Research

Cite this article: Blumensath T, Boardman R. 2015 Non-convexly constrained image reconstruction from nonlinear tomographic X-ray measurements. *Phil. Trans. R. Soc. A* **373**: 20140393.
<http://dx.doi.org/10.1098/rsta.2014.0393>

Accepted: 17 March 2015

One contribution of 11 to a theme issue 'X-ray tomographic reconstruction for materials science'.

Subject Areas:

image processing

Keywords:

compressed sensing, inverse problems, nonlinear constrained optimization, tomography

Author for correspondence:

Thomas Blumensath

e-mail: thomas.blumensath@soton.ac.uk

Non-convexly constrained image reconstruction from nonlinear tomographic X-ray measurements

Thomas Blumensath¹ and Richard Boardman²

¹ISVR Signal Processing and Control Group, and ²Engineering Materials, University of Southampton, Southampton SO17 1BJ, UK

The use of polychromatic X-ray sources in tomographic X-ray measurements leads to nonlinear X-ray transmission effects. As these nonlinearities are not normally taken into account in tomographic reconstruction, artefacts occur, which can be particularly severe when imaging objects with multiple materials of widely varying X-ray attenuation properties. In these settings, reconstruction algorithms based on a nonlinear X-ray transmission model become valuable. We here study the use of one such model and develop algorithms that impose additional non-convex constraints on the reconstruction. This allows us to reconstruct volumetric data even when limited measurements are available. We propose a nonlinear conjugate gradient iterative hard thresholding algorithm and show how many prior modelling assumptions can be imposed using a range of non-convex constraints.

1. Introduction

Tomographic X-ray imaging is used widely in areas as diverse as medical diagnostics and non-destructive material testing. Using sets of X-ray projection images taken from different directions, tomographic inverse algorithms reconstruct volumetric images of an object's internal structure. Most algorithms in use today rely on a linear X-ray attenuation model which assumes that the X-ray source produces mono-energetic X-ray photons. However, most laboratory sources produce polychromatic X-ray beams and so, approximate pre-processing algorithms are often used to correct for some of the nonlinear effects prior to volumetric reconstruction.

Figure 1. A simulated phantom made up of four materials which at 80 keV have mass attenuation coefficients of 0.2, 0.23, 2.77 and $4.46 \text{ cm}^2 \text{ g}^{-1}$, respectively. To show the difference between the two low-attenuating materials, the grey scale has been compressed (see colourbar).

While these pre-processing techniques often work well when imaging objects in which X-ray attenuation does not vary too widely, in settings where there is a mix of highly attenuating regions and low-attenuating regions, severe artefacts are introduced in the image reconstruction. Examples include metal implants imaged in human tissue or carbon composite components with metal inclusions.

As a motivating example, consider the simulated phantom shown in figure 1, which is made up of four materials, two materials with low X-ray attenuation factors that differ very little between materials (an aluminium outer ring with a silicon inner disc) and two highly attenuating materials (silver and gold inclusions). Figure 1 shows the mass attenuation for these materials at 80 keV. To differentiate the different materials, the greyscale has been compressed to better show the difference in the high- and low-attenuating regions (see the colourbar).

We simulated two different X-ray acquisitions, one with a monochromatic source at 80 keV and one with a polychromatic source with a maximum energy of 130 keV. Seven hundred and twenty fan beam projections were calculated for each system and reconstructed with the filtered backprojection (FBP) algorithm (using Matlab's image processing toolbox functions) and the algebraic reconstruction technique (ART; using the randomized Kaczmarz algorithm from the AIRTtools toolbox [1]). The results are shown in figure 2. While our compressed colour map shows reconstruction artefacts in the monochromatic reconstructions on the right, the left polychromatic reconstructions clearly show additional artefacts typical for polychromatic X-ray reconstructions. These artefacts can be seen as shadows between high-attenuating regions (the darker area in the centre of the reconstruction) and brighter streaks that tangentially connect the boundaries of different high-attenuating objects.

The realization that polychromatic X-rays lead to a nonlinear transmission models is typically known as beam-hardening [2]. Many different techniques for the correction of this effect have been proposed over the years, from physical filtering of the X-ray source spectrum to reduce the nonlinear relationship between path length and X-ray absorption [2] to direct inversions of the full nonlinear model (2.6).

For single material objects, intensity linearization can be performed [2] based on phantom scans [3], initial uncorrected reconstructions [3,4] or simulation studies [5]. Measurements are interpolated [6–9] and this process is often iterated [6,7,10]. Intensity corrections can also be calculated for multiple material objects [3,8,11–16]. The initial reconstruction is segmented which again provides initial path length estimates. This process can then also be iterated. For objects made up of two materials with known attenuation coefficients, it is also possible to instead estimate the volume fraction for each spatial location [6,17,18].

Figure 2. Reconstruction of the simulated phantom using two algorithms (the FBP algorithm and the ART). The reconstructions in (a,c) use simulations of a polychromatic X-ray source, whereas the reconstructions in (b,d) use a monochromatic source. To show the difference between the two low-attenuating materials, the grey scale has been compressed (see colourbar).

2. The nonlinear X-ray observation model

We will here propose an advanced reconstruction algorithm that explicitly models nonlinear X-ray attenuation. Furthermore, we will impose additional regularization constraints, which will not only be able to overcome the general ill-conditioning of the nonlinear inverse problem, but will also allow us to recover volumetric images when few X-ray measurements can be taken.

To a first approximation, ignoring X-ray scattering, the fraction of photons with energy E emitted from the X-ray source that travel towards a given detector element $I_0(\mathbf{r}, E)$ and the number of photons with energy E counted at the detector $I(\mathbf{r}, E)$ are a function of the mass attenuation coefficient $\mu(\mathbf{z}, E)$ which describes the likelihood with which X-ray photons with energy E are absorbed or scattered at location $\mathbf{z}$. We here use $\mathbf{r}$ to label the line between the X-ray source and the detector element. As most X-ray detectors cannot distinguish photons of different energies, the detector output integrates over all energy levels.

This leads to the following observation model used previously in [6,13,19–29] and [30]:

$$I(\mathbf{r}) = \int I(\mathbf{r}, E) dE = \int I_0(E) \exp \left(- \int_{\mathbf{r}} \mu(\mathbf{z}, E) d\mathbf{z} \right) dE, \quad (2.1)$$

where we assume that $I_0(\mathbf{r}, E)$ is independent of $\mathbf{r}$ so that we write $I_0(E)$ instead.¹

¹In practice, $I_0(\mathbf{r}, E)$ is a function of $\mathbf{r}$; however, the detector is normally calibrated so that photon counts are scaled to compensate for this non-uniformity in the source.

(a) The monochromatic model

For future reference, if we had a monochromatic source or if we were able to measure X-ray intensities $I(\mathbf{r}, E)$ in narrow energy bands, then we would have a simple linear relationship

$$-\ln \frac{I(\mathbf{r}, E)}{I_0(E)} = \int_{\mathbf{r}} \mu(\mathbf{z}, E) d\mathbf{z}. \quad (2.2)$$

This model is the basis for most traditional tomographic reconstruction algorithms such as FBP and algebraic reconstruction.

Algebraic methods start with a discretization of (2.2). Using $\mathbf{x}(E) \in \mathbb{R}^N$ as the unknown coefficient vector that approximates $\mu(\mathbf{z}, E)$ in a spatial basis and using the vector $\mathbf{a}(\mathbf{r}_i)$ to describe the integrals of each individual basis function along the path $\mathbf{r}_i$. Let

$$\bar{\mathbf{y}} = - \left[\ln \frac{I(\mathbf{r}_i, E)}{I_0(E)} \right]_{i \in [1, \dots, M]} \quad (2.3)$$

be the vector containing all our normalized and log-transformed projection measurements. If we furthermore stack all vectors $\mathbf{a}(\mathbf{r}_i)^T$ into a matrix $\mathbf{A}$, then we have the linear model

$$\bar{\mathbf{y}} \approx \mathbf{A}\mathbf{x}. \quad (2.4)$$

The bar over the vector $\bar{\mathbf{y}}$ here reminds us that this is the linear model for monochromatic X-ray interaction.

(b) The polychromatic model

To use the nonlinear model in (2.1) for reconstruction, we also discretize the photon energies E [21]

$$I(\mathbf{r}_i) \approx \sum_{l \in [1, \dots, L]} I_0(E_l) \exp(-\mathbf{a}(\mathbf{r}_i)^T \mathbf{x}(E_l)), \quad (2.5)$$

where the sum is now over L different energy bands E_l .

We again collect the normalized and log-transformed observations $I(\mathbf{r}_i)$ into a single vector

$$\mathbf{y} = -\log \left[\frac{I(\mathbf{r}_i)}{\sum_l I_0(E_l)} \right]_{i \in [1, \dots, M]}. \quad (2.6)$$

Using the normalized X-ray intensity vector $\mathbf{I} = \left[\frac{I_0(E_l)}{\sum_l I_0(E_l)} \right]_{l \in [1, \dots, L]}$, we thus have the measurement model

$$\mathbf{y} \approx -\log([e^{-\mathbf{A}\mathbf{X}}]\mathbf{I}), \quad (2.7)$$

where the exponent is evaluated component wise and where we introduce the matrix notation $\mathbf{X} = [\mathbf{x}(E_1), \dots, \mathbf{x}(E_L)] \in \mathbb{R}^{N \times L}$.

It is important to note that the columns in $\mathbf{X}$ are related to X-ray absorption at a certain X-ray energy level and so are the different entries in $\mathbf{I}$. To use this model for a fixed set of energies, we would thus either need to know the source intensities at these energy levels or estimate $\mathbf{I}$ at the same time as $\mathbf{X}$ from the measurements. If we did this, then we could incorporate additional knowledge about the X-ray absorption properties of materials known to be present in a sample. However, if we estimate $\mathbf{X}$ without any prior assumption on the materials' absorption properties, then quantization of the energy levels is to some extent arbitrary. Thus, instead of quantizing the energy levels, we might as well quantize X-ray intensities $\mathbf{I}$, which in turn would correspond to an unknown quantization of the energy levels. We can then estimate the average X-ray intensity and detector efficiency using a scan without the object or from a region of the detector not in the

Figure 3. Reconstruction of the simulated phantom from polychromatic X-ray measurements by optimizing equation (2.8).

shadow of the object. This estimate $\tilde{I} \approx \sum_i (I_0(E_i))$ can then be used in the denominator of (2.6) to normalize the observed X-ray intensities before taking the logarithm.

(c) Measurement errors

So far, we have modelled the X-ray physics using the Beer–Lambert law. While this model is motivated by the photon nature of X-rays, it does not directly model the quantized, probabilistic nature of X-ray generation and detection, which follow Poisson-like distributions. In certain settings, especially those in which few X-rays are measured along each X-ray path, the influence of the X-ray statistics becomes more important and it is then advantageous to directly model the statistical distribution to formulate the inverse problem. In other settings, where the number of measured X-ray photons becomes large, the central limit theory allows us to make Gaussian noise assumptions.

Additional uncertainty in the measurement comes from system noise and from discretization and quantization effects. Spatial discretization is further degraded as X-ray sources and detectors are not ideal points in space but cover extended areas, while X-ray intensities are stored digitally in quantized form. Scattering [31] is another significant process that can degrade X-ray projection images and lead to reconstruction artefacts.

In this paper, we will make Gaussian noise assumptions, though our approach is readily modified to use Poisson models when appropriate. We will start the development of our approach using a maximum-likelihood estimate, though again, extension to maximum *a posteriori* estimates follows similar lines. Assuming a fixed noise variance, the maximum-likelihood estimate for a Gaussian model with observations described by (2.7) is then equivalent to the following minimization problem:

$$\min_{\mathbf{X}} \|\mathbf{y} - \log([\mathbf{e}^{-\mathbf{A}\mathbf{X}}]\mathbf{I})\|^2. \quad (2.8)$$

To demonstrate the benefits of this nonlinear optimization problem over the linear model, we used the same simulated polychromatic X-ray measurements that were used on the left-hand side of figure 2 and optimized (2.8) using 200 iterations of a gradient descent algorithm with fixed step size. We initialized the matrix $\mathbf{X}$ as $\mathbf{x}(\boldsymbol{\phi} + \boldsymbol{\theta})^T$, where $\mathbf{x}$ is the ART reconstruction shown in figure 2 and where $\boldsymbol{\phi}$ and $\boldsymbol{\theta}$ are the photoelectric effect and Compton scattering basis functions proposed in [25].

The results are shown in figure 3, where we show the estimated mass attenuation averaged over energy. This average is calculated as

$$-\log([e^{-X}]\mathbf{I}), \quad (2.9)$$

where the exponent and logarithm are again evaluated component wise.

As can be seen, the beam hardening artefacts have been reduced though the performance is not as good as those achieved with monochromatic measurements. Part of this is due to the slow convergence of the gradient descent algorithm. This convergence is also the reason why we had to choose a sensible starting point for the algorithm, using a linear reconstruction to initialize the spatial distribution and a simple model of the energy-dependent X-ray attenuation to model the energy distribution.

3. Regularization with non-convex constraints

When using model (2.6) additional assumptions are often imposed. Let us write the matrix $\mathbf{X}$ as a factorization $\mathbf{X} = \mathbf{D}\mathbf{S}$, where $\mathbf{D}$ and $\mathbf{S}$ are matrices. The matrices $\mathbf{D}$ and $\mathbf{S}$ can now be used to model different assumptions. For example, the i th column in $\mathbf{D}$ could contain the density of material i within each of the voxels and the i th row in $\mathbf{S}$ could contain the discretized attenuation spectrum for this material. Thus, if we assume that there are no more than J materials, then the matrix $\mathbf{X}$ will have rank at most J . Alternatively, the formulation can be used to describe the model used in [23–26], where $\mathbf{S}$ has two rows, one representing a known basis function to model the photoelectric effect and one basis function to model Compton scattering [25].

Different model assumptions also often lead to additional constraints on $\mathbf{D}$. For example, the assumption in [21] that each voxel contains a single material leads to the constraint that $\mathbf{D}$ must have only a single entry per row. On the other hand, for the model where $\mathbf{S}$ contains basis functions for the photoelectric effect and for Compton scattering, Stonestrom *et al.* [24] proposed that the two coefficients in each row of $\mathbf{D}$ could be modelled using a single parameter. A similar parametrization of $\mathbf{D}$ is also used in [27].

(a) Non-convex regularization

The constraints on $\mathbf{X}$ discussed in the previous paragraph, namely the requirement that $\mathbf{X}$ should be low rank or that $\mathbf{D}$ should have one sparse columns can be enforced using non-convex regularization constraints. Regularization is a common approach used to stabilize tomographic reconstruction, preventing noise amplification and reducing artefacts due to non-ideal sampling. Traditional Tikhonov regularization [32] as well as now standard total variation (TV) regularization [33] typically lead to convex optimization problems for which increasingly efficient algorithms are being developed [34]. Non-convex constraints, on the other hand, while offering increased flexibility in the way prior information can be modelled and exploited for regularization, lead to non-convex and often combinatorial optimization problems, which pose significantly more challenges. However, the advent of compressed sensing has, over the last 10 years, shown that many non-convex optimization problems can nevertheless be solved efficiently using polynomial time algorithms [35,36]. Initially, compressed sensing was built around sparse, finite dimensional data models, though this has been extended to more general non-convex constraints and infinite dimensional models [37,38]. Another important extension has been the recent move away from the linear forward model of traditional compressed sensing to nonlinear forward models [39]. These recent innovations can now also be applied to X-ray tomographic reconstruction under nonlinear observation models, such as that of equation (2.7). Before we look at the computational approaches that can be used to solve this inverse problem, we first introduce a range of possible non-convex constraints that might be useful in the tomographic reconstruction setting. As we will see, our algorithm will require the computation of projections of matrices or vectors onto the closest element in the non-convex constraint set so that we will pay particular attention to how this projection can be done efficiently for the different constraints we introduce.

(b) Sparsity

Sparsity is a powerful constraint. In its simplest form, sparsity constraints restrict model parameters to be predominantly zero so that the effective degrees of freedom of a model can be much smaller than the number of model parameters.

(i) Sparsity of a vector

The canonical sparse model is a vector model, where vector entries are restricted to contain few non-zero entries. More formally, consider the vector $\mathbf{x} \in \mathbb{R}^N$ in Euclidean space. We can denote the number of non-zero entries in this vector by $\|\mathbf{x}\|_0$. A sparsity constraint could then be written as

$$\|\mathbf{x}\|_0 \leq K, \quad (3.1)$$

for some $K < N$. This is a non-convex constraint. If we have two K sparse vectors $\mathbf{x}_1$ and $\mathbf{x}_2$, then their convex combination $\lambda\mathbf{x}_1 + (1 - \lambda)\mathbf{x}_2$ is not guaranteed to be K -sparse. However, if both $\mathbf{x}_1$ and $\mathbf{x}_2$ have their non-zero elements at the same locations, then the linear combination will also be K -sparse. A K -sparse constraint thus restricts elements to lie on one of several subspaces. The set of all K -sparse vectors is thus a union of subspaces.

For an arbitrary vector $\mathbf{x}$, calculation of the K -sparse approximation $\hat{\mathbf{x}}_K$ with the smallest least-squares error can be done extremely efficiently. All that is required is to sort the entries in $\mathbf{x}$ by decreasing magnitude and set all but the K largest elements of $\mathbf{x}$ to zero. We will call the calculation of the best K -sparse approximation of a vector a *projection* of the vector onto the union of subspaces.

The power of sparsity constraints comes from the fact that they are not restricted to the canonical basis, but that we can restrict coefficients in any basis. By carefully choosing the right basis, we are thus often able to model and constrain different characteristics of $\mathbf{x}$. For example, for our tomographic X-ray problem, wavelet bases can provide good sparse models. Let Φ be a basis matrix and define $\mathbf{s}$ such that $\mathbf{x} = \Phi\mathbf{s}$. We then say that $\mathbf{x}$ is sparse in the basis Φ if $\mathbf{s}$ is sparse. If Φ is orthonormal, then it is again easy to commute the best K -sparse approximation of $\mathbf{x}$ in the basis Φ . This is done by calculating $\mathbf{s} = \Phi^T \mathbf{x}$, which is then again thresholded by keeping only the K largest entries in magnitude of $\mathbf{s}$. However, if Φ is no longer orthonormal, or if Φ is over-complete, that is Φ has more than N columns, then the estimation of the best K -sparse approximation in this basis becomes a combinatorial problem in general and suboptimal algorithms have to be used.

(ii) Sparsity in matrices

We are interested in regularizing the matrix valued optimization problem (2.8) and for matrices, sparsity constraints are even more flexible. Mirroring the vector sparsity model, the simplest constraint could be a restriction on the total number of non-zero entries; however the matrix structure reflects underlying data properties and so it makes sense to treat columns separately. We might for example restrict matrix rows or columns to be sparse. A K -row-sparse matrix is then a matrix where all but K rows contain only zero entries. Again, for an arbitrary matrix $\mathbf{X}$ to compute the best approximation with a K -row-sparse matrix in terms of the Frobenius norm, we simply order the rows by the size of their Euclidean norm. We then set all but the largest K rows to zero, thus again projecting the matrix onto the closest K -row-sparse matrix.

Extensions of this principle again allow for the use of basis matrices to capture row or column structure. For example, a matrix $\mathbf{X}$ is row sparse in a basis Φ if $\mathbf{X} = \Phi\mathbf{Z}$ with $\mathbf{Z}$ being row sparse. Similarly, a matrix $\mathbf{X}$ is column sparse in a basis Φ if $\mathbf{X} = \mathbf{Z}\Phi$ with $\mathbf{Z}$ being column sparse. Again, for orthonormal basis Φ , calculating the best column or row sparse matrix in Φ can be done efficiently.

(iii) Block and tree sparse models

Often, additional structure can be imposed on a sparse model, further reducing the size of the constraint set. Block sparse vector models group elements of a vector into either non-overlapping

or overlapping blocks. A K sparse block model then restricts the number of blocks that can contain non-zero entries. For arbitrary vectors $\mathbf{x}$, computing the best block sparse approximation is easy if the model has non-overlapping blocks, but there are no efficient methods for overlapping blocks.

A special subset of block sparse models is a tree sparse model where a tree structure defines the blocks. A K -tree-sparse vector is then a rooted subtree with no more than K nodes. Optimal tree sparse models are difficult to compute, but fast approximate algorithms are available [40].

(c) Rank constraints

For matrix models, another powerful non-convex constraint is the rank constraint [41]. By restricting matrices to have a rank of at most K , we in effect constrain the sparsity of the singular values. The difference from sparse models is however that the left and right singular vectors can vary continuously. While low-rank matrices with the same left and right singular vectors again lie on the same subspace, there are infinitely many potential left and right singular vectors and so there are infinitely many of these subspaces. Low-rank models thus restrict matrices to lie in one of infinitely many subspaces. For a general matrix $\mathbf{X}$, computing the best rank K approximation can be done using the singular value decomposition followed by a hard thresholding of the singular values based on their magnitude, which again computes the projection onto the union of subspaces model.

(d) Non-negative matrix factorization constraint

An extension of a low-rank constraint is the non-negative matrix factorization constraint. For a non-negative matrix $\mathbf{X}$, a non-negative matrix factorization is a decomposition of $\mathbf{X}$ into two positive factors $\mathbf{D} \in \mathbb{R}^{M \times K}$ and $\mathbf{S} \in \mathbb{R}^{K \times N}$. For K smaller than M and N , this is again a low-rank factorization; however, the positivity constraint on $\mathbf{D}$ and $\mathbf{S}$ imposes additional constraints. Computing the best approximation of a positive matrix as a product of two positive factors is not a straightforward task, yet approximate solutions can be computed using non-negative matrix factorization algorithms [42].

4. Iterative hard thresholding and nonlinear iterative hard thresholding

There are two common approaches to deal with non-convex constraints such as those discussed above. The strategy advocated in the original compressed sensing literature is convex relaxation. Here, a convex constraint set is specified that is as close as possible to the non-convex constraint. Surprisingly, in several sparse inverse problems, solving the relaxed convex problems is guaranteed to find good solutions to the original non-convex constraint. However, not all non-convex constraints lend themselves to efficient convex relaxation. Furthermore, the resulting convex algorithms are in many cases slower than alternative approaches.

The alternative to convex relaxation are so-called greedy algorithms. These methods solve simpler optimization problems in an iterative fashion. For sparse models, greedy pursuits are a fast family of methods that build up the sparse representation one (or a few) elements at a time. However, once chosen, a non-zero element is not normally un-selected so that the algorithm is unable to correct errors made in previous iterations. Furthermore, greedy pursuits are restricted to a small set of non-convex constraints. An alternative set of greedy methods are projection-based methods. We have highlighted above the way in which we can compute the best approximation of an element with an element of our constraint set. There are now a range of algorithms that iteratively project elements back onto the non-convex constraint while also trying to reduce the error between the observed measurements and the model. For general non-convex constraints and linear models, the compressed sensing matching pursuit (CoSaMP) [38,43] and the iterative hard thresholding (IHT) algorithms [37,44] are two common choices.

The work reported here is based on the IHT algorithm [44], its extension to nonlinear problems suggested in [39] and the use of the conjugate gradient acceleration suggested in [45].

(a) Iterative hard thresholding

The IHT algorithm was developed for the inversion of an underdetermined linear system under sparsity constraints. In particular, for given observations $\mathbf{y}$ and matrix Φ , IHT tries to solve the non-convex optimization problem

$$\min_{\mathbf{x}} \|\mathbf{y} - \Phi \mathbf{x}\|_2 \text{ such that } \|\mathbf{x}\|_0 \leq K. \quad (4.1)$$

This is done with the simple iteration

$$\mathbf{x}^{x+1} = P(\mathbf{x}^n + \mu \Phi^T(\mathbf{y} - \Phi \mathbf{x}^n)), \quad (4.2)$$

where $P(\cdot)$ calculates the best K sparse approximation of a vector. The step size μ has to be chosen carefully to prevent instability. Under certain conditions on Φ and μ , IHT is guaranteed to find near optimal solutions to the above optimization problem [44].

The extension of the IHT algorithm to nonlinear problems starts with the observation that $\Phi^T(\mathbf{y} - \Phi \mathbf{x}^n)$ in the above expression is proportional to the negative gradient of the unconstrained cost function $\|\mathbf{y} - \Phi \mathbf{x}\|_2$. Thus, if we were trying to optimize a more general nonlinear cost function

$$\min_{\mathbf{x}} f(\mathbf{x}) \text{ such that } \|\mathbf{x}\|_0 \leq K, \quad (4.3)$$

then we simply replace the appropriate gradient into the algorithm and use the iteration

$$\mathbf{x}^{x+1} = P(\mathbf{x}^n - \mu \nabla f(\mathbf{x}^n)), \quad (4.4)$$

again, for a suitable choice of μ and under certain conditions on $f(\mathbf{x})$, this algorithm is guaranteed to solve the non-convex optimization problem with a certain precision [39].

While the previous discussion of the IHT algorithm concentrated on sparse models, other non-convex constraints can be enforced. This is achieved through a suitable replacement of the projection operator $P(\cdot)$. For example, if we want to estimate a matrix with rank less than K , then $P(\mathbf{X})$ is the operator that finds the best rank K approximation to $\mathbf{X}$. As discussed above, this can be computed using the singular value decomposition.

(b) Conjugate gradient-based iterative hard thresholding for nonlinear observation models

Several approaches have been suggested to increase convergence of IHT type algorithms [46–48]. For the nonlinear X-ray model, we use the algorithm proposed in [39] for nonlinear compressed sensing problems; however, to improve convergence, we implement an acceleration similar to that proposed by Blanchard *et al.* [45], who developed a conjugate gradient-based variation of the linear IHT algorithm [44].

Our algorithm replaces the gradient descent step with a more general directional descent step. The update

$$\mathbf{x}^{x+1} = P(\mathbf{x}^n - \mu \nabla f(\mathbf{x}^n)) \quad (4.5)$$

is replaced by the more general update

$$\mathbf{x}^{x+1} = P(\mathbf{x}^n + \mu \mathbf{d}^n), \quad (4.6)$$

where $\mathbf{d}^n$ is a descend direction. The computation of the descend direction is inspired by the conjugate gradient methods used for nonlinear optimization. In the first iteration or during a conjugate gradient restart iteration, we use $\mathbf{d}^n = -\nabla f(\mathbf{x}^n)$. In other iterations, $\mathbf{d}$ is a combination of the current gradient and the previous descend direction

$$\mathbf{d}^n = -0.5 \nabla f(\mathbf{x}^n) + \beta \mathbf{d}^{n-1}, \quad (4.7)$$

where β is computed as $\beta = \frac{\|P_{\mathcal{U}}(\nabla f(\mathbf{x}^{[n]})\|_2^2}{\|(P_{\mathcal{U}}(\nabla f(\mathbf{x}^{[n-1]}))\|_2^2}$. We here use the projection operator $P_{\mathcal{U}}$ which projects the gradient at the current and the previous estimate onto the subspace $\mathcal{U}$, which is the subspace in

the union of subspace model in which the current estimate $\mathbf{x}^n$ lies. This is an important difference from the more standard conjugate gradient method. In effect, we assume that our optimization problem is restricted to the subspace $\mathcal{U}$. Whenever we change subspaces (that is, whenever the projection $P(\cdot)$ projects onto a different subspace), then our optimization problem changes so that we restart the conjugate gradient method by using $\beta = 0$ in the next iteration.

In this form, the algorithm can be used for non-convex constraint sets that are unions of linear subspaces. In this setting, the algorithm relies on four operations: the calculation of a cost function $f(\hat{\mathbf{x}})$, the evaluation of its gradient $g(\hat{\mathbf{x}})$, the nonlinear projection operator $P(\cdot)$ and the linear projection operator $P_{\mathcal{U}}$. For our matrix-valued problem with $\hat{\mathbf{x}} = \hat{\mathbf{X}}$, we here use

$$f(\hat{\mathbf{X}}) = \|\mathbf{y} + \log([\mathbf{e}^{-\mathbf{A}\hat{\mathbf{X}}}]\mathbf{I})\|^2 \quad (4.8)$$

and

$$g(\hat{\mathbf{X}}) = \nabla f(\hat{\mathbf{X}}). \quad (4.9)$$

To better understand the projection operators $P(\cdot)$ and $P_{\mathcal{U}}$, assume we have a vector $\mathbf{x} = [-1 \ 0.2 \ 4 \ -2 \ 0.9]^T$ and we use a 2-sparse non-convex constraint set. The projection of $\mathbf{x}$ onto this set is $P(\mathbf{x}) = [0 \ 0 \ 4 \ -2 \ 0]^T$. We can then also specify the subspace $\mathcal{U}$ as all those vectors that have non-zero entries in the third and fourth position, that is, all vectors that can be written as

$$\begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix}. \quad (4.10)$$

The projection $P_{\mathcal{U}}$ that projects onto this subspace simply sets the first, second and fifth entry in a matrix to zero. Alternatively, for matrix factorization models, we use a similar approach. Assume $\mathbf{X} \approx \mathbf{D}\mathbf{S}$, where $\mathbf{D}\mathbf{S}$ is the best approximation of $\mathbf{X}$ as a low-rank (or a non-negative low-rank) decomposition. We then use the matrix $\mathbf{D}$ to define the subspace onto which we can project any $\mathbf{X}$ using $\mathbf{D}(\mathbf{D}^T\mathbf{D})^{-1}\mathbf{D}^T\mathbf{X}$.

This leaves us with the problem of estimating a good step size μ . We here perform a line search in direction $\mathbf{d}$. This is done using a quadratic approximation to the line search cost function

$$\min_{\mu} f(\mathbf{x}^n + \mu P_{\mathcal{U}}(\mathbf{d})), \quad (4.11)$$

where we again only optimize within the subspace $\mathcal{U}$. The approximation of $f(\mathbf{x}^n + \mu P_{\mathcal{U}}(\mathbf{d}))$ fits a quadratic function to match the magnitude and gradient of $f(\mathbf{x}^n)$ evaluated at the current estimate $\mathbf{x}^n$ as well as the value and gradient of $f(\mathbf{x}^n + a P_{\mathcal{U}}(\mathbf{d}))$ for a given test step size a . Thus, the step size μ used is that which achieves the minimum of the quadratic function. In addition to the conjugate gradient restart due to the algorithm moving from one subspace to another, standard nonlinear conjugate gradient restart conditions are also monitored (such as those based on a gradient orthogonality condition or on maximum iteration counts). Note that for matrix factorization-based constraints, the subspace changes continuously between iterations so that a subspace-based restart would mean that we simply use gradient descend. In this case, it makes sense to only use gradient orthogonality conditions or maximum iteration counts to define the conjugate-gradient restarts.

The algorithm can thus be summarized as follows:

- Input: $\mathbf{X}^{[0]}$, P , f , and g
- $\mathbf{p}^{[0]} = 0$
- if $\|\mathbf{X}^{[0]}\| = 0$
 - $\mathcal{U} = \text{subspace}(P(g(\mathbf{X}^{[0]})))$
- else
 - $\mathcal{U} = \text{subspace}(P(\mathbf{X}^{[0]}))$
- $a = \|P_{\mathcal{U}}(g(\mathbf{X}^{[0]}))\|^{-1}$

- : Iterate ($n = 1, n++$) until some stopping criterion is met:
 - if $n = 1$ or if the conjugate gradient restart condition is satisfied
 - $\beta = 0$
 - else
 - $\beta = \frac{\|P_{\mathcal{U}}(g(\mathbf{x}^{[n]}))\|^2}{\|(P_{\mathcal{U}}(g(\mathbf{x}^{[n-1]})))\|^2}$
 - $\mathbf{d} = -0.5 g(\mathbf{X}^{[n]}) + \beta * \mathbf{p}^{[n-1]}$
 - $a = \text{linesearch}(f, g, \mathbf{d}, \mathcal{U}, a);$
 - $\mathbf{X}^{[n+1]} = P(\mathbf{X}^{[n]} + a\mathbf{d})$
 - $\mathcal{U} = \text{subspace}(\mathbf{X}^{[n+1]})$
 - $\mathbf{p}^{[n]} = \mathbf{d}.$

(c) Extension to mixed projections

The above algorithm assumes that P projects onto one of several linear subspaces. While many of the constraints we have discussed above fall into this category, other constraints or combinations of constraints will lead to more general projections. For example, we know that the X-ray attenuation coefficients are positive. It thus would make sense to include positivity constraints into our non-convex constraints. This raises two issues: firstly, a model with a positivity constraint is no longer a subspace model and secondly, projecting onto the union of a non-convex subspace model and the positive orthant is far from trivial in most cases. We here use what we call mixed projections in which we first project onto the subspace model and then project onto the second constraint. The second constraint does not have to be convex (as the positivity constraint) but can be a more general constraint. For example, we might want to constrain the reconstruction matrix to be sparse in the wavelet domain as well as have a low-rank structure. We do this by first projecting onto the sparse wavelet model using a simple wavelet thresholding approach. This is then followed by a low-rank approximation.

We found that inclusion of this additional projection generally improves performance. In the above algorithm, we insert this projection step into the algorithm as the second to last line, that is, after the projection onto the non-convex constraint and the estimation of the subspace. While this works well with a projection onto the positive orthant as well as the non-negative matrix decomposition, for the more general low-rank projection, we found that we were no longer optimizing within a subspace. In this setting, we thus used the conjugate gradient method to optimize in the full space, that is, we compute β as $(\|g(\mathbf{x}^{[n]})\|^2)/(\|g(\mathbf{x}^{[n-1]})\|^2)$, that is, using the unprojected direction vectors. Note however that we still conduct the line search within the subspace defined by the subspace projection.

5. Numerical results

(a) Artificial data

Our simulations are based on two phantoms. The disc phantom in figure 1 and the Shepp–Logan phantom. The disc phantom simulates a phantom with four materials, an aluminium ring with inner silicon disc, four silver inclusions in the aluminium ring and two gold inclusions in the silicon disc. For the Shepp–Logan phantom, we partitioned the original Shepp–Logan phantom into different areas, one for each of the different grey levels. Each of these areas was then assumed to be either a uniform mixture of three materials (aluminium, iron and silver), or was assumed to be one of these materials. For the first case, mixture coefficients are generated randomly, while in the second case, materials were assigned to regions randomly.

In all simulations, mass attenuation coefficients are taken from the NIST database. We model an X-ray source operating at 120 keV using a tungsten target. The detector is assumed to be caesium iodide with a linear response. Energy levels were quantized in bins with centre energies of 30, 35, 40, 45, 50, 52.5, 55, 57.5, 60, 62.5, 65, 67.5, 70, 77.5, 85, 92.5, 100, 112.5 and 115 keV.

Figure 4. Comparison of linear reconstruction methods on a challenging phantom from 40 projections. The original is an object of aluminium, silicon, gold and silver. FBP gives a reconstruction with severe artefacts and of the iterative solvers, the TV regularized method performs best.

Non-linear X-ray projections were generated on a different spatial grid from the grid used to generate the system matrix $\mathbf{A}$ for the reconstruction algorithm. For the reconstruction, we used an energy quantization of 10 levels with centre energies of 30, 40, 50, 55, 60, 65, 70, 85, 100 and 115 keV.

The disc phantom was measured using 40 cone-beam projections from equally spaced angles, covering 360° . Each projection was measured with a linear X-ray detector array with 181 sensor elements. Both the original phantom and the reconstruction were on a 128×128 spatial grid, with the difference that the simulated measurements used a grid rotated by 45° . Measurements from the Shepp–Logan based phantom were generated similarly with the difference that we only used 32 projections. For the reconstruction, we assumed knowledge of I_0 but did not use knowledge of the mass attenuation coefficients for the three materials. Observations are then calculated as

$$\mathbf{y} = -\log([e^{-\mathbf{A}\mathbf{D}\mathbf{S}}]I_0). \quad (5.1)$$

When simulating the Shepp–Logan phantom in which each homogeneous region consists of a mixture of the three materials, then $\mathbf{D} \in \mathbb{R}^{N \times 3}$ is a matrix with repeated randomly generated rows (with entries drawn from a uniform distribution and normalized to sum to 1), with a different row for each grey level in the Shepp–Logan phantom. To simulate the Shepp–Logan phantom in which each homogeneous region consists of a single material, the matrix $\mathbf{D}$ has a single non-zero entry per row, which is set to one. For each homogeneous region, the location of the non-zero entry is generated randomly. In both cases, $\mathbf{S} \in \mathbb{R}^{3 \times 19}$ is a matrix with the quantized mass attenuation coefficients for the three materials at the 19 different energy levels. I_0 is a vector that counts the average number of X-ray photons in a given energy band transmitted by the source and converted

Figure 5. Comparison of nonlinear reconstruction methods on a challenging phantom from 40 projections. The original is an object of aluminium, silicon, gold and silver. The nonlinear reconstruction performs better under certain constraints than the linear methods. The best-performing non-convex constraint is the combination of wavelet sparsity with a non-negative matrix factorization. When using wavelet-sparsity here, we use a row sparse matrix model as described in §3b(ii).

to visible light in the scintillator. We added Gaussian noise to the observations, normalized to give a signal to noise ratio (SNR) of 40 dB.

In order to be able to compare linear reconstruction methods which estimate a single X-ray attenuation image and nonlinear methods which estimate attenuation images for each energy level, we here compare the performance in terms of the algorithms' ability to estimate the average attenuation calculated as

$$-\log([e^{-X}]I). \quad (5.2)$$

The average attenuation can be compared directly to the estimates from the linear estimates, while for the nonlinear estimates, we also average the estimated attenuation coefficients over energy levels before comparison.

Results for the disc phantom are shown in figures 4 and 5 for the linear reconstruction and the nonlinear reconstructions, respectively. We here compared the FBP algorithm available through the Matlab image processing toolbox, the ART implemented in the AIRTools toolbox [1] (using the randomized Kaczmarz version of the algorithm) with our conjugate gradient algorithm that uses the linear model with two different constraints, wavelet sparsity and wavelet tree sparsity. We also used a total variation constrained inversion using the TVAL3 code (<http://www.caam.rice.edu/~optimization/L1/TVAL3/>). The algebraic method performs better than FBPs. Including additional constraints further improves performance, with the TV regularized results giving the highest SNR here, though the beam hardening artefacts are still

Figure 6. Distribution of results for different methods and for 10 different material assignments, where each area in the phantom is assigned a single material. The seven methods shown in red (light grey in print version) (lower case labels) use the nonlinear model in the reconstruction, while the other five methods (upper case labels) use the linear model. From top to bottom, the methods are constraint using (1) non-negative matrix factorization projections, (2) mixed wavelet sparsity and non-negative matrix factorization projections, (3) mixed wavelet sparsity and low-rank matrix factorization projections, (4) wavelet tree sparsity, (5) wavelet sparsity, (6) low-rank matrix projections, (7) positivity, (8) total variation regularization, (9) wavelet tree sparsity and (10) wavelet sparsity. The second to last results are those achieved with ART [1] and the last results are those for FBP. (Online version in colour.)

very strong. Also, none of the methods is able to allow us to distinguish between the aluminium and silicon.

The reconstructions with the nonlinear model are shown in figure 5. We here use our conjugate gradient method with a range of constraints, positivity, low rank, wavelet row-sparsity, wavelet-tree sparsity, a combination of wavelet row-sparsity and low rank, a combination of wavelet row-sparsity and non-negative matrix factorization and non-negative matrix factorization on its own. While some of the results have reduced beam-hardening artefacts, the artefacts are still visible here. However, the use of the matrix models has led to better SNR figures, especially for the model that combines wavelet row-sparsity with a non-negative matrix factorization. For the same constraint, the nonlinear model leads here to worse SNR results compared with the linear model. This is due to the fact that the nonlinear model has significantly more parameters. But once we start to use some of the constraints that are only available for matrix models, such as the non-negative matrix decomposition constraint, we see improvements over the linear model.

The results for the Shepp–Logan phantom are summarized in figures 6 and 7, where we show box plots of the SNR values achieved with the different methods for 10 different random instances of the two problem settings.

Again, while the non-convexly constrained linear approaches do in general work as well as if not better than the approaches based on the nonlinear cost using the same constraints, using more advanced constraints such as the non-negative matrix factorization can preform much better.

We found that we consistently got the best results either with a combination of wavelet sparsity in the spatial domain combined with an additional projection onto a non-negative matrix factorization or with the non-negative matrix factorization constraint alone. Looking at these

Figure 7. Distribution of results for different methods and for 10 different material assignments, where each area in the phantom is assigned a random mixture of all three materials. The seven methods shown in red (light grey in print version) (lower case labels) use the nonlinear model in the reconstruction, while the other five methods (upper case labels) use the linear model. From top to bottom, the methods are constraint using (1) non-negative matrix factorization projections, (2) mixed wavelet sparsity and non-negative matrix factorization projections, (3) mixed wavelet sparsity and low-rank matrix factorization projections, (4) wavelet tree sparsity, (5) wavelet sparsity, (6) low-rank matrix projections, (7) positivity, (8) total variation regularization, (9) wavelet tree sparsity and (10) wavelet sparsity. The second to last results are those achieved with ART [1] and the last results are those for FBP. (Online version in colour.)

results in the light of the fact that we are unable to calculate exact projections onto the union of these two constraints and the fact that our algorithm is primarily designed for subspace models (which the non-negative factorization is not) shows that the two constraints are capturing important complementary properties of the data. This demonstrates that there is a need to study these joint constraints in more detail in the future.

The only linear cost function-based result that is not matched by the nonlinear methods is the linear approach with the TV constraint. This suggests that the TV constraint is much more powerful for our phantoms than the related wavelet constraints we use here. This suggests the use of similar constraints also with our nonlinear model, which can be easily done by adding a TV regularization term to the nonlinear cost function. Combining this with the non-negative matrix factorization constraint is likely to offer additional benefits, though this is still under investigation.

(b) Real data

To demonstrate the applicability of the technique to real data, we applied it to two real datasets, comparing standard FBP to the best-performing approach in the previous experiment.

The first dataset was acquired at 360 kV. We used 32 projections and a line array detector. The scanned object was a model of an engine block made of aluminium. The longest dimension of the object was about 15 cm. Reconstruction of a two-dimensional cross section using FBPs is shown in figure 8. The reconstruction with a nonlinear model and a sparsity and non-negative matrix factorization constraint is shown in figure 9. Due to the extremely small number of projections, severe artefacts are visible in the FBP. Using our nonlinear model and a wavelet sparsity constraint, the streak artefacts have been removed but the wavelet constraint has introduced some smoothing.

Figure 8. FBP reconstruction from 32 projections of a slice through a model engine block made of aluminium.

Figure 9. Reconstruction with the nonlinear model and a sparsity and non-negative matrix factorization constraint. Reconstruction from 32 projections of a slice through a model engine block made of aluminium.

Figure 10. Detail of two reconstructions from 200 projections. (a) FBP and (b) nonlinear conjugate gradient reconstructions of a slice through a high-density workpiece with internal structures. Beam-hardening artefacts are visible in the FBP reconstruction which are not visible in the reconstruction achieved with our approach. (Online version in colour.)

The second set of X-ray projections was acquired with 200 kV. The test object consisted of two concentric tubes. The outer tube was acrylic and the inner tube aluminium. Within the inner tube, aluminium wires of different diameters were attached. We used 200 projections acquired with a flat panel detector, taking only the central slice for reconstruction. The FBP reconstruction is shown figure 10a, whereas our nonlinear reconstruction with a positivity constraint is shown in 10b. The red oval highlights beam-hardening artefacts visible in the FBP, which are not visible in the reconstruction achieved with our approach.

6. Conclusion

In this paper, we have proposed a projected conjugate gradient algorithm that is able to solve nonlinear inverse problems under a range of non-convex constraints. We have applied the method to the inversion of the nonlinear X-ray transmission model to solve X-ray computed tomography reconstruction. This not only allowed us to address the beam-hardening problem but also allowed us to reconstruct images from few projections. Extensive simulation results on synthetic data have shown that advanced non-convex constraints, such as a combination of wavelet sparsity and non-negative matrix factorization, can have significant benefits over standard positivity constraints alone. While these non-convex constraints do not guarantee globally optimal solutions, we could show that when we initialized our method with a good linear reconstruction, then the non-convexly constraint nonlinear reconstructions were often better. However, the improvement also comes at the cost of increased computational complexity. For real data, inverting the nonlinear model often took hundreds of iterations to achieve the best performance.

The fact that TV regularization outperforms even our best nonlinear model-based estimate shows the power to the TV constraint for objects with uniform material density as we simulated here. It is likely that the inclusion of such a TV regularization term in our nonlinear model will offer similar benefits, especially when combined with the non-negative matrix factorization constraint. Using joint constraints has also demonstrated clear benefits. While we enforced these here with consecutive projections, recent studies (not reported here) have shown that an approach that averages the projections onto the individual constraints tends to perform better than the approach used here, though a detailed study of how this works in the tomographic setting is still to be undertaken.

The main drawback of our approach currently is the slow convergence due to the complex nature of the non-convex cost function. This was here addressed with a sensible initialization approach, though additional benefits are likely if better optimization strategies are adopted. One possible improvement might be the use of a better line search strategy. The quadratic approximation does not always seem to give a good step size and other approaches are currently under investigation.

Funding statement. T.B. acknowledges support from EPSRC grant nos. EP/K037102/1, EP/J005444/1, the CCPi (funded through EPSRC grant no. EP/J010456/1) and the TSB (grant no. 101804) as well as a University of Southampton, Faculty of Engineering and the Environment 'New Frontiers Fellowship'.

References

1. Hansen PC, Saxid-Hansen S. 2012 AIR Tools—a MATLAB package of algebraic iterative reconstruction methods. *J. Comput. Appl. Math.* **236**, 2167–2178. (doi:10.1016/j.cam.2011.09.039)
2. Brooks RA, Di Chiro G. 1976 Beam hardening in X-ray reconstructive tomography. *Phys. Med. Biol.* **21**, 390–398. (doi:10.1088/0031-9155/21/3/004)
3. Van de Castele E, Van Dyck D, Sijbers J, Raman E. 2004 A model-based correction method for beam hardening artefacts in X-ray microtomography. *J. X-ray Sci. Technol.* **12**, 53–57.
4. So A, Hsieh J, Li JY, Lee TY. 2009 Beam hardening correction in CT myocardial perfusion measurement. *Phys. Med. Biol.* **54**, 3031–3050. (doi:10.1088/0031-9155/54/10/005)
5. Hammersberg P, Mangard M. 1998 Correction for beam hardening artefacts in computerised tomography. *J. X-ray Sci. Technol.* **8**, 75–93.
6. Herman GT. 1979 Correction for beam hardening in computed tomography. *Phys. Med. Biol.* **24**, 81–106. (doi:10.1088/0031-9155/24/1/008)
7. Herman G, Trivedi S. 1983 A comparative study of two postreconstruction beam hardening correction methods. *IEEE Trans. Med. Imaging* **2**, 123–135. (doi:10.1109/TMI.1983.4307626)
8. Hsieh J, Molthen RC, Dawson CA, Johnson RH. 2000 An iterative approach to the beam hardening correction in cone beam CT. *Med. Phys.* **27**, 23–29. (doi:10.1118/1.598853)
9. Gao H, Xing Y, Li S. 2006 Beam hardening correction for middle-energy industrial computerized tomography. *IEEE Trans. Nucl. Sci.* **53**, 2796–2807. (doi:10.1109/TNS.2006.879825)

10. Vedula V SVM, Munshi P. 2008 An improved algorithm for beam-hardening corrections in experimental X-ray tomography. *NDT E Int.* **41**, 25–31. (doi:10.1016/j.ndteint.2007.06.002)
11. Nalcioğlu O, Lou RY. 1979 Post-reconstruction method for beam hardening in computerised tomography. *Phys. Med. Biol.* **24**, 330–340. (doi:10.1088/0031-9155/24/2/009)
12. Ruegsegger P, Hangartner T, Keller HU, Hinderling T. 1978 Standardization of computed tomography images by means of a material-selective beam hardening correction. *J. Comput. Assist. Tomogr.* **2**, 184–188. (doi:10.1097/00004728-197804000-00012)
13. Elbakri IA, Fessler JA. 2002 Statistical image reconstruction for polyenergetic X-ray computed tomography. *IEEE Trans. Med. Imaging* **21**, 89–99. (doi:10.1109/42.993128)
14. Krumm M, Kasperl S, Franz M. 2008 Reducing non-linear artifacts of multi-material objects in industrial 3D computed tomography. *NDT E Int.* **41**, 242–251. (doi:10.1016/j.ndteint.2007.12.001)
15. Krumm M, Kasperl S, Franz M. 2008 Referenceless beam hardening correction in 3d computed tomography images of multi-material objects. In *World Conf. on Nondestructive Testing, October 2008*. Shanghai, China: Chinese Society for Non-destructive Testing.
16. Meagher JM, Mote Jr CD, Skinner HB. 1990 CT image correction for beam hardening using simulated projection data. *IEEE Trans. Nucl. Sci.* **37**, 1520–1524. (doi:10.1109/23.55865)
17. Yan CH, Whalen RT, Beaupre GS, Yen SY, Napel S. 2000 Reconstruction algorithm for polychromatic CT imaging: application to beam hardening correction. *IEEE Trans. Med. Imaging* **19**, 1–11. (doi:10.1109/42.832955)
18. Chen CY, Chuang KS, Wu J, Lin HR, Li MJ. 2001 Beam hardening correction for computed tomography images using a postreconstruction method and equivalent tissue concept. *J. Digit. Imaging* **14**, 54–61. (doi:10.1007/s10278-001-0003-2)
19. McDavid WD, Waggener RG, Payne WH, Dennis MJ. 1975 Spectral effects on three-dimensional reconstruction from X-rays. *Med. Phys.* **2**, 321–324. (doi:10.1118/1.594200)
20. McDavid WD, Waggener RG, Payne WH, Dennis MJ. 1977 Correction for spectral artifacts in cross-sectional reconstruction from X-rays. *Med. Phys.* **4**, 54–58. (doi:10.1118/1.594302)
21. Van Gompel G, Van Slambrouck K, Defrise M, Batenburg KJ, de Mey J, Sijbers J, Nuyts J. 2011 Iterative correction of beam hardening artifacts in CT. *Med Phys.* **38**, S36. (doi:10.1118/1.3577758)
22. Kachelriess M, Kalender WA. 2005 Improving PET/CT attenuation correction with iterative CT beam hardening correction. In *Proc. IEEE Nuclear Science Symp. Conf. Record, 23–29 October 2005*. (doi:10.1109/NSSMIC.2005.1596704)
23. Alvarez RE, Macovski A. 1976 Energy-selective reconstructions in X-ray computerised tomography. *Phys. Med. Biol.* **21**, 733–744. (doi:10.1088/0031-9155/21/5/002)
24. Stonestrom JP, Alvarez RE, Macovski A. 1981 A framework for spectral artifact correction in X-ray CT. *IEEE Trans. Biomed. Eng.* **28**, 128–141. (doi:10.1109/TBME.1981.324786)
25. De Man B, Nuyts J, Dupont P, Marchal G, Suetens P. 2001 An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Trans. Med. Imaging* **20**, 999–1007. (doi:10.1109/42.959297)
26. Menvielle N, Goussard Y, Orban D, Soulez G. 2005 Reduction of beam-hardening artefacts in X-ray CT. In *Proc. IEEE Engineering in Medicine and Biology Annual Conf., Shanghai, China*, pp. 1865–1868. (doi:10.1109/IEMBS.2005.1616814)
27. Elbakri IA, Fessler JA. 2003 Segmentation-free statistical image reconstruction for polyenergetic X-ray computed tomography with experimental validation. *Phys. Med. Biol.* **48**, 2453–2477. (doi:10.1088/0031-9155/48/15/314)
28. Olsen EA, Han KS, Pisano DJ. 1981 CT reprojection polychromaticity correction for three attenuators. *IEEE Trans. Nucl. Sci.* **28**, 3628–3640. (doi:10.1109/TNS.1981.4331811)
29. Van de Castele E, Van Dyck D, Sijbers J, Raman E. 2002 An energy-based beam hardening model in tomography. *Phys. Med. Biol.* **47**, 4181–4190. (doi:10.1088/0031-9155/47/23/305)
30. Elbakri IA, Fessler JA. 2001 Statistical X-ray-computed tomography image reconstruction with beam-hardening correction. *Proc. SPIE* **4322**, 1–12. (doi:10.1117/12.430961)
31. Ruehrnschopf EP, Klingenberg K. 2011 A general framework and review of scatter correction methods in X-ray cone-beam computerized tomography. Part 1: scatter compensation approaches. *Med. Phys.* **38**, 4296–4311. (doi:10.1118/1.3599033)
32. Tikhonov AN, Arsenin VY. 1977 *Solution of ill-posed problems*. Washington, DC: Winston and Sons.
33. Rudin L, Osher S, Fatemi E. 1992 Nonlinear total variation based noise removal algorithms. *Phys. D* **60**, 259–268. (doi:10.1016/0167-2789(92)90242-F)

34. Hämmäläinen K, Harhanen L, Hauptmann A, Kallonen A, Niemi E, Siltanen S. 2014 Total variation regularization for large-scale X-ray tomography. *Int. J. Tomogr. Simul.* **25**, 1–25.
35. Candès E, Romberg J, Tao T. 2006 Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory* **52**, 489–509. (doi:10.1109/TIT.2005.862083)
36. Donoho D. 2006 Compressed sensing. *IEEE Trans. Inform. Theory* **52**, 1289–1306. (doi:10.1109/TIT.2006.871582)
37. Blumensath T. 2011 Sampling and reconstructing signals from a union of linear subspaces. *IEEE Trans. Inform. Theory* **57**, 4660–4671. (doi:10.1109/TIT.2011.2146550)
38. Baraniuk R, Cevher V, Duarte M, Hegde C. 2010 Model-based compressive sensing. *IEEE Trans. Inform. Theory* **56**, 1982–2001. (doi:10.1109/TIT.2010.2040894)
39. Blumensath T. 2013 Compressed sensing with nonlinear observations and related nonlinear optimization problems. *IEEE Trans. Inform. Theory* **59**, 3466–3474. (doi:10.1109/TIT.2013.2245716)
40. Baraniuk RG. 1999 Optimal tree approximation with wavelets. *Proc. SPIE* **3813**, 196–207. (doi:10.1117/12.366780)
41. Candès E, Recht B. 2009 Exact matrix completion via convex optimization. *Found. Comput. Math.* **9**, 717–772. (doi:10.1007/s10208-009-9045-5)
42. Paatero P, Tapper U. 1994 Positive matrix factorization: a non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics* **5**, 111–126. (doi:10.1002/env.3170050203)
43. Needell D, Tropp JA. 2009 CoSaMP: iterative signal recovery from incomplete and inaccurate samples. *Appl. Comput. Harmon. Anal.* **26**, 301–321. (doi:10.1016/j.acha.2008.07.002)
44. Blumensath T, Davies ME. 2009 Iterative hard thresholding for compressed sensing. *Appl. Comput. Harmon. Anal.* **27**, 265–274. (doi:10.1016/j.acha.2009.04.002)
45. Blanchard J, Tanner J, Wei K. 2014 CGIHT: conjugate gradient iterative hard thresholding for compressed sensing and matrix completion. Technical report, University of Oxford. See <http://eprints.maths.ox.ac.uk/1833>.
46. Blumensath T. 2011 Accelerated iterative hard thresholding. *Signal Process.* **92**, 752–756. (doi:10.1016/j.sigpro.2011.09.017)
47. Foucart S. 2011 Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM J. Numer. Anal.* **49**, 2543–2563. (doi:10.1137/100806278)
48. Cevher V. 2011 An ALPS view of sparse recovery. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5808–5811. (doi:10.1109/ICASSP.2011.5947681)