *et al.,* was implemented [8]. This method is a two step approach which combines the ART algorithms with a TV minimization step. The first stage is a single iteration of ART where the positivity constraint is subsequently applied as opposed to after each individual row projection described above. The second stage of the algorithm is TV minimization using a fixed step, γ, and fixed iteration gradient descent approach. It should be noted that the gradient of the TV function with respect to each pixel is undefined, so a smooth approximation must be used where the amount of smoothing is controlled by the parameter δ [8].

For the experiments a number of parameters must be set which were determined using trial and error process. We set $\delta = 10^{-8}$ which is sufficiently small to stabilize the TV gradient without smearing any of the edges, the gradient descent has a fixed step length $\gamma = 0.05$ for 200 iterations whilst the total number of iterations was set to 100 which was sufficiently large for convergence.

4.4.3. Discrete Algebraic Reconstruction Technique (DART) DART is a practical discrete tomographic approach built upon the results of an ART approach [23]. The method relies on prior knowledge of the number of materials and their corresponding grey values to discretized an initial ART reconstruction. The method focuses on iteratively improving the boundaries between the segmented parts of the reconstruction which are known to be poorly resolved particularly when a limited number of projections are used [23]. At each iteration the edge pixels are identified along with a randomly selected subset of image pixels which are updated using the previous ART method [23].

One of the major source of errors in the original DART algorithm was that the grey values of the pixels must be known a priori. In practice, true pixel grey values are affected by a number of factors, including the source energy, the spectrum of the source and whether any filters are used, and are therefore difficult to predict. The so called Projection Distance Minimization (PDM) approach was developed to optimize the grey values and segmentation threshold at each iteration [24]. PDM is a two stage optimization approach which minimizes the least squares difference between the measured projections and those calculated using the currently segmented image [24]. The first stage optimizes the segmentation thresholds followed by the pixel grey values using the Nelder-Mead simplex algorithm [24].

For the numerical examples, the Simultaneous Iterative Reconstruction Technique (SIRT) was used for the initial reconstruction with 500 iterations and positivity constraints applied at each iteration. A total of 200 iterations were used for DART and 100 iterations of SIRT for the edge update phase. At each DART iteration a PDM update was computed for the image segmentation.

4.4.4. FNSR The number of iterations and size of the regularization filter were determined following the results of Section 4.2 and Section 4.3. A total of 50 iterations was deemed optimal with a regularization filter of 3×3 in the case of Phantom 1 and 5×5 for all other Phantoms (Figure 3).

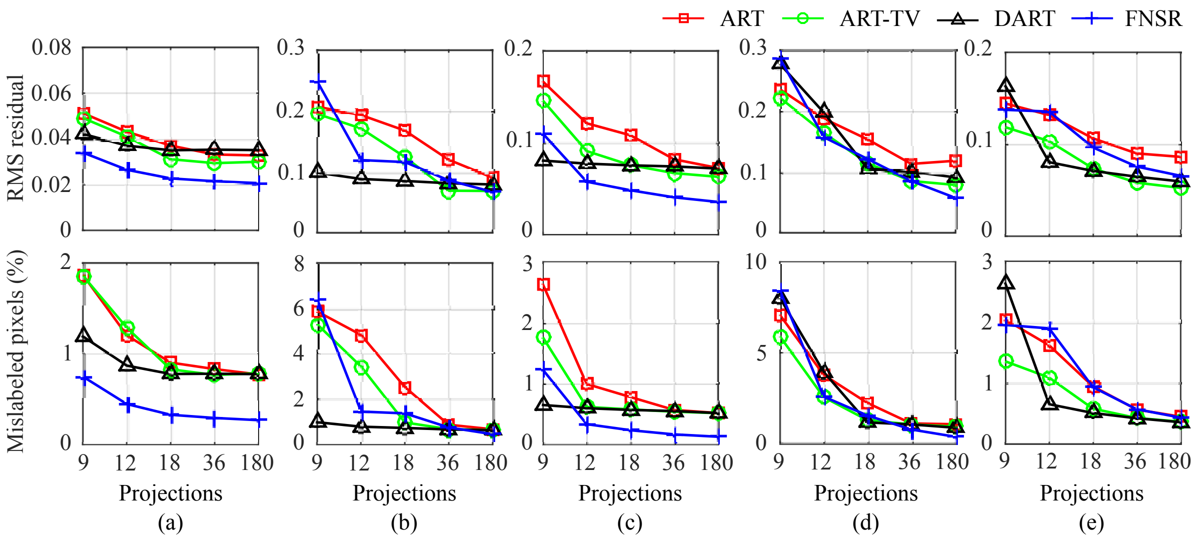

Figure 8. RMS error and misidentified pixels of the reconstruction methods using a limited number of projections. (a) Phantom 1. (b) Phantom 2. (c) Phantom 3. (d) Phantom 4. (e) Phantom 5.

Figure 8 summarizes the RMS error and the percentage of mislabeled pixels as a function of projections for all phantoms. The results show that FNSR consistently produces more accurate reconstructions than the ART and ART-TV algorithms for all projections. In the case of Phantoms 1 and 3 and for 36 projections or more FNSR routinely outperforms DART. Table 1 provides an overview of the average ranking of the various algorithms relative to one another as a function of the number of projections. For 12 or more projections FNSR is ranked either 1^{st} or 2^{nd} within the algorithms tested and falling to 3^{rd} for 9 projections. These results suggest that FNSR required at least 12 projections or more in order to obtain an accurate reconstruction of the object.

Table 1. Rankings of the RMS and mislabeled pixels of the reconstructions using a limited number of projections averaged over all phantoms. The final column overall ranking of the reconstruction method. A score of 1 has the lowest RMS error/number of mislabeled pixels and 4 the highest.

RMS ranking

Method	9 Projections	12 Projections	18 Projections	36 Projections	180 Projections	Averag
ART	4	4	4	4	4	4
ART-TV	1	3	3	1	3	2
DART	2	2	1	3	2	3
FNSR	3	1	2	2	1	1

Mislabeled pixel ranking

Method	9 Projections	12 Projections	18 Projections	36 Projections	180 Projections	Averag
ART	4	4	4	4	4	4
ART-TV	1	3	3	1	3	3
DART	2	1	1	3	2	1
FNSR	3	1	2	1	1	2

Examples of reconstructions as a function of the number of projections are seen in Figure 9. The examples clearly show that the FNSR and DART reconstructions produce similar results for Phantoms 1 and 2 with significantly less smearing than seen in ART and ART-TV methods. In Phantom 4 some of the image details, particularly the ellipsoid shape hole at the center of the image has been overly smoothed by the regularizing filter. This smoothing may also bee seen in the rounding of some of the edges of Phantom 5. These images reiterate the importance of the trade off between iteration number and filter size.

Figure 10 compares the run time of each algorithm as a function of the number of projections for Phantom 3. All four algorithms scale approximately linearly as the number of projections increase. In the case of the ART methods the run times increase as a function of algorithm complexity. The run time of DART is significantly affected by the PDM based segmentation with each iteration spending the majority of time in Nelder-Mead simplex optimization [24]. FNSR is at least an order of magnitude faster than the other methods with the 180 projection reconstruction having a similar run time to an ART reconstruction with 9 projections.

The main deviation between FNSR and DART occurs at 12 projections or fewer (Figure 8 and Table 1). However, the reduction in run time of more than an order of magnitude allows the FNSR method to allow for additional projection data to be acquired without compromising processing speed, whilst improving the reconstruction quality. Relative to DART and ART-TV based methods FNSR produces equivalent or better quality reconstructions with a significantly reduced processing time.

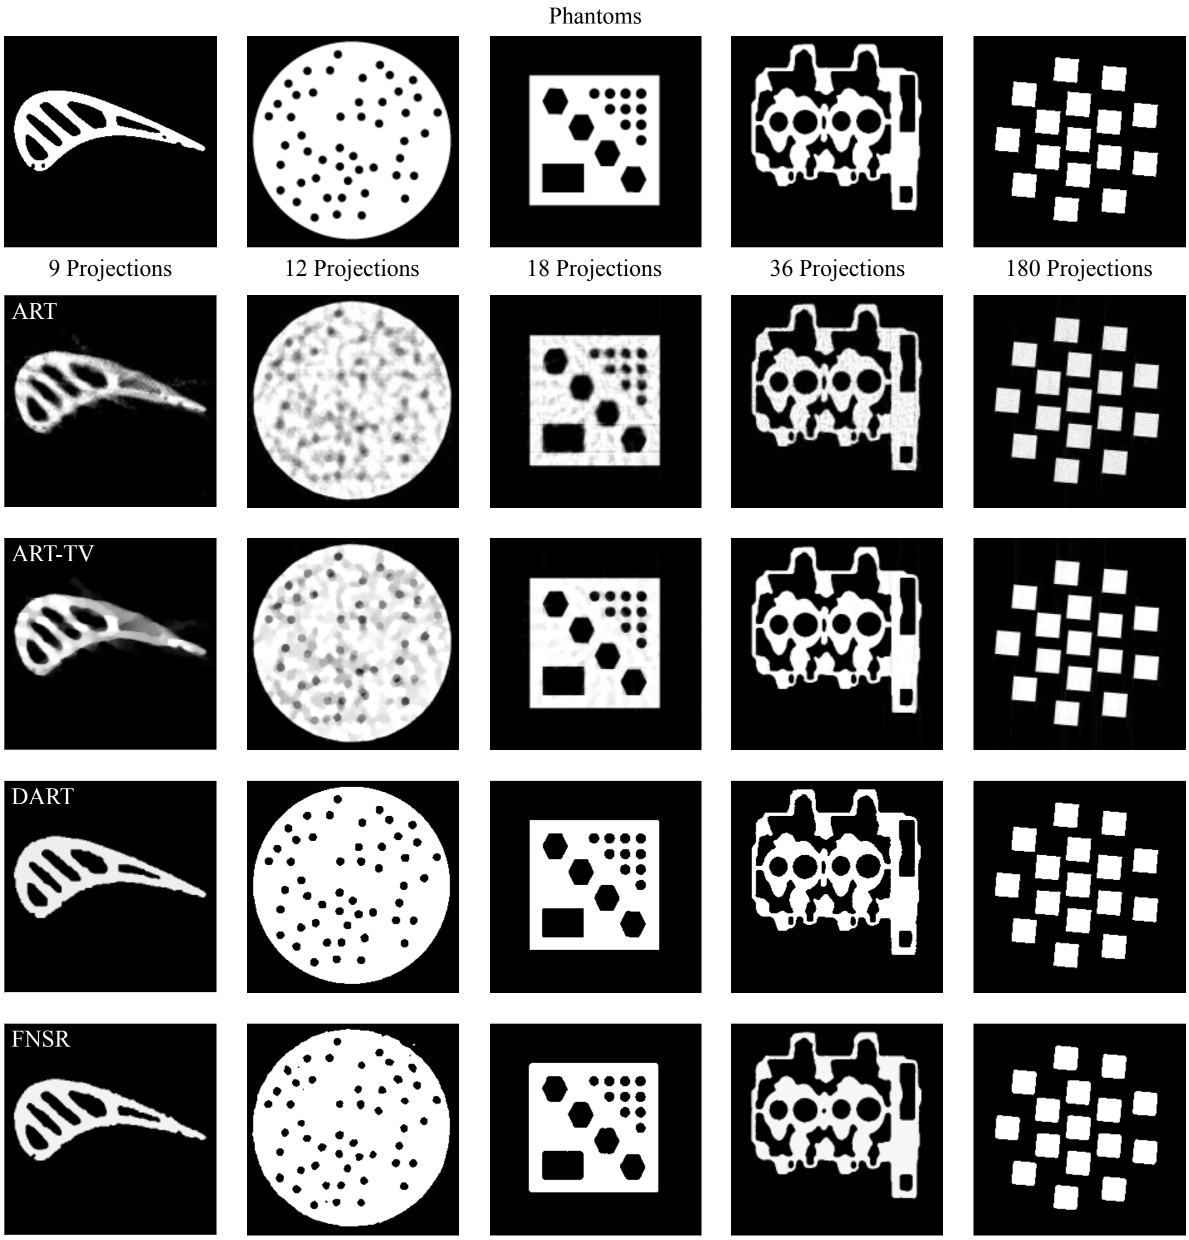

Figure 9. Comparison of reconstructions using a limited number of projections (columns) for ART (row 2), ART-TV (row 3), DART (row 4) and FNSR (row 5)

4.5. Convergence properties

The convergence behaviour of iterative methods is important to understand in order to assess the quality and computational stability of the method. FNSR is a heuristic reconstruction algorithm so a formal proof about the convergence conditions of the algorithm cannot be provided. Thus far the numerical experiments have shown that FNSR can accurately reconstruct a variety of phantoms with varying number of projections (Section 4.4) but the result is dependent on the number of projections (Section 4.2) and filter size (Section 4.3).

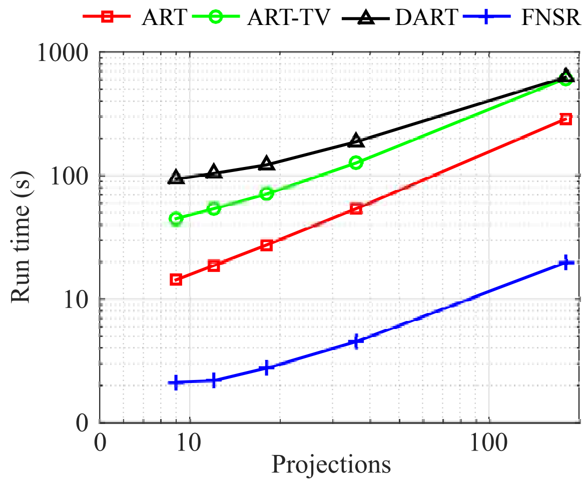

Figure 10. Comparison of the run times of the reconstruction methods using a limited number of projections for Phantom 3.

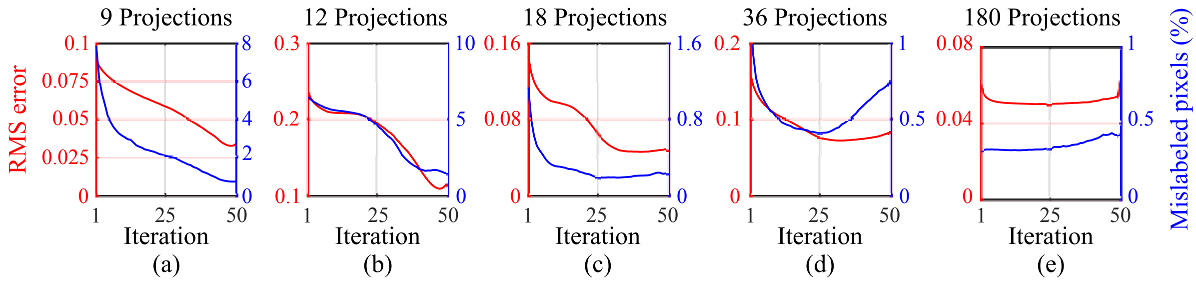

Figure 11. RMS error and mislabeled pixels of FNSR as a function of iteration using a limited number of projections. (a) Phantom 1. (b) Phantom 2. (c) Phantom 3. (d) Phantom 4. (e) Phantom 5.

The convergence properties of FNSR were tested using Phantoms 1–5 by comparing the RMS error and the number of mislabeled pixels as a function of iteration number for a varying number of projections. A total of 50 iterations were used and a 3×3 size median filter used for Phantom 1 and 5×5 filter for all other phantoms. As shown in Figure 11 FNSR generally displays smooth convergence for RMS error and mislabeled pixels although absolute convergence cannot be guaranteed. In all cases the RMS error and number of mislabeled pixels slightly increased at the final iterations as the binarisation sequence is completed. This slight increase is attributed to unlabeled pixels on the edge of the object taking definitive values.

The iteration history of mislabeled pixels for Fig 11 (d) displays a significant deviation from other convergence histories. This deviation can be attributed to the use of the incorrect number of iterations and filter size for that particular phantom. In this particular example the optimum number of iterations are 28 considering mislabeled pixels and 33 for RMS error. The results of the incorrect number of iterations and filter size can be seen in Figure 9 where the central ellipsoid hole is partially filled and image edges have been overly smoothed. This example reiterates the importance of the number of iterations in controlling the convergence of the reconstruction.

In the case of 180 projections we can see that the convergence history is flat with some deviation at the first and last iteration associated with the binarisation process (Figure 11). This type of behavior is not unexpected since there is sufficient data available to populate the 2D image k-space with the final result being effectively a binarised filtered back projection of the data.

5. Experimental Validation

Experimental X-ray CT data were acquired using a collimated 140 kV, 225 μA source with a 4 mm copper filter and recorded on a linear array of 2048, 0.415mm long pixels with exposure times of 1s [39]. A total of 5000 projections with an even angular sampling of 0.72° through 360° were acquired. Using all of 5000 projections, a tomographic slice of the turbine blade was obtained using the industry standard FDK method, which was post processed to remove any beam hardening artefacts and segmented [39]. This image will be used as a reference to compare the FNSR reconstructions against.

Prior to applying FNSR the fan beam data were preprocessed to remove any beam hardening using an empirically defined amplitude correction curve [39]. The current implementation of the FNSR is applicable to parallel ray datasets and as such the projection data were resampled to generate a limited projection parallel ray datasets. The projections correspond to an evenly sampled dataset from 0° to 180° with the number of projections the same as that used in the previous numerical examples (9, 12, 18, 36 and 180). The fan beam data were subsampled to the desired number of projections and rebinned to generate a parallel ray dataset [3]. In the process of rebinning only data which correspond to the desired projection angles are retained which leads to an uneven sampling of the detector and subsequently the introduction of interpolation errors. We note that the NUFFT has been successfully and accurately applied to the transformation of fan to parallel beam data [29, 30] and could readily applied instead of rebining and interpolation.

For the experimental reconstructions a total of 10 iterations was selected based on the results of the iterations test (Section 4.2) and a 3×3 regularization filter was used. Figure 12 shows the reconstruction of the experimental turbine blade dataset. It can clearly be seen that FNSR produces accurate images up to 18 projections, which is similar to results obtained by Jones and Huthwaite [39], who investigated the performance of different ART and TV methods for different industrial CT datasets. The RMS error and number of mislabeled pixels in the experimental data (Figure 13) are slightly higher than those observed for the numerical experiments on Phantom 1 using FNSR (Figure 8). For 18 or more projections the errors in the experimental reconstructions are comparable with those observed by the ART based reconstruction of numerical data (Figure 8). When fewer that 18 projections are used, the FNSR experimental based reconstruction and errors significantly deviate from those observed in the numerical examples.

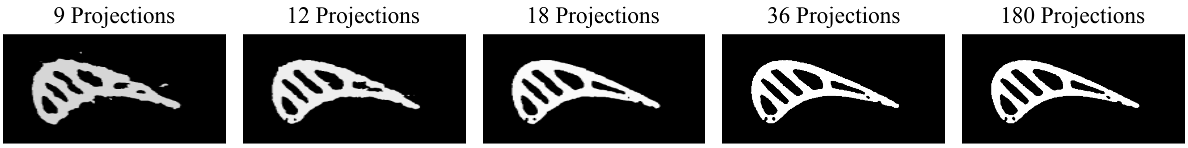

Figure 12. FNSR reconstruction of experimental turbine blade data as a function of the number of parallel projections.

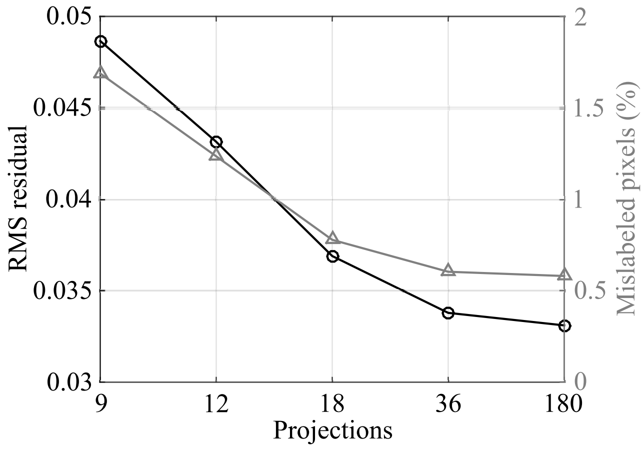

Figure 13. RMS and misidentified pixels between the reference image and FNSR as a function of the number of parallel projections.

The quality of the FNSR results are dependent on the number of iterations and the regularization filter, both of which are affected by noise. Any noise in the projections will introduce additional high frequencies into the image k-space thus degrading the image. To alleviate this high frequency noise either additional preprocessing of the projections are necessary, or an increase in the size of the FNSR regularization filter, which will limit the method to accurately reconstruct the small features and object edges. Additionally, beam hardening in the projection data can also limit FNSR and other binary reconstruction methods in the production of accurate images.

The number of iterations and the size of the smoothing filter will vary depending on the specific type of image regularisation used (Figure 1), the imaged object and the SNR. In an environment where similar objects are to be imaged e.g., on a production line the parameters would only need to be adjusted once. Based on the presented examples where a binary regularisation approach was used (Sections 4.2 and 4.3) we suggest setting the default reconstruction parameters to fewer than 50 iterations with a median filter of 3×3 or 5×5 pixels.

6. Conclusions

We have presented a fast and accurate iterative reconstruction algorithm for X-ray CT using a limited number of projections called Fourier Null Space Regularization (FNSR). The method uses an innovative approach to regularize the reconstruction and compensate for any missing projections by explicitly updating the image null space with values derived from a filtered image from the previous iteration. The speed of the method is achieved by directly applying the Fourier Slice Theorem using the NUFFT to compute the frequency spectrum of the projection data at their exact positions in the corresponding image k-space. Furthermore, the FNSR algorithm permits the use of any image regularization allowing both binary and non-binary x-ray images to be recovered. In line with many industrial X-ray CT applications we apply FNSR to binary tomographic reconstructions where a binary steering regularization method is used to drive the solution towards a discrete image.

The comparison of the reconstruction algorithms tested demonstrated that the FNSR method outperformed ART for all projections and produces comparable or better results to DART for 12 or more projections. The numerical experiments also highlight the significant reduction in computational time of more than an order of magnitude achieved by FNSR, relative to the other methods. The significant reduction in the computational processing time achieved by FNSR would permit more data to be acquired thus improving reconstruction quality without compromising acquisition or processing speeds.

The numerical experiments highlight the key role the number of iterations have in the accuracy of the reconstruction; too few and the reconstruction contains smearing artefacts from the lack of projections, whilst too many leads to fine ‘noise’ associated with the mislabeling of pixels close to the final segmentation threshold. The findings of the iteration test suggest that the number of iterations should be as low as possible, typically less than 50 and an additional regularizing filter should be included at the end of each iteration. A simple median filter with a window size of 3×3 or 5×5 was found to produce the best result in removing the noise associated with pixel mislabeling whilst preserving the edge and internal features of the object. Improvements to the overall FNSR algorithm could be made to the binarisation process by using spatially varying thresholding e.g., [40] and a more advanced filtering operation e.g., [41].

FNSR was validated using an industrial X-ray CT dataset of a turbine blade. The data were acquired in a fan beam configuration and resampled to generate a series of evenly spaces parallel ray datasetsfrom $0°$ to $180°$. FNSR produced accurate reconstructions for 18 or more projections with error estimates which are comparable to the numerical examples, confirming the applicability of the algorithm in practice. Future improvement for the practical application of the algorithm would be the extension of the method to non parallel ray acquisition geometries e.g., fan beam.

Acknowledgments

GAJ is funded by EPSRC RCNDE 3 (EP/L022125/1) and PH is funded by EPSRC Early Career Fellowship (EP/M020207/1). The authors would like to thank Rolls-Royce for providing the turbine blade and the Manufacturing Technology Centre for acquiring the turbine blade X-ray data. We would also like to thank 2 anonymous reviewers and editorial board for their constructive comments and suggestions for the manuscript.

[1] Kruth JP, Bartscher M, Carmignato S, Schmitt R, De Chiffre L, Weckenmann A. Computed tomography for dimensional metrology. CIRP Annals-Manufacturing Technology. 2011;60(2):821–842.

[2] Baruchel J, Buffiere JY, Maire E. X-ray tomography in material science. 2000;.

[3] Kak AC, Slaney M. Principles of Computerized Tomographic Imaging. IEEE Press; 1988.

[4] Natterer F. The mathematics of computerized tomography. SIAM; 2001.

[5] Ramlau R, Ring W. A Mumford–Shah level-set approach for the inversion and segmentation of X-ray tomography data. Journal of Computational Physics. 2007;221(2):539–557.

[6] D C, E S. An Introduction to Bayesian Scientific Computing: Ten Lectures on Subjective Computing. vol. 2. Springer Science & Business Media; 2007.

[7] Rudin LI, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. Physica D: Nonlinear Phenomena. 1992;60(1):258–268.

[8] Sidky EY, Kao CM, Pan X. Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT. Journal of X-ray Science and Technology. 2006;14(2):119–139.

[9] Romanov M, Dahl AB, C YDP, Hansen. Simultaneous tomographic reconstruction and segmentation with class priors. Inverse Problems in Science and Engineering. 2016;24(8):1432–1453.

[10] Kjer HM, Dong Y, Hansen PC. User-Friendly Simultaneous Tomographic Reconstruction and Segmentation with Class Priors. In: International Conference on Scale Space and Variational Methods in Computer Vision. Springer; 2017. p. 260–270.

[11] Storath M, Weinmann A, Frikel J, Unser M. Joint image reconstruction and segmentation using the Potts model. Inverse Problems. 2015;31(2):025003.

[12] Kadu A, van Leeuwen T, Batenburg KJ. A parametric level-set method for partially discrete tomography. In: International Conference on Discrete Geometry for Computer Imagery. Springer; 2017. p. 122–134.

[13] Herman GT, Kuba A. Discrete tomography: Foundations, algorithms, and applications. Springer; 1999.

[14] Herman GT, Kuba A. Advances in discrete tomography and its applications. Springer; 2007.

[15] Alpers A, Gritzmann P, Thorens L. Stability and instability in discrete tomography. In: Digital and image geometry. Springer; 2001. p. 175–186.

[16] Gardner RJ, Gritzmann P, Prangenberg D. On the computational complexity of reconstructing lattice sets from their X-rays. Discrete Mathematics. 1999;202(1-3):45–71.

[17] Gritzmann P, Vries SD, Wiegelmann M. Approximating binary images from discrete X-rays. SIAM Journal on Optimization. 2000;11(2):522–546.

[18] Batenburg KJ. An evolutionary algorithm for discrete tomography. Discrete applied mathematics. 2005;151(1-3):36–54.

[19] Baumann J, Kiss Z, Krimmel S, Kuba A, Nagy A, Rodek L, et al. Discrete tomography methods for nondestructive testing. In: Advances in discrete tomography and its applications. Springer; 2007. p. 303–331.

[20] Valenti C. A genetic algorithm for discrete tomography reconstruction. Genetic Programming and Evolvable Machines. 2008;9(1):85–96.

[21] Batenburg KJ. A network flow algorithm for reconstructing binary images from continuous X-rays. Journal of Mathematical Imaging and Vision. 2008;30(3):231–248.

[22] Censor Y. Binary steering in discrete tomography reconstruction with sequential and simultaneous iterative algorithms. Linear Algebra and its Applications. 2001;339(1-3):111–124.

[23] Batenburg KJ, Sijbers J. DART: a practical reconstruction algorithm for discrete tomography. IEEE Transactions on Image Processing. 2011;20(9):2542–2553.

[24] van Aarle W, K J Batenburg J, Sijbers. Automatic parameter estimation for the discrete algebraic reconstruction technique (DART). IEEE Transactions on Image Processing. 2012;21(11):4608–4621.

[25] Dabravolski A, Batenburg KJ, Sijbers J. A multiresolution approach to discrete tomography using

DART. PloS one. 2014;9(9):e106090.
[26] Xu F, Mueller K. Accelerating popular tomographic reconstruction algorithms on commodity PC graphics hardware. IEEE Transactions on nuclear science. 2005;52(3):654–663.
[27] Dutt A, Rokhlin V. Fast Fourier transforms for nonequispaced data. SIAM Journal on Scientific computing. 1993;14(6):1368–1393.
[28] Greengard L, Leslie, Lee JY. Accelerating the nonuniform fast Fourier transform. SIAM review. 2004;46(3):443–454.
[29] Matej S, Fessler JA, Kazantsev IG. Iterative tomographic image reconstruction using Fourier-based forward and back-projectors. IEEE Transactions on medical imaging. 2004;23(4):401–412.
[30] O'Connor YZ, Fessler JA. Fourier-based forward and back-projectors in iterative fan-beam tomographic image reconstruction. IEEE transactions on medical imaging. 2006;25(5):582–589.
[31] Hansen PC. Discrete Inverse Problems: Insight and Algorithms. SIAM; 2010.
[32] Müller JL, Siltanen S. Linear and Nonlinear Inverse Problems with Practical Applications. SIAM; 2012.
[33] Deal MM, Nolet G. Nullspace shuttles. Geophysical Journal International. 1996;124(2):372–380.
[34] Huthwaite P, Zwiebel AA, Simonetti F. A new regularization technique for limited-view sound-speed imaging. IEEE transactions on ultrasonics, ferroelectrics, and frequency control. 2013;60(3):603–613.
[35] Shi F, Huthwaite P. Ultrasonic wave-speed diffraction tomography with undersampled data using virtual transducers. IEEE transactions on ultrasonics, ferroelectrics, and frequency control. 2018;65(7).
[36] Natterer F. Fourier reconstruction in tomography. Numerische Mathematik. 1985;47(3):343–353.
[37] De Francesco S, Silva A. Fourier methods in tomography, a review. Electrónica e Telecomunicações. 2003;3(8):818–826.
[38] Ching DJ, Gürsoy D. XDesign: an open-source software package for designing x-ray imaging phantoms and experiments. Journal of synchrotron radiation. 2017;24(2):537–544.
[39] Jones GA, Huthwaite P. Limited view X-ray tomography for dimensional measurements. NDT & E International. 2018;93:98–109.
[40] Bradley D, Roth G. Adaptive thresholding using the integral image. Journal of graphics tools. 2007;12(2):13–21.
[41] Perona P, Malik J. Scale-space and edge detection using anisotropic diffusion. IEEE Transactions on pattern analysis and machine intelligence. 1990;12(7):629–639.