



Swansea University  
Prifysgol Abertawe



## Cronfa - Swansea University Open Access Repository

---

This is an author produced version of a paper published in:

*NDT & E International*

Cronfa URL for this paper:

<http://cronfa.swan.ac.uk/Record/cronfa48855>

---

### **Paper:**

Jones, G. & Huthwaite, P. (2018). Limited view X-ray tomography for dimensional measurements. *NDT & E International*, 93, 98-109.

<http://dx.doi.org/10.1016/j.ndteint.2017.09.002>

Released under the terms of a Creative Commons Attribution License (CC-BY).

---

This item is brought to you by Swansea University. Any person downloading material is agreeing to abide by the terms of the repository licence. Copies of full text items may be used or reproduced in any format or medium, without prior permission for personal research or study, educational or non-commercial purposes only. The copyright for any work remains with the original author unless otherwise specified. The full-text must not be sold in any format or medium without the formal permission of the copyright holder.


Permission for multiple reproductions should be obtained from the original author.

Authors are personally responsible for adhering to copyright and publisher restrictions when uploading content to the repository.

<http://www.swansea.ac.uk/library/researchsupport/ris-support/>

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

# NDT and E International

journal homepage: [www.elsevier.com/locate/ndteint](http://www.elsevier.com/locate/ndteint)

## Limited view X-ray tomography for dimensional measurements

G.A. Jones<sup>\*</sup>, P. Huthwaite

Department of Mechanical Engineering, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK

### ARTICLE INFO

#### Keywords:

X-ray CT  
Tomography  
Imaging  
Limited view reconstruction

### ABSTRACT

The growing use of complex and irregularly shaped components for safety-critical applications has increasingly led to the adoption of X-ray CT as an NDE inspection tool. Standard X-ray CT methods require thousands of projections, each distributed evenly through 360° to produce an accurate image. The time consuming acquisition of thousands of projections can lead to significant bottlenecks. Recent developments in medical imaging driven by both increasing computational power and the desire to reduce patient X-ray exposure have led to the development of a number of limited view CT methodologies. Thus far these limited view algorithms have been applied to basic synthetic data derived from simple medical phantoms. Here, we use experimental data to rigorously test the capability of limited view algorithms to accurately reconstruct and precisely measure the dimensional features of an additive manufactured sample and a turbine blade. Our findings highlight the importance of prior information in producing accurate reconstructions capable of significantly reducing X-ray projections by at least an order of magnitude. In the turbine blade example a dramatic reduction in projections from 5000 to 24 was observed while still demonstrating the same level of accuracy as standard CT methods. The findings of the study also suggest the importance of sample complexity and the presence of sparsity in the X-ray projections in order to maximise the capabilities of these limited algorithms. With the ever increasing computational power, limited view CT algorithms offer a method for reducing data acquisition time and alleviating manufacturing throughput bottlenecks without compromising image accuracy and quality.

### 1. Introduction

Modern engineering is increasingly utilising complex components. Turbine blades, for example, feature complex cooling channels and highly optimised curved surfaces, and the rise of additive manufacturing has given huge potential for extremely complex shapes. Such shapes present significant inspection challenges to traditional NDE techniques, as these features can obscure defects or manufacturing errors. X-ray computational tomography (CT) is one of the few technologies capable of non-destructively measuring both the external and internal features of a component [1]. Numerous CT approaches exist, but within industry they commonly consist of a static X-ray source and a movable detector which is perpendicular to the source. The sample to be CT scanned is placed on a movable disk which rotates through 360° allowing multiple X-rays projections to be captured. Standard X-ray CT methods require thousands of projections, each regularly and evenly distributed through 360° to produce an accurate image [2,3]. Once an accurate tomographic image is generated it can be used to assess the specimen for flaws, quality control

and undergo dimensional analysis by comparison with CAD e.g., [1].

One of the major downsides of X-ray CT is the time consuming data acquisition process which can lead to significant bottlenecks. To alleviate these bottlenecks in throughput, companies may be forced to purchase additional X-ray CT capability at great cost or reduce individual X-ray exposure times lowering the signal-to-noise ratio and image quality. Spurred by the ever increasing power of computers and the increasing flexibility of graphics cards for general purpose computing [4], a variety of limited view tomographic techniques capable of generating high quality images with less data have been developed for medical purposes e.g., [5–13]. Although much theoretical work has been conducted over the past decade on limited view tomography the algorithms developed have been tested on simple synthetic examples (e.g., Shepp - Logan phantom [2]) typically with parallel ray geometry, which poorly mimics true industrial applications where ray path geometries and noise are an issue.

Routine medical X-ray CT applications are limited to disease or trauma detection and the precise measurements are often not required as

<sup>\*</sup> Corresponding author.

E-mail addresses: [glenn.jones@imperial.ac.uk](mailto:glenn.jones@imperial.ac.uk) (G.A. Jones), [p.huthwaite@imperial.ac.uk](mailto:p.huthwaite@imperial.ac.uk) (P. Huthwaite).

<http://dx.doi.org/10.1016/j.ndteint.2017.09.002>

Received 30 March 2017; Received in revised form 15 August 2017; Accepted 8 September 2017

Available online 12 September 2017

0963-8695/© 2017 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

other, higher resolution and targeted, diagnostics methodologies (e.g., MRI) may be applied. In contrast, industrial CT requires precise high resolution and high contrast images to provide meaningful dimensional measurements for quality control purposes. These differences in the objectives of medical and industrial X-ray CT may prevent the direct transfer of the newly developed imaging algorithms. Here we survey and quantify the performance of limited view tomographic algorithms developed for medical applications to reconstruct and precisely measure the dimensional features of a turbine blade and a simple additive manufactured sample using industrial X-ray data and setup.

## 2. X-ray CT theory

Traditionally X-ray tomographic reconstructions are computed using the filtered back-projection approach or its cone beam equivalent the FDK method [2,3]. The main advantages of the filtered back-projection method are its speed of computation, since it relies on the computation of the Fast Fourier Transform and its inverse, and the low computational memory requirements. Despite these advantages the filtered back projection approach does suffer from a number of drawbacks primarily associated with the data acquisition, where many hundreds or thousands of projections that are uniformly distributed over  $360^\circ$  are required to produce accurate reconstructions [2]. The incorporation of prior knowledge (e.g., material properties and geometry) to aid in the reconstruction is usually available in NDE applications but is often difficult to accomplish.

An alternative method to filtered back-projection imaging is to relate the measured projection data to a set of unknown image pixels via a set of algebraic equations [2,10,16]. The measured amplitude of a monochromatic X-ray through an object is given [2] by

$$I = I_0 \exp\left(-\int_r \mu(x,y) ds\right), \quad (1)$$

where  $I$  is the measured X-ray intensity at the detector,  $I_0$  is the intensity of the monochromatic X-ray source and  $\int_r \mu(x,y) ds$  is the ray-path integral through the object with radiographic attenuation  $\mu(x,y)$ . Equation (1) may be discretised and re-written as

$$-\log \frac{I}{I_0} = \sum_{i=1}^n a_i \nu_i, \quad (2)$$

where  $i$  indicates the pixel number,  $a_i$  is the weighting of each pixel based on the length of the X-ray raypaths crossing each pixel and  $\nu$  is the attenuation value of the pixel. For  $m$  observations (2) may be put into matrix form  $b = Ax$  where  $b \in \mathbb{R}^m$  are the X-ray projections,  $A \in \mathbb{R}^{m \times n}$  is a matrix of pixel weights that relates the image to the data projections and

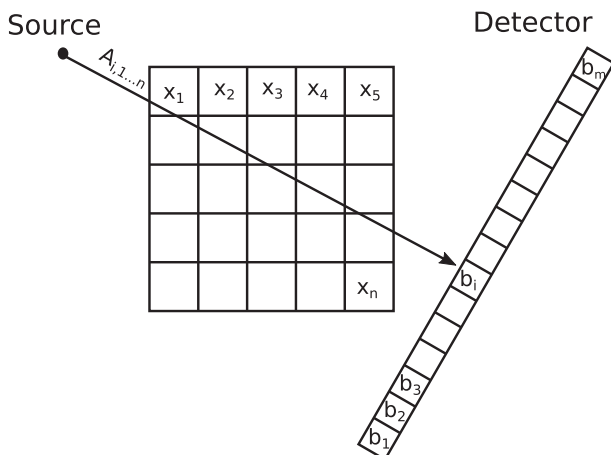


Fig. 1. Illustration of a single ray-path ( $A_{i,n}$ ) passing through a  $5 \times 5$  pixel image array and recorded on the  $i^{\text{th}}$  detector bin.

is often called the projection matrix and  $x \in \mathbb{R}^n$  are the image pixels (Fig. 1). These systems of equations are usually contaminated by noise, underdetermined and singular and therefore must be solved in a least squares sense [2,10].

The main advantage of formulating the tomographic reconstruction as a least squares inverse problem over the more routinely used filtered back-projection method is its flexibility if large numbers of uniformly sampled data over  $360^\circ$  are unavailable from which to produce an adequate reconstruction [2]. In addition, prior knowledge such as positivity ( $x \geq 0$ ) or the minimum and maximum values of each pixel ( $a \leq x \leq b$ ) may be easily incorporated as constraints. The major downside of the algebraic reconstructions are their relatively slow compute times and solution convergence. However, with the ever increasing power of computers and the advent of graphic card processors the application of least squares inversion for X-ray tomography is rapidly becoming viable, particularly if the acquisition time is reduced by obtaining fewer projections e.g., [5].

### 2.1. Algebraic tomographic imaging

The system of equations to be solved in CT tomographic imaging are typically large ( $n = 262144$  for  $512 \times 512$  pixel image) so iterative least squares methods must be used to effectively solve the problem [2,10,16]. A series of algebraic reconstruction methods have been developed over the years which iteratively solve the least squares the problem (for example see Refs. [2,16] for a review). Here, we shall consider two routinely applied reconstruction techniques, the Algebraic Reconstruction Technique (ART) and the Simultaneous Iterative Reconstruction Technique (SIRT), which solve the iterative least squares problem by a series of forward- and back-projections.

The ART method is a row action method where the  $k^{\text{th}}$  iteration of the image is estimated by sweeping through each row of the matrix  $A$  and projecting the solution onto orthogonal hyperplanes e.g., [2,16]. These hyperplanes are defined as  $b_i - a_i^T x^{[k(i-1)]}$  where  $b_i$  is the  $i^{\text{th}}$  component of the data vector  $b$ ,  $a_i$  is the  $i^{\text{th}}$  row of  $A$  written as a column vector and  $x^{[k(i-1)]}$  is the image vector from the  $k^{(i-1)}$  iteration. The update to the image vector is computed by orthogonally back-projecting the hyperplane by multiplying it by  $a_i$ . Practically, the process corresponds to calculating a residual between the measured and the estimated projection made from forward modelling through the current image, then back-projecting this residual to update the image. A single iteration  $x^{[k(i)]}$  is completed once the solution has been updated for all rows of  $A$ . After each row iteration its prior constraints on the solution such as positivity may be applied. The ART algorithm has been shown to have a convergence history which initially improves the solution to better approximations of the true image but at later iterations diverges away from this. This convergence history is known as semiconvergence and has been shown in the ART case to be very fast obtaining a solution in just a few iterations (Algorithm 1 [16]).

#### Algorithm 1 Algebraic Reconstruction Technique (ART)

```

1:  $x^{[0]} = 0$ 
2: for  $k=1, \dots, K$  do
3:   for  $i = 1, \dots, M$  do
4:      $x^{[k(i)]} = x^{[k(i-1)]} + \frac{b_i - a_i^T x^{[k(i-1)]}}{\|a_i\|_2^2} a_i$ 
5:     ▷ Optional positivity ( $x > 0$ ) or box constraint  $x \in [c, d]$ 
6:     if pos = True then
7:       if  $x^{[k+1]} < 0$  then  $x^{[k+1]} = 0$ 
8:     else if box = True then
9:       if  $x^{[k+1]} < c$  then  $x^{[k+1]} = c$ 
10:      else if  $x^{[k+1]} > d$  then  $x^{[k+1]} = d$ 
11:    $x^{[k+1]} = x^{[k(M)]}$ 

```

The second reconstruction procedure considered is SIRT which performs the  $k^{\text{th}}$  update to image vector by simultaneously using all projection data (Algorithm 2). The procedure is similar to the ART approach where the  $k^{\text{th}}$  solution is projected onto a series of hyper-planes  $b - Ax^{[k]}$ , which act to regularise the solution [2]. The image vector is updated by orthogonal back-projection of the projected solution  $x^{[k(i)]}$ . The back-projection is weighted by the diagonal matrices  $C$  and  $R$  with  $c_{ij} = 1/\sum_i A_{ij}$  and  $r_{ii} = 1/\sum_j A_{ij}$ , which compensate for the number of rays that cross each image pixel. The SIRT algorithm converges to the solution in many more iterations than ART but has the advantage of better multi-core parallelism. The implementation of SIRT in the manuscript assumes positivity throughout.

---

**Algorithm 2** Simultaneous Iterative Reconstruction Technique (SIRT)
 

---

```

1:  $x^{[0]} = 0$ 
2: for  $k=1, \dots, K$  do
3:    $x^{[k+1]} = x^{[k]} + CA^T R(b - Ax^{[k]})$ 
4:   ▷ Optional positivity ( $x > 0$ ) or box constraint  $x \in [c, d]$ 
5:   if pos = True then
6:     if  $x^{[k+1]} < 0$  then  $x^{[k+1]} = 0$ 
7:   else if box = True then
8:     if  $x^{[k+1]} < c$  then  $x^{[k+1]} = c$ 
9:     else if  $x^{[k+1]} > d$  then  $x^{[k+1]} = d$ 

```

---

## 2.2. Total variation regularisation

Total variation (TV) reconstructions were first introduced in the image processing literature to the problem of image denoising [11]. The method is based on the principle of compressed sensing whereby an understanding of the signal sparsity may be exploited to perfectly reconstruct the signal of interest with fewer samples than those required by the Shannon-Nyquist theorem [14,15]. A TV reconstruction assumes that the gradient of the signal is sparse making the method ideally suited to reconstructing signals which have sharp edges. The TV of a function  $f(x)$  is defined as e.g., [10,16]:

$$TV(f(x)) = \int_{\Omega} |\nabla f(x)| dx. \quad (3)$$

For the tomographic reconstruction problem we reorder the vector of image pixels  $x$  such that it now represents a 2D array i.e., the image  $x \in \mathbb{R}^{r \times c}$  where  $r$  and  $c$  are the total numbers of row and column pixels respectively. The TV norm of the discretised image may therefore be defined as e.g., [10,12,15]:

$$\|x\|_{TV} = \sum_{i,j} \sqrt{|x_{ij} - x_{i-1,j}|^2 + |x_{ij} - x_{i,j-1}|^2}, \quad (4)$$

where,  $i = 1 \dots r$  and  $j = 1 \dots c$ .

Numerous algorithms have been developed in the recent years which minimise the TV of an image e.g., [4,6–13]. However, many of these algorithms are dedicated to denoising or deblurring an image where the forward operator  $A$  has a specific structure which is efficiently exploited (e.g., identity matrix) [9]. For CT tomographic reconstructions in general, the matrix  $A$  has no exploitable structure and the large nature of the problem to be solved means that a number of the methods developed for TV minimisation are unsuitable for this problem. Iterative gradient descent based methods provide an attractive approach of CT tomographic reconstruction with TV minimisation. The gradient of Equation (4) with respect to each pixel is undefined and thus the following smooth approximation must be used [9,10,12]:

$$\nabla \|x\|_{TV} \approx \frac{(x_{ij} - x_{i-1,j}) + (x_{ij} - x_{i,j-1})}{\sqrt{\delta + (x_{ij} - x_{i-1,j})^2 + (x_{ij} - x_{i,j-1})^2}} \frac{x_{i+1,j} - x_{ij}}{\sqrt{\delta + (x_{i+1,j} - x_{ij})^2 + (x_{i+1,j} - x_{i+1,j-1})^2}} \frac{x_{ij+1} - x_{ij}}{\sqrt{\delta + (x_{ij+1} - x_{ij})^2 + (x_{ij+1} - x_{i-1,j+1})^2}}, \quad (5)$$

where  $\delta > 0$  is small. The selection of  $\delta$  is vital for an accurate reconstruction where values too large lead to smooth reconstructions, and too small a value may result in convergence problems [10]. The value of  $\delta$  is typically estimated via trial and error for each TV algorithm and application. In this article we shall consider the TV reconstruction algorithm of Sidky et al. [12] and the gradient descent method of Barzilai and Borwein [17] with TV regularisation [9,10].

The method of Sidky et al. [12] combines the ART algorithm with a TV minimisation step to solve the following optimisation problem:

$$\underset{x}{\text{minimise}} \quad \|x\|_{TV} \quad \text{subject to} \quad Ax = b, \quad (6)$$

in an iterative manner. The first stage of the method is to minimise  $Ax = b$  in a least squares sense using a single iteration of ART (Algorithm 1). It should be noted that the inclusion of the optional positivity constraint is applied at the end of each ART iteration as opposed to after each row iteration as described in Algorithm 1. The next step in the method is to minimise the TV of the ART reconstruction using a fixed step,  $\gamma$ , gradient descent method. This two step process is repeated for a fixed number of iterations (Algorithm 3). Throughout this manuscript this algorithm will be referred to as ART-TV.

---

**Algorithm 3** ART-TV
 

---

```

1:  $x^{[0]} = 0$ 
2:  $\gamma = 0.05$  ▷ Initialise step length
3: for  $k=1, \dots, K$  do
4:   for  $i = 1, \dots, M$  do ▷ ART algorithm
5:      $x^{[k(i)]} = x^{[k(i-1)]} + \frac{b_i - a_i^T x^{[k(i-1)]}}{\|a_i\|_2^2} a_i$ 
6:    $x^{[ART]} = x^{[k(M)]}$ 
7:    $x^{[k]} = x^{[k(M)]}$ 
8:   if  $x^{[k]} < 0$  then  $x^{[k]} = 0$  ▷ Applying positivity constraint
9:    $d_A = \|x^{[ART]} - x^{[k]}\|_2$ 
10:  for  $j = 1, \dots, N_{\text{grad}}$  do ▷ Gradient TV minimisation
11:     $\hat{v} = \frac{\nabla \|x^{[k(j-1)]}\|_{TV}}{\|\nabla \|x^{[k(j-1)]}\|_{TV}\|}$ 
12:     $x^{[k(j)]} = x^{[k(j-1)]} - \gamma d_A \hat{v}$ 
13:   $x^{[k+1]} = x^{[k(N_{\text{grad}})]}$ 

```

---

The second TV reconstruction algorithm introduces the TV as a regularisation term into a least squares problem and the goal is to minimise:

$$\|Ax - b\|_2^2 + \alpha \|x\|_{TV}, \quad (7)$$

where  $\alpha > 0$  is the regularisation parameter that controls the weighting between the two terms in the equation [10,16]. Equation (7) is solved using a gradient descent method where the step length parameter is estimated using the Barzilai and Borwein method [17]. The inclusion of the Barzilai and Borwein step length provides a computationally cheap approximation to the Newton step, thus improving the rate of convergence and being less affected by ill-conditioning than fixed step gradient methods [9,10]. The Barzilai and Borwein gradient descent method does not guarantee that the objective function decreases with each step. To

overcome this and to guarantee that the objective function does decrease over  $N$  iterations, a non-monotone line search using a backtracking procedure is implemented [9] (Algorithm 4).

---

**Algorithm 4** Gradient descent with TV regularisation
 

---

```

1:  $x^{[0]} = 0$ 
2:  $\epsilon = 0.0001$ 
3:  $p_1 = 0.0001$  ▷ Initialise  $1^{st}$  step length
4:  $N = 10$ 
5: for  $k=1, \dots, K$  do
6:    $\nabla f(x^{[k]}) = \nabla \|Ax^{[k]} - b\|_2^2 + \alpha \nabla \|x^{[k]}\|_{TV}$ 
7:   if  $k > 1$  then
8:      $p_k = \frac{(x^{[k]} - x^{[k-1]})^T (x^{[k]} - x^{[k-1]})}{(x^{[k]} - x^{[k-1]})^T (f^{[k]} - f^{[k-1]})}$  ▷ Barzilai and Borwein step
9:      $\delta = 0.95$ 
10:     $\bar{x} = x^{[k]} - \beta p_k \nabla f(x^{[k]})$  ▷ Gradient TV minimisation
11:     $\hat{f} = \max\{f(x^{[k]}), f(x^{[k-1]}), \dots, f(x^{[k-N]})\}$ 
12:    ▷ Non-monotone line search
13:    while  $f(\bar{x}) \geq \hat{f} - \epsilon \nabla f(x^{[k]})^T (x^{[k]} - \bar{x})$  do
14:       $\beta = \beta \delta$ 
15:       $\bar{x} = x^{[k]} - \beta p_k \nabla f(x^{[k]})$ 
16:     $x^{[k+1]} = \bar{x}$ 

```

---

Typically the value of  $\alpha$  may be estimated by trialing a number of values and producing a trade-off plot, known as an L-curve which compares the solution norm  $\|x\|_{TV}$ , and the residual norm  $\|Ax - b\|_2$  e.g., [10,16]. The value of  $\alpha$  is then selected as ‘bend’ or ‘knee’ in the curve and provides the best trade-off between the solution and residual norms [16].

### 3. Dimensional measurements

The CT algorithms introduced previously are assessed for their applicability to limited view NDE for dimensional measurement with the results compared to traditional filtered back-projection images, manufacturers’ tolerances and physical measurements. In the following tests particular attention is drawn to the accuracy of both internal and external edges of the samples for varying angular sampling ( $1^\circ, 5^\circ, 10^\circ, 15^\circ, 20^\circ$ ). The test also determines at which angular sampling rate ( $1^\circ, 5^\circ, 10^\circ, 15^\circ, 20^\circ$ ) each tomographic method fails to accurately reconstruct the samples.

#### 3.1. Turbine blade

The experimental X-ray CT data for a turbine blade were acquired using a collimated 410 kV, 225  $\mu$ A source and a 4 mm copper filter. The X-rays were recorded on a linear array consisting of 2048, 0.415 mm long pixels, with an exposure time of 1s per frame where each illumination is averaged over 2 frames. The turbine blade was illuminated by 5000 projections with an even angular sampling of  $0.072^\circ$  through  $360^\circ$ . A tomographic image of the turbine blade was obtained using the industry standard FDK method [3] and post processed for any beam hardening artefacts (Fig. 2).

Typically turbine blade wall thickness measurements are made perpendicular to the surface so as to limit small angle errors which scale proportionally to  $\sec(\theta)$  where  $\theta$  is the deviation angle from the normal. However, the complex geometry of the turbine blade means that only a very few walls have parallel surfaces. Given the ambiguity in defining measurement lines perpendicular to the turbine blade walls, here we follow the approach used in other imaging studies where the FDK image is used as a baseline to compare the results of the other reconstruction algorithms e.g., [2,10,12,15,16]. Comparing the reconstruction against a




Fig. 2. Filtered back-projection image of turbine blade. The three red lines indicate the cross-sections through the blade.

baseline image provides a consistent and robust method of quantifying the accuracy of each reconstruction which is independent of any geometrical errors. The accuracy of the reconstructions, in particular, the location of the blade edges and the pixel amplitudes, are assessed by comparing three cross-sections through the turbine blade (Fig. 2).

In order to eliminate artefacts associated with beam hardening an empirical correction was estimated which revises the measured amplitude of the detector based on the travel path through the turbine blade. The correction was estimated by fitting a third order polynomial such that the difference between the expected amplitudes computed by ray-tracing through the corrected filtered back-projection image (Fig. 2) and those measured on the central detector bin are minimised (Fig. 3).

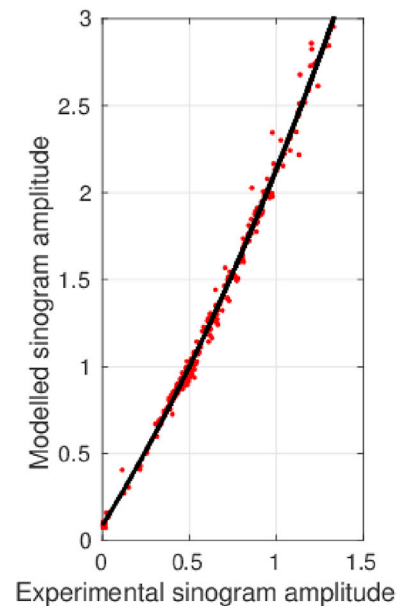


Fig. 3. Comparison of modelled and experimental sinogram amplitudes recorded at the central detector bin (red points). The black line is least squares empirical correction applied to the experimental sinogram amplitude data to account for beam hardening artefacts.

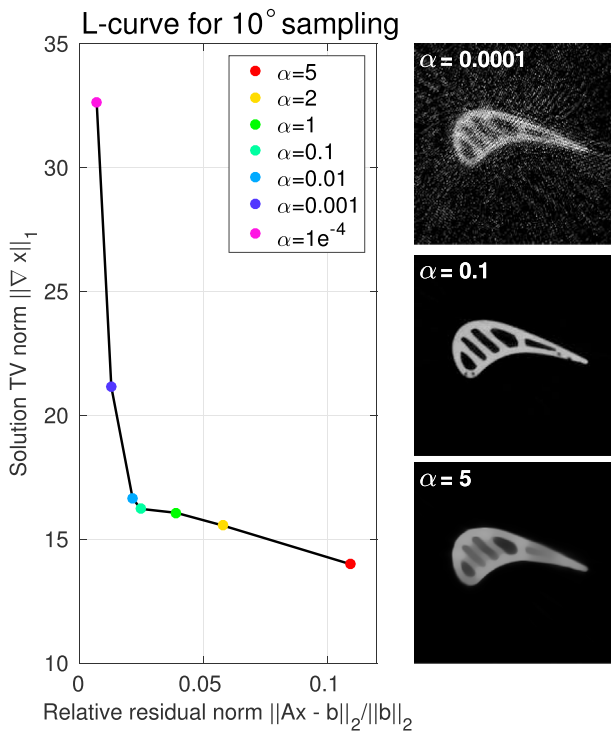


Fig. 4. Trade off curve to determine the parameter  $\alpha$  in Algorithm 4 for tomographic reconstructions with a  $10^\circ$  sampling. The images on the right hand side are turbine blade reconstructions with varying values of  $\alpha$ . Note that larger and smaller values of  $\alpha$  lead to smoother and coarser reconstructions respectively. The optimum value of  $\alpha$  is determined by the ‘bend’ in the curve at 0.1.

Prior to testing the algorithms a number of user defined parameters had to be set; the most important of these are  $\delta$  and  $\alpha$  in Algorithms 3 and 4, and the maximum number of iterations,  $K$ , for all algorithms. The value of  $\delta$  in Algorithms 3 and 4 was determined following a trial and error approach whereby a number of reconstructions of the turbine blade were computed and visually compared with the desired control image. Following this process  $\delta$  was set as  $10^{-8}$  and  $10^{-5}$  for Algorithms 3 and 4 respectively. The regularisation parameter,  $\alpha$ , that controls the weighting between the two terms in Equation (7) must be determined prior to the

analysis. A series of reconstructions were computed using  $\alpha = 1^{-4}, 0.01, 0.1, 1, 2, 5$  with the solution norm and relative residual norm plotted on a trade off curve (Fig. 4). The optimum value of  $\alpha$  is that which adequately reduces the both solution norm and residual norm which is determined by the ‘bend’ or knee in the trade off curve. Following these tests  $\alpha = 0.1$  was deemed to produce the best reconstructions over a number of different angular sampling regimes. The maximum numbers of iterations are algorithm dependent and were determined via a trial and error process.

Fig 5 shows the image reconstructions for an angular sampling of  $1^\circ, 10^\circ$  and  $20^\circ$ . A notable feature of the reconstructions is the significant amounts of noise in the ART, ART with positivity and SIRT reconstructions; this noise is routinely referred to as ‘salt and pepper’ noise which is commonly visible in these reconstructions [2]. Both the TV algorithms produce homogeneous turbine blade image amplitudes in the  $10^\circ$  sampling case whilst displaying smearing towards the tail of the blade and within its internal cavities (Fig 5). This smearing in the images at  $20^\circ$  is less apparent in the ART with positivity and SIRT where the main features of the turbine blade have been preserved.

Fig. 6 shows the reconstructed cross-sections A-A\*, B-B\*, C-C\* for the tomographic inversions with an angular sampling of  $1^\circ, 5^\circ, 10^\circ, 15^\circ, 20^\circ$ . Due to the noise in ART and SIRT reconstructions, it is difficult to assess the quality of each of the reconstructions shown in Fig. 6. The ART method produced adequate cross-sections for angular sampling of  $1-5^\circ$  and is unable to produce meaningful results for greater angular sampling. However, the addition of prior knowledge in the form of positivity constraints improves the reconstructions with all of the different angular sampling methods, producing meaningful reconstructions which closely resemble the true cross-sections.

The TV algorithms produce accurate reconstructions of the turbine blade cross-sections up to an angular sampling of  $10^\circ$ . For an angular sampling of  $15^\circ$  the shape of the cross-sections are correct, however, there are some discrepancies in the amplitudes of the reconstructions particularly in cross-sections A – A\* and C – C\* (Fig. 6). At an angular sampling of  $20^\circ$  the TV methods are unable to accurately reconstruct the cross-sections particularly towards the tail of the blade.

Fig. 7 compares the RMS residual of the control image (Fig. 2) with the tomographic reconstructions for the different choices of angular sampling. The ART has the largest RMS image residual. The gradient descent with TV regularisation (Algorithm 4) produces the best image match of the TV algorithms for angular sampling of  $1-15^\circ$  with the ART-

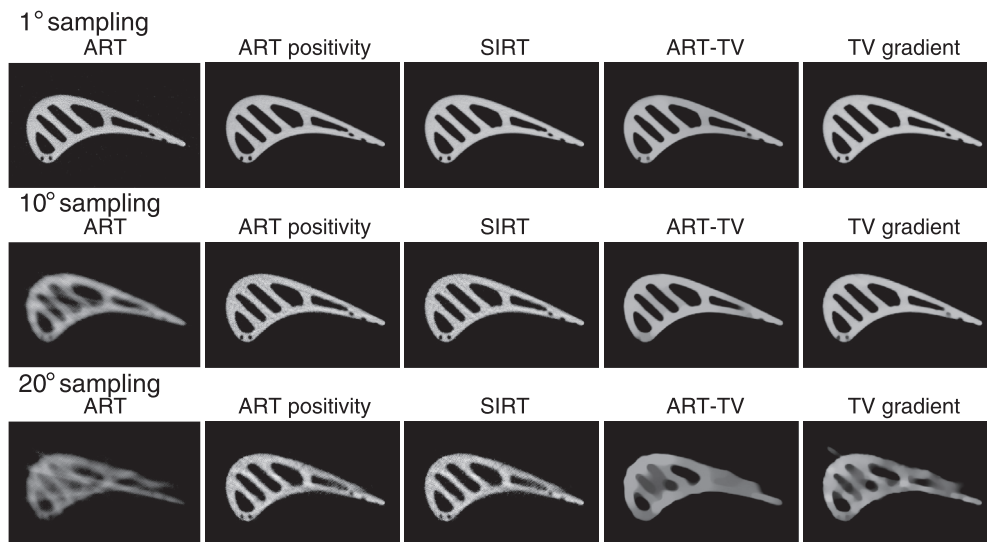


Fig. 5. Turbine blade reconstructions for  $1^\circ, 10^\circ$  and  $20^\circ$  angular sampling. The ‘salt and pepper’ noise is clearly visible in the ART and SIRT images whilst those generated by the TV method have a single value. Note the increased smearing towards the tail of the turbine blade in the TV reconstructions whilst this smearing is reduced in the ART with positivity and SIRT reconstructions.

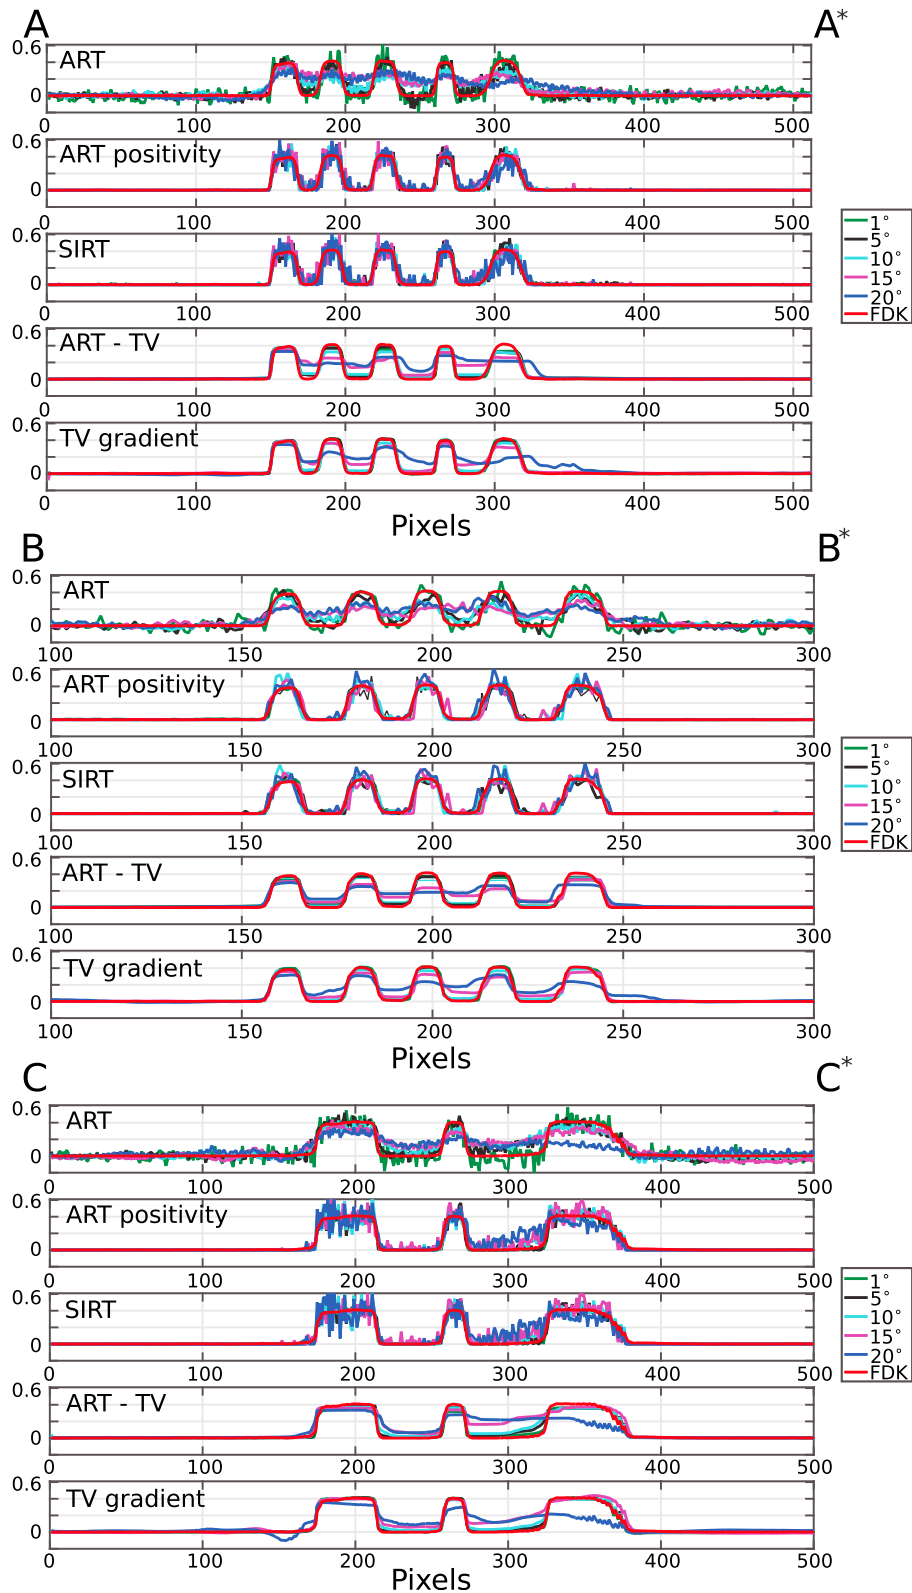


Fig. 6. Tomographic reconstruction of cross-sections A – A\*, B – B\* and C – C\*, for ART, ART with positivity, SIRT, the ART-TV and TV gradient descent method with angular sampling 1°, 5°, 10°, 20°.

TV algorithm (Algorithm 3) producing a slightly improved result at 20°. Both the ART (Algorithm 1) and ART-TV (Algorithm 3) produce RMS residuals which are worse for 1° than for 5° which may be explained by a slight discrepancy in the amplitude of the reconstruction associated with the ‘salt and pepper’ noise.

The final stage of the turbine blade analysis consists of comparing the accuracy of the reconstruction algorithms in calculating the thickness of each part of the turbine blade and comparing this thickness value with the FDK solution [3]. The thickness of each part of the turbine blade was computed along each cross-section, A – A\*, B – B\*, C– C\* with an

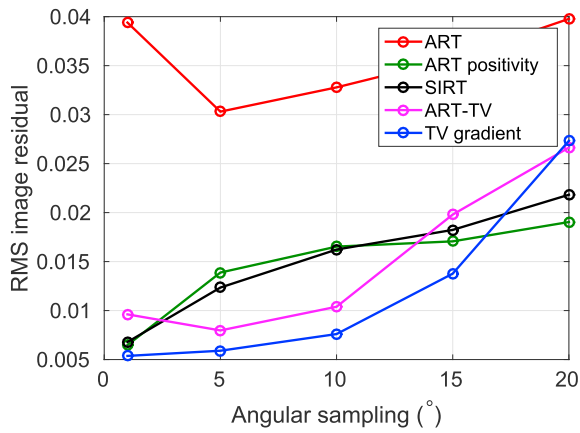


Fig. 7. RMS residual between the control turbine blade and the tomographic reconstructions for angular sampling of 1°, 5°, 10°, 15°, 20°.

automated 1D Canny edge detection algorithm [18]. The automated edge detection algorithm convolves the cross-section of the turbine blade reconstruction with the first derivative of a normalised Gaussian kernel. The edges are defined as the local maxima or minima in the convolved function [18]. This edge detection algorithm has low error rate and works well with noisy data which is the case in the ART, ART with positivity and SIRT (Fig. 6). For the edge detection analysis a Gaussian kernel of 9 pixels with a standard deviation of 2 pixels was deemed optimal. Fig 8 shows the location of the automated thickness estimates and  $\pm 5$  pixel error on the filtered back-projection image (Fig. 2) along with a matrix which indicates the residual between the edge location picked for the reconstruction algorithms and the control as function of the angular sampling.

In all three cross-sections the ART algorithm without positivity performs the poorest of the reconstruction algorithms, with only reliable edge locations being obtained for a sampling of 1°. The inclusion of positivity constraints on the ART and SIRT greatly improves the edge location estimates with a maximum edge residual of  $\pm 2$  pixels ( $\pm 0.23$  mm) up to 10° angular sampling for both methods. Both TV algorithms produce accurate edge detection of less than  $\pm 2$  pixels ( $\pm 0.23$  mm) for an angular sampling of up to 15°. To further understand the cause in the degradation in the image estimates, the relative contrast between each edge was calculated. The relative contrast is defined as:

$$C = \frac{|x_{RMS}^{obj} - x_{RMS}^{noise}|}{x_{RMS}^{obj}} \times 100\%, \quad (8)$$

where  $x_{RMS}^{obj}$  and  $x_{RMS}^{noise}$  are the RMS estimates of the object and noise within a 5 sample window around a valid edge. Fig. 9 summarises the relative contrast for each iterative method as a function of angular sampling. The graphs clearly highlight a decrease in the relative contrast with increasing angular sampling separations. For each cross-section the ART with positivity constraints and SIRT algorithms produce the most stable results with a contrast between 70 and 90%. The TV methods have a contrast of greater than 80% up to a sampling of 10°, before dropping sharply to less than 60% for angular sampling of 10°. These reductions in contrast of the edges with increasing angular sampling correlates with the degradation in the reconstructed images (Fig. 5) and an increase in error for the dimensional analysis (Fig. 8).

### 3.2. Additive manufactured sample

For the additive manufacturing test case a simple example was manufactured consisting of a 70 mm by 70 mm square polymer block with hexagonal voids with a length of 6.93 mm and a separation of 2 mm between each hexagon (Fig. 10). Experimental X-ray CT data were

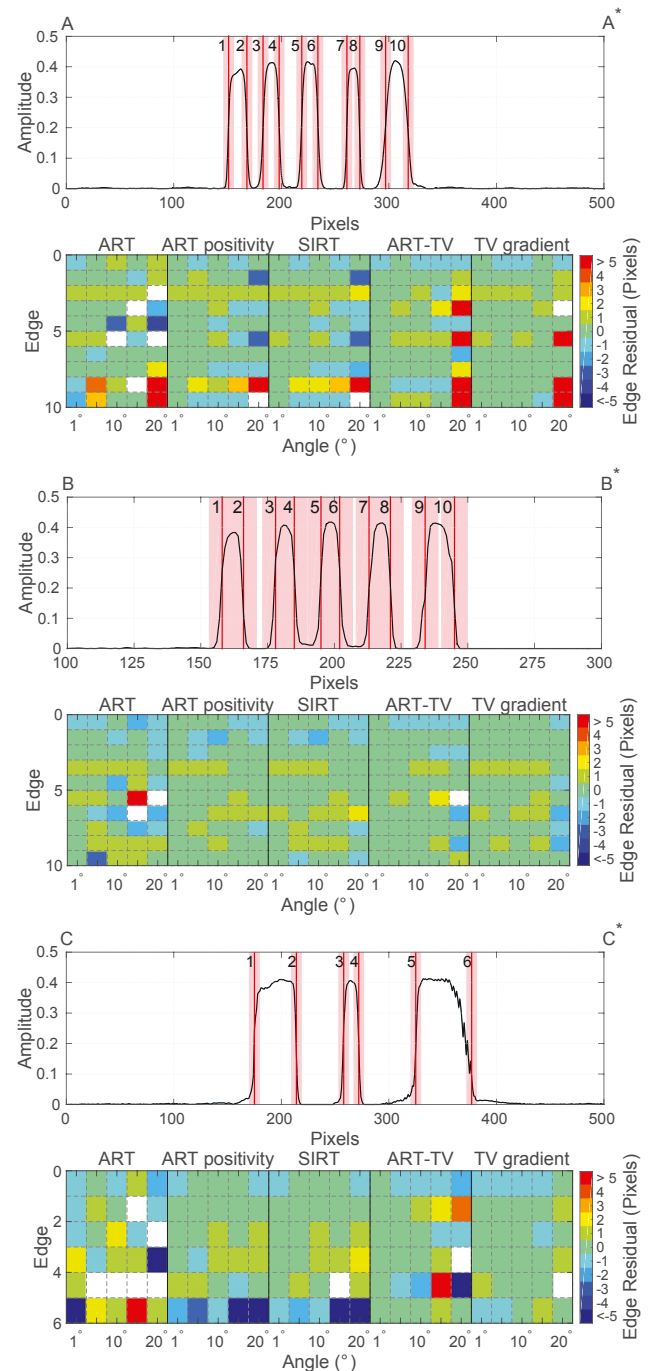


Fig. 8. Top panel shows the control cross section (A – A\*, B – B\* and C – C\*) through the turbine blade. The red lines indicate the location of the automatically determined edges and the shading indicates the location of  $\pm 5$  sample error. The bottom panel shows the picking error in pixels for each of the edges with angular sampling of 1°, 5°, 10°, 15°, 20° for the different tomographic algorithms. The white squares represent an edge which was not picked.

acquired with a 120 kV, 220  $\mu$ A source and a flat panel detector consisting of 2304 by 2304, 0.05 mm pixels, running in continuous mode with an angular velocity of 0.24°/s. In this study we only consider 2D image reconstructions and as such the input data for each algorithm is chosen to be the X-rays recorded on the central horizontal pixels of the detector. As in the previous study a number of user defined parameters were set;  $\delta$  in Algorithms 3 and 4 were determined by trial and error and set to  $10^{-8}$  and  $10^{-6}$  for each algorithm respectively. The parameter  $\alpha$  in Algorithm 4 was again determined using a trade off curve and  $\alpha = 0.1$



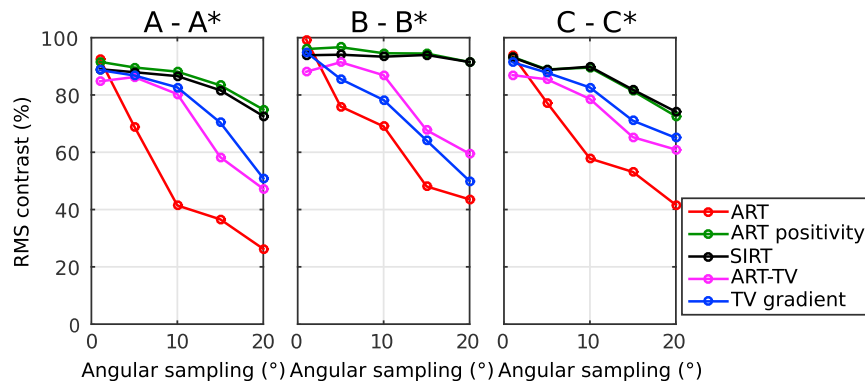


Fig. 9. RMS percentage contrast ratio for the detected edges as a function of angular sampling.




Fig. 10. Schematic illustration of the additive manufactured polymer sample. The dashed line indicates the location of the cross-section through the sample.

was deemed to produce the best reconstructions over a number of different angular sampling regimes.

The reconstructed images of the sample for 1°, 10° and 20° are shown in Fig. 11. A circular ring artefact is clearly visible in all images and is associated with the sample not being entirely contained within the X-ray beam. The ART reconstruction once again produces the poorest image at all angular sampling with significant ‘salt and pepper’ noise present. The inclusion of positivity constraints dramatically improves the ART reconstruction and reduces the image noise. The SIRT reconstruction is similar to the ART with positivity constraints but with reduced noise content. The TV reconstructions produce images with homogeneous pixel values in the 1° reconstructions. In all cases the quality of reconstruction is degraded particularly at the centre of the object where the images are blurred. This blurred patch grows with increasing sampling up to 20° where the image reconstructions have no discernible features present (Fig. 11). The location of this blurred patch at the centre of the reconstruction is unexpected since this region has the greatest number of


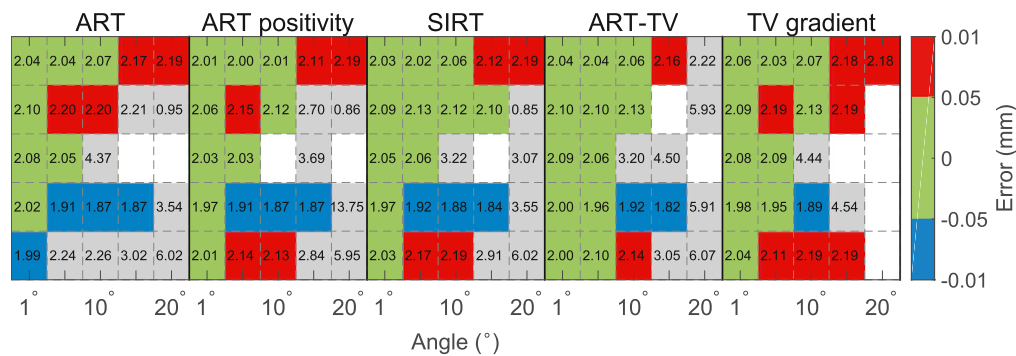


Fig. 11. Additive manufactured sample reconstruction for 1°, 10° and 20° degree angular sampling. The circular ring artefact associated with the sample moving in and out of the view of the detector can clearly be seen on all reconstructions. Once again ‘salt and pepper’ noise is clearly visible in the ART and SIRT images whilst those generated by the TV method have a more homogeneous pixel value. Note the increased smearing at the centre of the sample at 10° and 20° degree angular sampling.



**Fig. 12.** Matrix of wall thickness estimates in mm between the hexagons along a cross-section through the middle of the sample as a function of angular sampling. The numbers in each block are the measured thickness estimates for each reconstruction. The colours represent the difference between the estimated thickness and those measured using the digital micrometers. Values which are coloured green are within the precision bounds of the measured thickness ( $\pm 0.05$  mm). Blue and red colours represent thickness estimates that are up to twice the precision bounds of the physical measurements and are within the manufacturers quoted build resolution ( $\pm 0.2$  mm). The gray coloured squares represent values that are outside the manufacturers build resolution for the sample. White squares represent thickness measurements that could not be reliably computed. The values in each square represent the thickness estimates computed from each reconstruction.

crossing X-rays which is sufficient to produce an accurate reconstruction. It was expected that errors with the reconstruction would have occurred towards the edges of the sample where the X-ray sampling is limited.

To qualitatively assess the capability of the various algorithms to provide accurate dimensional assessments of the sample thickness, measurements were made along a cross-section through the sample (Fig. 10). For the dimensional analysis the internal wall thickness between each of the hexagons along the cross-section was measured using digital callipers and these values were compared with estimates calculated from the cross-section of each reconstruction. The thickness estimates were calculated by computing the edge in each cross-section and associating each of the edges that have an up-going gradient with the nearest down-going edge. Prior to this the 2D images were smoothed using a Gaussian filter of 9 pixels with a standard deviation of 2 pixels. The automatic Canny 1D edge detection algorithm [18] was applied in order to compute each of the edges of the hexagon. The Gaussian kernel used to compute each of the edges had a standard deviation of 2 pixels and a window width of 9 samples. A comparison with the physically measured hexagon wall thickness and the build tolerance of the sample is shown in Fig. 12.

It can be clearly seen that the ART reconstruction is unable to generate accurate thickness measurements which are within the tolerance of the physical thickness measurements. The ART with positivity, SIRT and both TV algorithms are capable of producing accurate measurements for  $1^\circ$  sampling. At  $5^\circ$  only ART-TV is capable of producing thickness measurements within the tolerance of the physical measurements whilst the ART, SIRT and TV gradient methods produce estimates within the build tolerance of the sample. This inability of the methods to

accurately determine the material thickness between the hexagons for sampling rates greater than  $1^\circ$  is also seen in the degradation in the relative contrast (Equation (8)) of the edges with increasing sampling rate (Fig. 13). For this dataset the relative contrast drops to below 70% by a sampling rate of  $10^\circ$  for all methods.

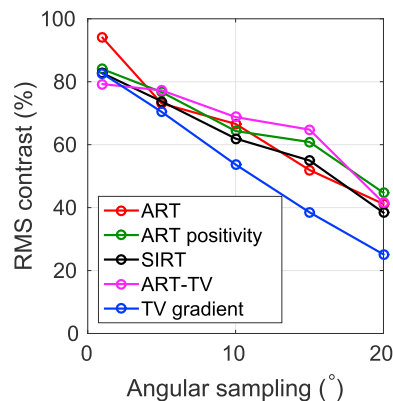
## 4. Discussion

### 4.1. Overview

We have shown that inversion based methods are capable of producing accurate image reconstructions and dimensional measurements with more than an order of magnitude reduction in the number of X-ray projections compared to filtered back-projection methods. The inclusion of prior information either as positivity constraints or as an edge preserving (TV) regularisation greatly improved the quality of the reconstructed image when compared to the unconstrained reconstruction method. Inclusion of this prior information reduced the number of evenly spaced projections necessary for an accurate reconstruction of the turbine blade from 5000 to 36 in the case of ART and SIRT algorithms (Algorithms 1 and 2; Fig. 8) and 24 projections in the TV examples (Algorithms 3 and 4; Fig. 8). However, in the additive manufactured example, such dramatic reduction in the number of projections required to accurately characterise the sample was not observed. 360 evenly spaced projections were required in the ART, SIRT and both TV algorithms (Fig. 13). In both examples we note that the number of projections required to accurately recover the image of the object is greater than the 20 projections suggested in theoretical experiments e.g., [12]. The variability in results between the examples presented indicate that experimental setup and sample complexity plays a significant role on the accuracy of the recoverable image for a given angular sampling.

### 4.2. Image improvement with prior information

In both of the reconstruction experiments presented here, significant improvement to limited view reconstructions are possible by the inclusion of basic prior information into the inverse problem. In NDE applications of X-ray CT, additional information is present (e.g., expected boundary, number of materials present in the specimen etc.) and could be incorporated into any reconstruction algorithm. In addition to knowledge associated with the object to be imaged, statistical information about the expected data noise may also be incorporated to aid the reconstruction process. However, great care must be taken when incorporating prior knowledge into the inversion so as to avoid erroneous information corrupting the final image. Prior knowledge should help direct the optimisation problem whilst also honouring the measured



**Fig. 13.** RMS percentage contrast ratio for the detected edges as a function of angular sampling.

data [19].

#### 4.3. Experimental setup and sample complexity

There was a clear difference in the success of the limited view algorithms between the two experiments presented. Significantly, a blurred patch was observed at the centre of the reconstructions for angular sampling rates greater than  $1^\circ$  in the additive manufactured case. The location of this smearing was unexpected since this region has the greatest X-ray density, which should be sufficient to produce an accurate reconstruction. However, inverse CT imaging is an ill-posed problem whereby small perturbations in the data can have a significant impact on the higher spectral frequencies of the reconstruction, which are generally amplified [10,16]. The amplification of these higher frequencies is often damped by some form of regularisation, which in itself will have an effect on the image. In addition, CT problems are generally under- or mixed-determined meaning some regions of the image may be well resolved whilst others are unstable where multiple solutions exist [20]. We also note that a high X-ray density does not guarantee that the pixel attenuation values are well resolved, particularly where data cannot discriminate between changes in one or more pixels e.g., in the case where multiple parallel rays traverse a pair of pixels, the variations in the measured X-ray attenuation of each X-ray will occur if either or both pixels are changed [20]. In the case of the additive manufactured sample we attribute the main cause of the smearing in the area of greatest X-ray density to the ill-posed nature of the imaging problem.

In order to obtain a satisfactory reconstruction of the object, sufficient spectral frequencies of the data must be included which themselves are a function of the object we wish to CT. Therefore, we attribute the variability in the reconstruction quality between the turbine blade and the additive manufactured sample to the relative complexity of both objects. The metrics used to define the complexity of the objects are the percentage length of an internal feature or structure to its full length and number of internal voids in the object. We define a complex structure as one with low percentage wall lengths whilst having a higher number of voids. In the turbine blade case, the average wall thickness along the three cross sections A – A\*, B – B\*, C – C\* relative to the total blade length is 6.25% and the number of internal voids is 8. In the additive manufactured case wall thickness relative to the total length of the object is 2.88%, with 48 internal voids. Based on the complexity criteria the turbine blade is a less complex object when compared with the additive manufactured sample.

To understand the role of sample complexity on the variability of the results two synthetic datasets were created from  $1024 \times 1024$  segmented reconstructions of the additive manufactured sample. The first utilised the full segmented  $1^\circ$  sampled reconstruction and the second dataset was generated using a circular mask to reduce the object's complexity such that it closely mimicked that of the turbine blade (Fig. 14). We note that in the simplified synthetic example the relative length of the internal walls of the object increased to 4.5% whilst the number of internal voids reduced to approximately 22. The reconstructions were performed on a coarser grid  $512 \times 512$  to mimic the continuous material penetrated by the X-rays and avoid so called 'inverse crimes' where the model grid is the same in both the modelling of the data and inversion [10].

Fig. 14 shows the results of the synthetic reconstructions using the SIRT algorithm (Algorithm 2) with  $10^\circ$  angular sampling. The central area of the full segmented image produced similar smearing artefacts to those observed with the experimental data (Fig. 11). The cross-section of the fully segmented image highlights the inability of the algorithm to accurately reconstruct the amplitude of the central wall and the generation of secondary peaks and smearing at the wall edge. These, in turn, caused the algorithms to fail to generate meaningful wall thickness estimates in the experimental data for angular sampling greater than  $1^\circ$  (Fig. 13). Increasing the number of iterations enhanced the amplitude of the central wall whilst also reducing the edge effects. However, due to the ill-posed nature of the tomographic problem, increasing the number of iterations resulted in the over-fitting and enhancement of any data-noise on the resulting image [16].

The reconstruction of the circular synthetic dataset did not exhibit the smearing artefacts seen in the fully segmented example (Fig. 14). The cross-section through the object showed that the inversion accurately reconstructed both the amplitude and edge location of the object. The improvement in the reconstruction of the circular synthetic dataset relative to the full synthetic dataset highlights the importance of object complexity on the capacity of limited view algorithms to accurately reconstruct the object of interest for dimensional analysis. In addition to the sample complexity the experimental setup and the material properties of the sample play a vital role in the successful application of limited view algorithms. The acquisition of sparse X-ray data where a large number of detector bins have values close to zero, such as in the turbine blade and circular synthetic example, also aid in the successful application of the limited view algorithms. The sparse dataset allows the reconstruction algorithm to focus on the region of interest as seen in the circular synthetic example. Finally, the X-ray source and material properties of the sample play a vital role in controlling the contrast of the

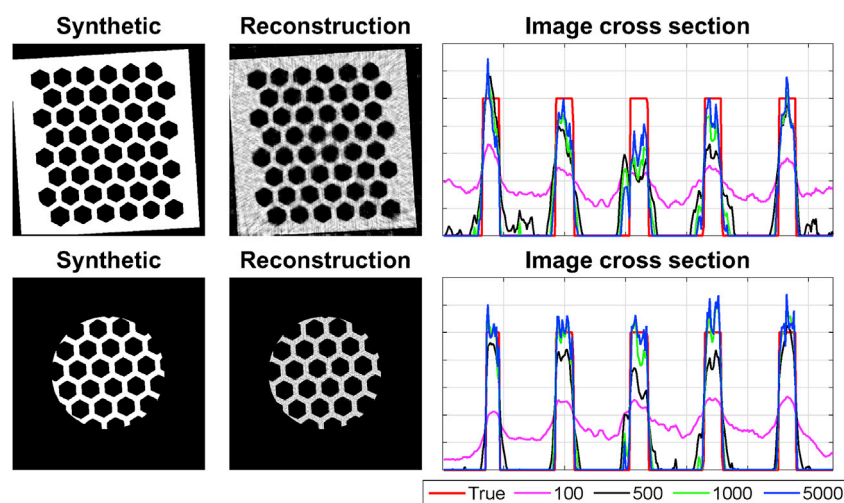


Fig. 14. Comparison of the synthetic sample reconstruction using SIRT algorithm with  $10^\circ$  angular sampling. Note the secondary peaks and amplitude mismatch at the central wall estimate in the synthetic example of the full sample. This mismatch is reduced by increasing the number of iterations. Reducing the size of the sample such that all raypaths homogeneously sample the area of interest, removes the amplitude mismatch and secondary lobes along the central wall.

measured data which subsequently dictates how sharp a material edge may be resolved in the image.

These results and those of the turbine blade and synthetic examples suggest that a simple sample with a minimum wall thickness of 4.5% of the total length of the sample and the acquisition of sparse data where a large number of detector bins have values which are close to zero are necessary to successfully utilise the limited view algorithms discussed.

#### 4.4. Extension to 3D

The algorithms presented here were tested using 2D examples. Each method can be easily scaled to 3D without any structural change to the linear equations  $b = Ax$ . However, the addition of another dimension to the problem leads to a significant increase in the data vector which can increase from  $m$  to  $m^2$  when a single linear detector up-scales to a square flat panel detector, whilst the number of image pixels (or voxels) also increases dramatically. For example, the upscaling of 2D to a 3D image would increase the number of pixels from  $n = 262144$  to  $n = 134217728$  for 512 pixels in each dimension. Based on our experimental findings the improvement in image quality and the reduction in the number of projections required when applying the TV methods versus the ART and SIRT could be outweighed by the need to acquire more data due to the increased complexity of the algorithms.

Cheap graphic processing units (GPU), which possess large scale parallel architecture, are ideally suited to highly vectorised computational problems and must be used in order to upscale the X-ray CT inverse problem to 3D. So called block iterative methods e.g., [21] combine the fast semi-convergence of ART with the multicore properties of SIRT rendering them suitable for large scale inverse problems on GPU architecture. The resulting reconstruction from these block methods are essentially an average of the ART and SIRT results, and based on our results, are an ideal candidate for an inversion based 3D X-ray CT.

In addition to reducing the number of projections required to accurately generate a CT image and improving the hardware, enhancements to the speed of the discussed iterative algorithm could be achieved by changing the structure of the model matrix  $A$  [21,22,23]. The algorithms discussed back-project the calculated solution at each iteration onto a series of hyperplanes and for a consistent system, the solution is the point of intersection of the hyperplanes. To accelerate convergence of ART and SIRT, the rows of the matrix  $A$  should be ordered in such a way that each successive hyperplane is orthogonal to the previous one. In a practical sense this amounts to ordering the projection data so that the current projection data is as independent as possible compared to the previous projection [21,22,23]. A number of approaches have been proposed in the ordering of the rows of  $A$  e.g. [22,23], however similar convergence acceleration may be achieved by a randomised ordering scheme [21].

## 5. Conclusions

Much work has been conducted over the past decade on applying compressive sensing methods to medical X-ray CT as a way of reducing patient radiation exposure. These methods have yet to be applied or rigorously tested to industrial X-ray CT where reduced data acquisition times are desirable to improve manufacturing throughput. This study surveyed and rigorously tested the capability of a number of limited view CT algorithms to accurately reconstruct and precisely measure the dimensional features of an additive manufactured sample and turbine blade. Unlike previous studies where basic synthetic data derived from simple phantoms have been used, the examples studied have used experimental data acquired in an industrial setting.

Our findings highlight the importance of prior information, either as pixel value constraints or as an edge preserving regularisation term, in improving the image reconstruction with limited angular sampling when compared with the unconstrained reconstruction method. All methods were capable of reducing the number of X-ray projections to produce an

accurate reconstruction of the sample by an order of magnitude or more. In the turbine blade example the number of projections were reduced from 5000 to 36 for the case of positivity constrained ART and SIRT algorithms and to 24 projections in the TV examples. However, in the additive manufactured example this dramatic reduction in the number of projections necessary to produce accurate dimensionality measurements was not observed (3000–360 in the positivity constrained ART, SIRT and TV cases). This variability in the number of projections necessary to produce an accurate reconstruction highlights the additional importance of the experimental setup and sample complexity in reducing X-ray projections. It was identified that simple structures which have a minimum feature size of 4.5% relative to the length of the object and the acquisition of sparse X-ray data are necessary to maximise the limited view algorithms presented. The methods presented may be easily extended from 2D to 3D with no significant change in the underlying inverse problem. With the rapid application and design of complex and irregularly shaped components for safety-critical applications X-ray CT will continue to grow as a NDE inspection tool. Limited view CT algorithms offer a method of reducing data acquisition time and alleviate manufacturing throughput bottlenecks which may arise without compromising the reconstruction accuracy.

## Acknowledgements

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. The authors would like to thank Prof Michael Lowe for his helpful discussion and critical reading of the manuscript and Gabor Gubicza for designing the additive manufactured specimen.

## References

- [1] Kruth JP, Bartscher M, Carmignato S, Schmitt R, De Chiffre L, Weckenmann A. Computed tomography for dimensional metrology. *CIRP Annals-Manuf. Technol* 2011;60(2):821–42.
- [2] Kak AC, Slaney M. Principles of computerized tomographic imaging. IEEE Press; 1988.
- [3] Feldkamp LA, Davis LC, Kress JW. Practical cone-beam algorithm. *JOSA A* 1984; 1(16):612–9.
- [4] Xu F, Mueller K. Accelerating popular tomographic reconstruction algorithms on commodity pc graphics hardware. *IEEE Trans Nucl Sci* 2005;52(3):654–63.
- [5] Han X, Bian J, Eaker DR, Kline TL, Sidky EY, Ritman EL, et al. Algorithm-enabled low-dose micro-ct imaging. *IEEE Trans Med Imaging* 2011;30(3):606–20.
- [6] Beck A, Teboulle M. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Trans Image Process* 2009;18(11): 2419–34.
- [7] Chambolle A. An algorithm for total variation minimization and applications. *J Math Imaging Vis* 2004;20(1):89–97.
- [8] Goldfarb D, Yin W. Second-order cone programming methods for total variation-based image restoration. *SIAM J Sci Comput* 2005;27(2):2005.
- [9] Jensen TL, Jørgensen JH, Hansen PC, Hansen SH. Implementation of an optimal first-order method for strongly convex total variation regularization. *BIT Numer Math* 2012;52:329–56.
- [10] Müller JL, Siltanen S. Linear and nonlinear inverse problems with practical applications. SIAM; 2012.
- [11] Rudin LI, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. *Phys D Nonlinear Phenom* 1992;60(1):258–68.
- [12] Sidky EY, Kao C-M, Pan X. Accurate image reconstruction from few-views and limited-angle data in divergent-beam ct. *J X-ray Sci Technol* 2006;14:119–39.
- [13] Vogel CR, Oman ME. Iterative methods for total variation denoising. *SIAM J Sci Comput* 1996;17(1):227–38.
- [14] Donoho DL. Compressed sensing. *IEEE Trans Inf Theory* 2006;52(4):1289–306.
- [15] Candès EJ, Romberg J, Tao T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans Inf Theory* 2006;52(2):489–509.
- [16] Hansen PC. Discrete inverse problems: insight and algorithms. SIAM; 2010.
- [17] Barzilai J, Borwein JM. Two-point step size gradient methods. *IMA J Numer Anal.* 1988;8(1):141–8.
- [18] Canny J. A computational approach to edge detection. *IEEE Trans Pattern Anal. Mach Intell* 1986;8(6):679–98.
- [19] Scales JA, Tenorio L. Prior information and uncertainty in inverse problems. *Geophysics* 2001;66(2):389–97.
- [20] Rawlinson N, Fichtner A, Sambridge M, Young MK. Seismic tomography and the assessment of uncertainty. *Adv Geophys* 2014;55:1–76.

- [21] Sørensen HH, Hansen PC. Multicore performance of block algebraic iterative reconstruction methods. *SIAM J Sci Comput* 2014;36(5):C524–46.
- [22] Herman GT, Meyer LB. Algebraic reconstruction techniques can be made computationally efficient. *IEEE Trans Med Imaging* 1993;12(3):600–9.
- [23] Guan H, Gordon R. A projection access order for speedy convergence of art (algebraic reconstruction technique): a multilevel scheme for computed tomography. *Phys Med Biol* 1994;39(11):2005–22.