Regularization by Denoising Sub-sampled Newton Method for Spectral CT Multi-Material Decomposition

Alessandro Perelli and Martin S. Andersen

Technical University of Denmark (DTU Compute)

Abstract

Spectral Computed Tomography (CT) is an emerging technology that enables to estimate the concentration of basis materials within a scanned object by exploiting different photon energy spectra. In this work, we aim at efficiently solving a model-based maximum-a-posterior problem to reconstruct multi-materials images with application to spectral CT. In particular, we propose to solve a regularized optimization problem based on a plug-in image-denoising function using a randomized second order method. By approximating the Newton step using a sketching of the Hessian of the likelihood function, it is possible to reduce the complexity while retaining the complex prior structure given by the data-driven regularizer. We exploit a non-uniform block sub-sampling of the Hessian with inexact but efficient Conjugate gradient updates that require only Jacobian-vector products for denoising term. Finally, we show numerical and experimental results for spectral CT materials decomposition.

1 Introduction

Spectral Computed Tomography (CT) is a developing technique for diagnostic and therapy in medical imaging which can gather enhanced physical parameters information compared to conventional CT by exploiting the dependence of the attenuation coefficients of an object with distinct photon energy spectra. This information leads to the estimation of images of basis materials [1], which can be used for quantitative imaging [2], reducing beam hardening artifacts and improving the contrast-to-noise ratio [3]. Recently, spectral CT with photon counting detectors has attracted increasing interest.

In contrast to conventional transmission CT, photon counting detectors can resolve the X-ray photon energy and allow to acquire registered energy selective images [4] with effective reduction of the electronic readout noise, which is of crucial importance for low-dose applications. However, material decomposition algorithms typically lead to a degradation of the signal-to-noise ratio (SNR) and noise amplification [5] compared to the unprocessed spectral images. This is still a fundamental problem of spectral CT since it limits the usability of material selective images.

Most spectral CT algorithms separate the process of material decomposition and image reconstruction. Material decomposition is either performed in the projection space prior to image reconstruction [6, 7] or an image based decomposition algorithm is applied after reconstructing the images corresponding to different photon energy spectra [8, 9, 10, 11] by using tensor factorization [12] or dictionary learning [13]. Separating these steps is sub-optimal because the full statistical information contained in the spectral tomographic measurements cannot be exploited. Statistical iterative reconstruction (SIR) techniques provide a data consistent approach to obtaining material

selective images using a statistical model of the projection measurements and incorporating prior knowledge about the reconstructed material images. Unfortunately, the loss function induced by the forward model which directly maps the expected spectral projection measurements and the material images is non-convex and difficult to minimize; convex approximations of the loss function such as convex-concave techniques [14, 15, 16] or semi-empirical models [17, 18] have been proposed but suffer of high computational cost because of the double-loop nature of the solver.

The main challenges which are attracting considerable interest are related to the fact that the image to be recovered lies in an high-dimensional space and to leverage more sophisticated or learned data-dependent priors which are needed since the estimation problem is generally ill-posed. Regarding the issue of reducing the per-iteration computational cost to solve the reconstruction in the high-dimensional setting, numerous first order solvers which leverage random numerical linear algebra have been deployed in different machine learning and imaging applications. Well established algorithms like Stochastic Gradient Descent (SGD) [19, 20], order subsets [21] and Alternating Direction Method of Multiplier (ADMM) have gained considerably attention in spectral CT [22, 23]. Unfortunately, the iteration complexity of first order methods relies significantly on the condition number of the problem like SGD which has a sub-linear convergence rate to only a neighborhood of the solution [24]. Instead, the information related to the curvature is decisive to improve the convergence by reducing the condition number and to obtain physical meaningful and accurate quantitative estimation. Therefore, accounting for a greedy selection of parameters and step-size is crucial to design first-order algorithms for imaging like CT image reconstruction [21, 25]. Deterministic second order methods for solving regularized optimization problems with Generalized Linear Models (GLM) have been widely studied [26, 27] but despite the superior, linear or super-linear, convergence rate of Newton methods compared to first order methods one weakness relies on the high computational cost of calculating the Hessian matrix in high dimensional imaging application [26, 27]. The Newton method requires at each iteration to solving exactly a linear system with complexity $\mathcal{O}(mn^2)$, with m the dimension of the measurement vector and n the dimension of the input vector, and therefore it is practical only when the dimension n is low or if the Hessian can be factorized in time linear to the input dimension, like by the fast Fourier Transform (FFT) [28]. Instead inexact methods solve approximately the linear system by iterative algorithms like Conjugate Gradient (CG) which does not require to calculate the full Hessian matrix but only operate through vectors products [29]; this is particularly appealing in imaging systems like CT where there are fast implementations of the operator that represents the physical forward model.

In this work, a new second-order algorithm for statistical iterative material decomposition (Denoising-IHS) is developed which enjoys an improved accuracy and computation trade-off. We propose a computational efficient quantitative estimation with second order and parameters-free methods based on randomized linear algebra and regularization induced by denoising. We exploit the recent development of randomized sketches [30] and sub-sampling methods [31] for the Hessian matrix to improve upon the computational cost-per-iteration. Unfortunately, these methods are efficient in the regime where the measurement dimension m is much larger than the input dimension n, i.e. $m \gg n$ [32], since the theoretical bound on the embedding of the sketched Hessian matrix requires a number of samples of order $\mathcal{O}(n)$. Furthermore, these methods have been mainly analyzed for minimizing smooth and convex cost function and the interesting case of using denoiser-based regularizers has not been deeply studied. Therefore, we aim at developing an algorithm for spectral CT model where the dimension of the measurements is of the same order as the dimension of the image to be reconstructed, i.e. $m \approx n$ and exploit denoising function to capture the complex structure of the material images to be reconstructed.

1.1 Related works

Important contributions related to sketching methods for solving optimization problems refer to the Newton-Sketch algorithm [33] where the Newton equations are solved using an unbiased estimator of the true Hessian obtained by performing a dimensionality reduction using random sketching matrices. Different types of sketching such as Sub-Gaussian matrices, Hadamard or Fourier transforms, sparse Count Sketch matrices can be employed and have been analyzed [33]. Furthermore, sub-sampling Newton methods [34, 35, 36, 37] are based on random row sub-sampling according to a uniform or non-uniform probability distribution [36, 38, 31]. It is interesting to note that random matrix sub-sampling schemes can be considered a particular type of sketching and they are more efficient especially in imaging problems where the linear forward model has a precise, deterministic structure, like the Random transform for CT imaging. Furthermore, since the Radon matrix $\mathbf{A}$ is an high dimensional tall and sparse matrix, the count sketch might be efficient because of $\mathcal{O}(\text{nnz}(\mathbf{A}))$ multiplication complexity, where $\text{nnz}(\mathbf{A})$ stands for the number of non-zero elements of the matrix $\mathbf{A}$. Unfortunately, the count sketch requires a bound on the sketch size of order $\mathcal{O}(n)$, with n the dimension of the vectorized image input, which is practically useless in CT scenarios where n is of the same order of the number of measurements m.

Our main interest is to develop a non-uniform sub-sampling Newton algorithm for CT imaging by incorporating data-dependent regularizer which is becoming state-of-the art in model-based iterative reconstruction. In imaging applications where the input prior is difficult to model, powerful regularization techniques are based on data-driven models or by exploiting denoising functions, like the Plug-and-play (PnP) approach [39] or the regularization by denoising (RED) [40, 41, 42]. It is interesting how both approaches allow to integrate in the iterative solver either deterministic denoisers, like BM3D [43], NLM, TV or deep learning CNN-based denoisers [44]. The RED approach is based on an explicit regularization term [40] and scalable first-order solvers for the RED problem in high dimensions have been proposed [45]. As far as the authors are aware, the RED approach has been investigated with first order solvers but not with Newton methods mainly because by default it would require to explicitly generate the Jacobian matrix of the denoiser which is infeasible in high-dimensions.

1.2 Main Contribution and Relation to Existing Works

In this work, we develop a second order method (Denoising-IHS) which combine the iterative sketching [33] for the log-likelihood function and the RED regularization which can be implemented by a generic denoiser through the score matching formulation with application to spectral CT material decomposition. We propose an effective choice for both the sketching scheme and the solver in order to reduce the computational complexity for imaging problems characterized by physical-based convolutional operators.

The overall algorithmic time complexity is $T = n_{outer} \cdot (t_{samp} + t_{solve})$ where t_{samp} is the time needed at each iteration to construct the sketching, t_{solve} is the time to solve the linear sub-problem inexactly, and n_{outer} is the total number of iterations which is affected by the convergence rate. Considering t_{samp}, taking inspiration from [37], we develop a random non-uniform sub-sampling of the measurement where the probabilities are calculated according to a penalized block-leverage score metric. While penalizing the ridge leverage scores allows to reduce the sketch size with provable convergence, this metric exploits the CT measurements block structure which is defined as a product of the number of views and the number of detectors. This means that we randomly sub-sample over the acquired angles and the ridge leverage score associated with each

view is calculated over a block constituted by the number of detectors. We take advantage of the convolutional structure of the Radon transform to estimate the leverage scores through the Fourier decomposition and reduce the complexity to $\mathcal{O}(n \log n)$. Regarding the computational time to solve the sub-problem t_{solve}, we employ the CG algorithm since it does not require to store any gradient or Jacobian matrices and all operations can be performed using implicitly operators instead of matrix form. We develop a method to estimate the denoiser's Jacobian product through automatic differentiation for CNN-based denoisers. Finally, for the term n_{outer}, we provide convergence analysis of the proposed method.

1.3 Notation

Matrices or discrete operators and column vectors are written respectively in capital and normal boldface type, i.e. $\mathbf{A}$ and $\mathbf{a}$ to distinguish from scalars and real-valued variables written in normal weight. $(\cdot^T)$ refers to the transpose of a matrix and $\mathbf{1}$ refers to a vector of ones. Non-random quantities and random realizations are not distinguished typographically while random variable are written with capital letters. The conditional probability density function of $\mathbf{y}$ given $\mathbf{x}$ is denoted by $p(\mathbf{y}|\mathbf{x})$. A Gaussian random variable $\mathbf{x}$ with vector mean $\mathbf{a}$ and isotropic variance b is denoted by $\mathbf{x} \sim \mathcal{N}(\mathbf{a}, b\mathbf{I})$. $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{b}^T \mathbf{a}$ refers to the vector inner product.

1.4 Structure of the Paper

The remainder of this paper is structured as follows: in Section 2, the spectral CT mathematical model with the Gaussian approximation of the noise is described and Section 3 analyses the optimization problem together with a review of the Hessian sketching Newton method and the formulation of the regularization by denoising approach. The proposed algorithms is presented in Section 4 which contains the details of the denoising-based regularizer, the spectral CT data loss and an analysis of the CG inner solver. The further sections describe the crucial blocks of the Denoising-IHS algorithm: the approach to estimating the denoiser directional derivatives needed in the CG solver is described in Section 5 while the sketching strategy based on block leverage scores sub-sampling of the Hessian matrix of the data loss function is presented in Section 6. Finally, in Sections 7 and 8 we show the numerical simulation and real acquisition on different sets of images and we assess the algorithm in terms of accuracy and computation.

2 Spectral X-ray CT Model

In this Section, we describe the CT physical measurement process in the continuous domain with the spectrum of the X-ray source beams composed by different energy intervals. The spectral X-ray mathematical model is based on the Beer's law which provides the X-ray intensity after transmission, p_i^k of the i-th ray at the k-th energy interval (bin) as follows

$$p_i^k = \int_E S^k(E) e^{-\int_{l \in L_i} \mu(\vec{\kappa}(l), E) dl} dE + \eta_i^k, \quad \begin{cases} i & = 1, 2, \ldots, N_d \times N_p, \\ k & = 1, 2, \ldots, N_b, \end{cases} \tag{1}$$

where p_i^k is the X-ray intensity of the i-th detector pixel in the k-th detector bin, E is the photon flux density, N_d is the number of detector pixels, N_p is the number of projections, $\mu(\vec{\kappa}(l), E)$ denotes the linear attenuation coefficient that is related to the position function $\vec{\kappa}(l)$ and the energy level E, η_i^k is the Gaussian error term for the i-th element in the k-th energy bin and $S^k(E)$ represents

the photon flux density for the k-th detector bin, which is the number of incident photons at the energy E in the k-th energy window.

In the discrete domain, we map the continuous spatially and energy dependent distribution $\mu(\vec{\kappa}(l), E) \in L_2(\mathbb{R})$ by using a basis material decomposition for representing the object

$$\mu(\vec{\kappa}(l), E) \approx \sum_{m=1}^{N_m} c_m(E) \sum_{j=1}^{N_v} \phi_j(\vec{\kappa}) x_{j,m} \tag{2}$$

where $c_m(E)$ indicates the linear attenuation coefficient of the m-th material at the energy level E, $\phi_j(\vec{\kappa})$ defines the j-th among N_v basis functions associated with a discrete sampling on a $\sqrt{N_v} \times \sqrt{N_v}$ Cartesian grid, N_m represents the number of basis materials and $x_{j,m}$ is the weight fraction of the m-th material in the j-th image voxel/pixel. According to the basis material parameterization in Eq. (2), the line integral in Eq. (1) becomes a summation:

$$\int_{l \in L_i} \mu(\vec{\kappa}(l), E) dl \approx \sum_{m=1}^{N_m} c_m(E) \sum_{j=1}^{N_v} x_{j,m} \int_{l \in L_i} \phi_j(\vec{\kappa}(l)) dl = \sum_{m=1}^{N_m} \sum_{j=1}^{N_v} r_{i,j} x_{j,m} c_{m,e} \tag{3}$$

where $r_{i,j}$ represents the i, j element of the system matrix describing the discrete summation along the i-path from source through object at pixel position j onto each detector. Repeating this over all lines defines the full view linear tomographic system matrix $\mathbf{R} = [r_{ij}]$; by physical design, $\mathbf{R}$ is a sparse and non-negative matrix with dimensions $N_d N_p \times N_v$. For a discrete set of energies of dimension N_e, the discrete model (1) can be written as

$$p_i^k = \sum_{e=1}^{N_e} s^k e^{\sum_{m=1}^{N_m} \sum_{j=1}^{N_v} r_{i,j} x_{j,m} c_{m,e}} + \eta_i^k \tag{4}$$

The source spectrum is modeled as a vector $\mathbf{s} \in \mathbb{R}^{N_e \times 1}$ with non-negative entries and the detector response is a positive matrix $\mathbf{D} \in \mathbb{R}^{N_b \times N_e}$ whose dimension is the product of the number of bins N_b and the number of discrete energies N_e. Therefore, the spectrum for all the energy bins $b = 1, \ldots, N_b$ is the matrix obtained as Hadamard point-wise product $\mathbf{S} = \mathbf{D} \odot \mathbf{s} \in \mathbb{R}^{N_b \times N_e}$. Both the source spectrum and the detectors response matrix are assumed to be known.

In matrix notation, $\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_l, \ldots \mathbf{x}_{N_m}] \in \mathbb{R}^{N_v \times N_m}$ denotes the matrix obtained as columns concatenation of the vectorized image $\mathbf{x}_l$ of the $l = 1, \ldots, N_m$ basis materials. Therefore, $r_{i,j}$, and $c_{m,e}$ are collected in a matrix form and p_i^k, s_e^k, η_i^k are concatenated with respect to their energy windows. The corresponding matrix equation of (4) can be represented as

$$\mathbf{P} = \exp\left(-\mathbf{R}\mathbf{X}\mathbf{C}^T\right) \mathbf{S} + \mathbf{N} \tag{5}$$

with $\mathbf{C}$ a matrix of the size $N_e \times N_m$ where each entry corresponds to $c_{e,m}$, the linear attenuation coefficient of the energy e, and the m-th material. We vectorize the nonlinear matrix equation (5) on both sides and linearize to avoid a tensor form for computing second order derivatives and we follow the mathematical model derived in [46]. In the forward problem, we use the full spectrum, and the matrix $\mathbf{S}$ is rectangular while for solving the inverse problem, we choose the average in each energy window to represent the corresponding energy spectrum, so $N_b = N_e$ and the matrix $\mathbf{S}$ in is a nonsingular diagonal matrix. Therefore, we multiply $\mathbf{S}^{-1}$ and vectorize on both sides of (5) and using the properties of Kronecker products we obtain

$$\left(\mathbf{S}^{-T} \otimes \mathbf{I}\right) \mathbf{p} = \exp\left(-\left(\mathbf{C} \otimes \mathbf{R}\right) \mathbf{x}\right) + \left(\mathbf{S}^{-T} \otimes \mathbf{I}\right) \mathbf{n} \tag{6}$$

where $\mathbf{p} = \mathrm{vec}\,(\mathbf{P})$, $\mathbf{x} = \mathrm{vec}\,(\mathbf{X})$ and $\mathbf{n} = \mathrm{vec}\,(\mathbf{N})$. Denoting $\widetilde{\mathbf{p}} = (\mathbf{S}^{-T} \otimes \mathbf{I})\,\mathbf{p}$ and $\widetilde{\mathbf{n}} = (\mathbf{S}^{-T} \otimes \mathbf{I})\,\mathbf{n}$ and taking the logarithm, we have $\log\,(\widetilde{\mathbf{p}} - \widetilde{\mathbf{n}}) = -\,(\mathbf{C} \otimes \mathbf{R})\,\mathbf{x}$.

Using a first order Taylor expansion at $\widetilde{\mathbf{p}}$, $\log\,(\widetilde{\mathbf{p}} - \widetilde{\mathbf{n}}) = \log\,(\widetilde{\mathbf{p}}) - \mathrm{diag}\,(\widetilde{\mathbf{p}})^{-1}\,\widetilde{\mathbf{n}} + \mathcal{O}\,(\|\widetilde{\mathbf{n}}\|^2)$, we obtain the following approximation $\mathbf{y} \approx (\mathbf{C} \otimes \mathbf{R})\,\mathbf{x} - \mathrm{diag}\,(\widetilde{\mathbf{p}})^{-1}\,\widetilde{\mathbf{n}}$ where $\mathbf{y} = -\log\,(\widetilde{\mathbf{p}})$. Assuming Gaussian noise $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathrm{diag}(\mathbf{p}))$, we have

$$\mathbf{y}|\mathbf{x} \sim \mathcal{N}\big((\mathbf{C} \otimes \mathbf{R})\,\mathbf{x}, \boldsymbol{\Sigma}\big) \tag{7}$$

where the inverse covariance matrix $\boldsymbol{\Sigma}^{-1}$ is expressed as

$$\boldsymbol{\Sigma}^{-1} = \mathrm{diag}\,(\widetilde{\mathbf{p}})\,\Big(\mathbf{S} \otimes \mathbf{I}\Big)\mathrm{diag}\,(\mathbf{p})^{-1}\,(\mathbf{S}^T \otimes \mathbf{I})\,\mathrm{diag}\,(\widetilde{\mathbf{p}}) \tag{8}$$

$\mathbf{P}$ is a matrix which represents the number of photons of each energy window in the corresponding column and each element of $\mathbf{P}$ is a positive integer. From expression (8), we can see that the $\boldsymbol{\Sigma}^{-1}$ is diagonal since we have assumed that $\mathbf{S}$ is diagonal.

We consider the MAP estimation problem for the spectral X-ray CT model with Gaussian approximation. Given the measurement vector $\mathbf{y} \in \mathbb{R}^m$ of dimension $m = N_d \cdot N_p \cdot N_b$, we consider the data-fit loss function $f(\mathbf{x}, \mathbf{y}) : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ as a measure of the distance between the input parameter vector $\mathbf{x} \in \mathbb{R}^n$, of dimension $n = N_v \cdot N_m$, and the measurement vector $\mathbf{y}$, with $m > n$. The loss function is given by the following empirical negative log-likelihood (NLL)

$$f(\mathbf{x}, \mathbf{y}) \;=\; -\sum_{i=1}^{m} \log p(y_i|\mathbf{x}) \;\approx\; \frac{1}{2}\sum_{i=1}^{m} \sigma_i^{-1}\big(\mathbf{a}_i\mathbf{x} - y_i\big)^2 \;=\; \frac{1}{2}\|\mathbf{A}\mathbf{x} - \mathbf{y}\|_{\boldsymbol{\Sigma}^{-1}}^2 \tag{9}$$

where $\mathbf{A} = \mathbf{C} \otimes \mathbf{R}$. Given the likelihood function in Eq. (9) with $\mathbf{x} \in \mathbb{R}_+^n$, it results that

$$\nabla_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) \;=\; \mathbf{A}^T \boldsymbol{\Sigma}^{-1}\,(\mathbf{A}\mathbf{x} - \mathbf{y})\,, \qquad \nabla_{\mathbf{x}}^2 f(\mathbf{x}, \mathbf{y}) \;=\; \mathbf{A}^T \boldsymbol{\Sigma}^{-1} \mathbf{A} \tag{10}$$

It is important to note that in the region of physical interest, i.e. $a_i^T \mathbf{x} \in \mathbb{R}_+$, $\nabla_{\mathbf{x}}^2 f(\mathbf{x}, \mathbf{y})$ is bounded.

3 Optimization problem

The MAP estimator can be formulated as the following optimization problem

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x} \in \mathbb{R}_+^n} f(\mathbf{x}, \mathbf{y}) + \rho_{\boldsymbol{\sigma}, \nu}(\mathbf{x}) \tag{11}$$

where $f(\mathbf{x}, \mathbf{y}) = -\log p(\mathbf{y}|\mathbf{x})$ represents the likelihood loss function as described in Eq. (9) and $\rho_{\boldsymbol{\sigma}, \nu}(\mathbf{x})$ is the regularization term which defines a prior $\rho_{\boldsymbol{\sigma}, \nu}(\mathbf{x}) = -\log p(\mathbf{x}, \nu)$ on the vector $\mathbf{x}$ to be estimated, with ν the noise level. In the following, we will consider the constrained minimization over the set of non-negative vectors $\mathbf{x} \in \mathbb{R}_+^n$. We focus on a particular non-convex regularizer which can be explicitly defined according to a denoising (RED) function dependent on a set of parameters $\boldsymbol{\sigma}$. The regularization term and the gradient can be expressed as

$$\rho_{\boldsymbol{\sigma}, \nu}(\mathbf{x}) \;=\; \frac{1}{2\nu}\mathbf{x}^T\big(\mathbf{x} - D_{\boldsymbol{\sigma}}(\mathbf{x})\big), \qquad \nabla\rho_{\boldsymbol{\sigma}, \nu}(\mathbf{x}) \;=\; \frac{1}{\nu}\big(\mathbf{x} - D_{\boldsymbol{\sigma}}(\mathbf{x})\big) \tag{12}$$

if the denoiser $D_{\boldsymbol{\sigma}}(\cdot)$ is locally homogeneous, i.e. $(1 + \epsilon)D_{\boldsymbol{\sigma}}(\mathbf{x}) = D_{\boldsymbol{\sigma}}((1 + \epsilon)\mathbf{x})$, $\forall \mathbf{x}, \epsilon > 0$ and the its Jacobian is symmetric, $\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x})] = \mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x})]^T$. When the denoiser does not follows the above conditions, the relationship between $\rho_{\boldsymbol{\sigma}}(\mathbf{x})$ and the gradient expression in Eq. (12) does not hold. In the following Section 3.1, an interpretation of RED for generic denoisers is summarized which is based on the score-matching by denoising framework [41].

3.1 Denoising Score Matching

In medical imaging applications, the true prior $p(\mathbf{x})$ is not available and therefore it is not possible to compute directly the MMSE estimator $D_{MMSE,\nu}(\cdot)$. Since calculating $D_{MMSE,\nu}$ is infeasible in high dimensions, the main idea is to approximate the optimal MMSE estimator by a generic denoiser, $D_{\boldsymbol{\sigma}}(\cdot)$ parametrized over $\boldsymbol{\sigma}$, even for denoisers not locally homogeneous or with non-symmetric Jacobian. By utilizing $p(\mathbf{x}, \nu)$ as prior of $\mathbf{x}$ for the regularization term $\rho(\mathbf{x}) = -\log p(\mathbf{x}, \nu)$ in (11), it is necessary to evaluate the expression for $\nabla_{\mathbf{x}}\rho(\mathbf{x}) = \nabla_{\mathbf{x}} \log p(\mathbf{x}, \nu)$ and $\nabla_{\mathbf{x}}^2 \rho(\mathbf{x})$. Instead, using an efficient denoiser on the class $\mathcal{D} \triangleq \{D_{\boldsymbol{\sigma}} : \boldsymbol{\sigma} \in \boldsymbol{\Sigma}\}$, although an analytic expression for $p(\mathbf{x}, \nu)$ is either not available or complex to handle, from an algorithmic perspective the prior is not needed but only the gradient and its Jacobian. Therefore, we exploit the approach based on approximating the score instead of the prior directly by invoking the connection between denoising and the score, i.e. the gradient of the log-prior, and its Hessian [47]. It was shown in [48] that if we choose a function such that

$$\theta_{\boldsymbol{\sigma}}(\mathbf{x}) \quad = \quad \frac{D_{\boldsymbol{\sigma}}(\mathbf{x}) - \mathbf{x}}{\nu}, \qquad \frac{\partial \theta_{\boldsymbol{\sigma}}(\mathbf{x})}{\partial \mathbf{x}} \quad = \quad \frac{1}{\nu}\Big(\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x})] - \mathbf{I}\Big) \tag{13}$$

where $\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x})]$ is the Jacobian matrix of the denoising function of dimension $n \times n$, then for any $D_{\boldsymbol{\sigma}}$ and $\mathbf{x}$, the score-matching error is connected to the denoiser approximation error as follows

$$\left\|D_{\boldsymbol{\sigma}}(\mathbf{x}) - D_{MMSE,\nu}(\mathbf{x})\right\|^2 = \left\|\theta_{\boldsymbol{\sigma}}(\mathbf{x}) - \nabla_{\mathbf{x}} \log p(\mathbf{x}, \nu)\right\|^2 \quad = \quad \left\|\frac{\partial \theta_{\boldsymbol{\sigma}}(\mathbf{x})}{\partial \mathbf{x}} - \nabla_{\mathbf{x}}^2 \log p(\mathbf{x}, \nu)\right\|^2 \tag{14}$$

It is important to stress that the equivalence between Eq. (13) and (14) holds also in the case where the parameter vector of the denoiser $\boldsymbol{\sigma}$ is not optimal. The Jacobian matrix $\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x})]$ of dimension $n \times n$ is computationally prohibitive to calculate and store at each iteration t of the solver in high dimensions. In Section 5 we propose a way to approximate the Jacobian-vector products required within the proposed iterative sub-sampled Newton solver.

3.2 Review of the Newton Method

In this section we build the analysis of the sub-sampled Newton method by first summarizing the Newton method. Let $g : \mathbb{R}^n \to \mathbb{R}$ be a closed, convex, and twice-differentiable function that is bounded below. Given the current iterate $\mathbf{x}^t \in \mathcal{C}$, the Newton method generates a new iterate $\mathbf{x}^{t+1}$ by performing a constrained minimization of the second-order Taylor expansion

$$\mathbf{x}^{t+1} = \arg \min_{\mathbf{x} \in \mathcal{C}} \left[\frac{1}{2}\langle \mathbf{x} - \mathbf{x}^t, \nabla_{\mathbf{x}^t}^2 g(\mathbf{x}^t, \mathbf{y})(\mathbf{x} - \mathbf{x}^t)\rangle + \langle \nabla_{\mathbf{x}^t} g(\mathbf{x}^t, \mathbf{y}), \mathbf{x} - \mathbf{x}^t\rangle\right] \tag{15}$$

where for the traditional Newton's method $\nabla_{\mathbf{x}^t}^2 g(\mathbf{x}^t, \mathbf{y})$ represents the full $n \times n$ Hessian matrix. In the following derivations, we simplify the notation of the gradient and the Hessian with $\nabla g(\mathbf{x}^t)$ and $\nabla^2 g(\mathbf{x}^t)$ respectively, hiding the dependence over $\mathbf{y}$. By setting the gradient respect to $\mathbf{x}^t$ equal to zero, it follows the fixed-point update for the exact full Newton's method

$$\nabla^2 g(\mathbf{x}^t)(\mathbf{x} - \mathbf{x}^t) \quad = \quad -\nabla g(\mathbf{x}^t) \quad \Rightarrow \quad \mathbf{x}^{t+1} \quad = \quad \mathbf{x}^t - \alpha^t \left[\nabla^2 g(\mathbf{x}^t)\right]^{-1} \nabla g(\mathbf{x}^t) \tag{16}$$

The two forms of updates presented in (16) highlight that 2 different types of solvers can be employed; in the first case, we solve a linear system in order to estimate the vector $\mathbf{p}^t = \mathbf{x} - \mathbf{x}^t$

and this is generally obtained by using the CG algorithm. Alternatively, it is possible to directly invert the square $n \times n$ Hessian matrix. The step-size α^t in (16) can be assumed constant, $\alpha^t = 1, \ \forall t \in \mathbb{N}$.

In the following Section, we describe the proposed Denoising-IHS algorithm which uses a stochastic approximation of the Hessian $\nabla^2 g(\mathbf{x}^t)$ in order to compute $\mathbf{p}^t$ with reduced computational cost, instead of calculating an $n \times n$ Hessian matrix and solving the linear system in (16) exactly which would be remarkably expensive in terms of time. We assume that the gradient $\nabla g(\mathbf{x}^t)$ is computed exactly and the Hessian is approximated by random projection to lower dimension $s < n$ by applying a sketching matrix $\mathbf{S}^t \in \mathbb{R}^{s \times n}$ such that $\frac{1}{s}\mathbb{E}\big[(\mathbf{S}^t)^T \mathbf{S}^t\big] = \mathbf{I}_n$.

4   Proposed Denoising-based Sub-sampled Newton Method

The Denoising-IHS algorithm minimizes the original MAP-driven regularized composite cost function of the form $g(\mathbf{x}) = f(\mathbf{x}) + \rho_{\boldsymbol{\sigma}}(\mathbf{x})$ in Eq. (11) and we utilize a partially sketched Newton update which means that a sketch of the Hessian $\nabla^2_{\mathbf{x}} f(\mathbf{x}, \mathbf{y})$ is performed while retaining the exact form of the Hessian associated with the regularizer, i.e. $\nabla^2 \rho_{\boldsymbol{\sigma}}(\mathbf{x})$. The reason behind this choice is while it is possible to reduce the computational complexity related to the loss function by using fast non-uniform rows sampling, for generic black-box denoising function the computational cost of applying the sketching to the regularizer is higher than using the original function. The algorithm assumes that the square root of the Hessian matrix of dimension $m \times n$, indicated as $[\nabla^2 f(\mathbf{x})]^{\frac{1}{2}}$, can be computed for the CT model in (10), it follows that $[\nabla^2 f(\mathbf{x}^t)]^{\frac{1}{2}} = \boldsymbol{\Sigma}^{-\frac{1}{2}} \mathbf{A}$
being $\boldsymbol{\Sigma}$ a diagonal matrix. The sketched Hessian of the loss function $f(\mathbf{x})$ is constructed as

$$\mathbf{H}^t_f(\mathbf{x}^t) \ = \ \left(\mathbf{S}^t \left[\nabla^2 f(\mathbf{x}^t)\right]^{\frac{1}{2}}\right)^T \mathbf{S}^t \left[\nabla^2 f(\mathbf{x}^t)\right]^{\frac{1}{2}} \ = \ \mathbf{G}(\mathbf{x}^t)^T \mathbf{G}(\mathbf{x}^t), \qquad \mathbf{S}^t \in \mathbb{R}^{s \times n} \qquad (17)$$

with $\mathbf{G}(\mathbf{x}^t) = \mathbf{S}^t \left[\nabla^2 f(\mathbf{x}^t)\right]^{\frac{1}{2}}$, while the full Hessian of the denoising-based regularizer is denoted by $\mathbf{H}^t_\rho(\mathbf{x}^t) = \nabla^2 \rho_{\boldsymbol{\sigma}}(\mathbf{x}^t)$ with $\mathbf{H}^t_\rho(\mathbf{x}^t) \succeq 0$ a positive semi-definite matrix. Eq. (17) represents an unbiased estimator for the full Hessian of the loss

$$\mathbb{E}_t \left[\left(\left[\nabla^2 f(\mathbf{x}^t)\right]^{\frac{1}{2}}\right)^T (\mathbf{S}^t)^T \mathbf{S}^t \left[\nabla^2 f(\mathbf{x}^t)\right]^{\frac{1}{2}}\right] = \nabla^2 f(\mathbf{x}^t) \qquad (18)$$

where the expectation is taken over the iteration index t conditioned over the random sketches selected at previous iterations $[1, \ldots, t-1]$. The splitting leads to the partially sketched update

$$\mathbf{x}^{t+1} \ = \ \arg\min_{\mathbf{x} \in \mathcal{C}} \left[\frac{1}{2}(\mathbf{x} - \mathbf{x}^t)^T \mathbf{H}^t(\mathbf{x}^t)(\mathbf{x} - \mathbf{x}^t) + \ \langle \nabla g(\mathbf{x}^t), \mathbf{x} - \mathbf{x}^t \rangle\right] \qquad (19)$$

$$\mathbf{H}^t(\mathbf{x}^t) \ = \ \mathbf{H}^t_f(\mathbf{x}^t) + \mathbf{H}^t_\rho(\mathbf{x}^t) = \left(\mathbf{S}^t[\nabla^2 f(\mathbf{x}^t)]^{\frac{1}{2}}\right)^T \mathbf{S}^t[\nabla^2 f(\mathbf{x}^t)]^{\frac{1}{2}} + \nabla^2 \rho_{\boldsymbol{\sigma}}(\mathbf{x}^t)$$

By substituting into (19) the definition for the loss in (9) and the denoising-based regularizer in (13), the expressions for the gradient and the Hessian become

$$\nabla g(\mathbf{x}^t) \ = \ \nabla f(\mathbf{x}^t, \mathbf{y}) + \nabla \rho_{\nu,\boldsymbol{\sigma}}(\mathbf{x}^t) \ = \ \mathbf{A}^T \boldsymbol{\Sigma}^{-1}\left(\mathbf{A}\mathbf{x}^t - \mathbf{y}\right) + \frac{1}{\nu}\left(D_{\boldsymbol{\sigma}}(\mathbf{x}^t) - \mathbf{x}^t\right)$$

$$\mathbf{G}(\mathbf{x}^t) \ = \ \mathbf{S}^t[\nabla^2 f(\mathbf{x}^t)]^{\frac{1}{2}} \ = \ \mathbf{S}^t \left[\boldsymbol{\Sigma}^{-\frac{1}{2}}\mathbf{A}\right]$$

$$\mathbf{H}^t(\mathbf{x}^t) \ = \ \mathbf{G}(\mathbf{x}^t)^T \mathbf{G}(\mathbf{x}^t) + \frac{1}{\nu}\left(\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)] - \mathbf{I}\right) \qquad (20)$$

The update is $\mathbf{x}^{t+1} = \mathbf{x}^t + \mathbf{p}^t$ where $\mathbf{p}^t$ is obtained as solution of the system of equations

$$\mathbf{H}^t(\mathbf{x}^t)\mathbf{p}^t = -\nabla g(\mathbf{x}^t) \tag{21}$$

We consider to solve the system (21) approximately using the Conjugate gradient method. We focus on the analysis of the accuracy of the solution and the complexity. The inner update is

$$\mathbf{x}^{t+1} = \mathbf{x}^t + \mathbf{p}_r^t \tag{22}$$

where $\mathbf{p}_r^t$ is the approximate solution after r inner iterations of the CG solver. We consider the case where the gradient is fully computed which is a realistic scenario in case of separable loss function, and the Hessian is non-uniformly sampled. To perform an analysis of the complexity of the Sketch Newton-CG method, we remind that for a symmetric positive definite linear system the convergence of CG is not uniform but depends on the difference between the minimum and maximum eigenvalues. In detail, the convergence rate of the CG method varies at every iteration depending on the spectrum $\{\lambda_1 \leq \lambda_2 \leq, \ldots, \leq \lambda_n\}$ of the positive definite matrix $\mathbf{H}^t(\mathbf{x}^t)$. After r steps of the CG method applied to the linear system in (21), the iterate p_r^t satisfies

$$\|\mathbf{p}_r^t - \mathbf{p}_*^t\|_{\mathbf{C}}^2 \leq \left(\frac{\lambda_{n-r} - \lambda_1}{\lambda_{n-r} + \lambda_1} \right)^2 \|\mathbf{p}_0^t - \mathbf{p}_*^t\|_{\mathbf{C}}^2, \quad \mathbf{C} = \mathbf{H}^t(\mathbf{x}^t)$$

where $\mathbf{p}_*^t$ indicates the exact solution and $\|\mathbf{x}\|_{\mathbf{C}}^2 \triangleq \mathbf{x}^T \mathbf{C} \mathbf{x}$. The linear CG algorithm allows to solve the system (21) using only matrix/operator-vector products with $\mathbf{H}^t(\mathbf{x}^t)$. therefore, we use CG as an Hessian-free inner solver which does not require to store the full $n \times n$ Hessian matrix. It is possible to obtain sketched Hessian-vector products by performing two matrix-vector products with respect to the sketched square-root Hessian. In particular, we compute for each inner iteration $i \in [1, \ldots, r]$ of the CG solver the following variable vector

$$\mathbf{v}_i = \mathbf{H}^t(\mathbf{x}^t)\mathbf{p}_i = \left(\mathbf{G}(\mathbf{x}^t)^T \mathbf{G}(\mathbf{x}^t) + \frac{1}{\nu} \left(\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)] - \mathbf{I} \right) \right) \mathbf{p}_i \tag{23}$$

using an auxiliary vector $\mathbf{z}$ and performing two steps as follows

$$\mathbf{z}_i = \mathbf{G}(\mathbf{x}^t)\mathbf{p}_i, \quad \mathbf{v}_i = \mathbf{G}(\mathbf{x}^t)^T \mathbf{z} + \frac{1}{\nu} \mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]\mathbf{p}_i - \frac{1}{\nu}\mathbf{p}_i \tag{24}$$

Furthermore, by using the CG solver the matrix $\mathbf{H}^t(\mathbf{x}^t)$ is fixed through the inner iterations (being t the index of the outer iteration as in Algorithm 1 and i the index of the CG solver). In the following, we will not explicitly indicate the index i associated with the vector $\mathbf{p}$ for notation simplicity. Since $\mathbf{H}^t(\mathbf{x}^t)$ depends by the Jacobian of the denoiser, as shown in (20), it would be unfeasible to use gradient based solvers which requires to calculate the Jacobian at each inner iteration. Instead, with CG only directional derivatives of the denoiser, i.e. $\mathbf{J}$, are needed to be computed and not the full Jacobian; in Section 5 we present fast methods to estimate accurately the directional derivative of learned CNN-based denoisers.

Although in the worst case, CG will require n iterations to converge (thus requiring the evaluation of n products $\mathbf{H}^t(\mathbf{x}^t)\mathbf{p}^t$), the behavior CG can achieve significant progress in the minimization of the objective function after a reduced number of iterations r. The following Theorem outlines that the proposed Denoising-IHS algorithm 1 can achieve quadratic convergence, with constants C_1, C_2, C_3 dependent on the minimum and maximum eigenvalues of the full Hessian $\bar{\mathbf{H}}(\mathbf{x}^t)$) and the condition number κ of the problem as defined in the proof in Appendix A

Theorem 4.1. *Let $\mathbf{x}^t$ be the iterates generated by the non-uniform sampling Newton-CG algorithm with $r < n$ CG iterations are performed. Then,*

$$\mathbb{E}_t\|\mathbf{x}^{t+1} - \mathbf{x}_*\| \leq C_1\|\mathbf{x}^t - \mathbf{x}_*\|^2 + (C_2 + C_3)\|\mathbf{x}^t - \mathbf{x}_*\|^2 \tag{25}$$

Algorithm 1: Sub-sampled Newton method with Regularization by Denoising

Initialization: set $t = 0$, $\mathbf{r}^0 = \mathbf{0}$, $\mathbf{x}^0 = \mathbf{0}$, $\tau_x^0 = 1$

—- Outer iterations —-

for $t = 1, \ldots, T_{max}$ **do**

    1) Compute the leverage-scores based non-uniform sampling distribution $\{p_i\}_{i=1}^p$ (6.1)

    2) Compute the sketched Hessian of the loss:

        $\mathbf{G}^t(\mathbf{x}^t) = \mathbf{S}^t[\nabla^2 f(\mathbf{x}^t)]^{\frac{1}{2}} = \mathbf{S}^t \left[\mathbf{\Sigma}^{-\frac{1}{2}} \mathbf{A} \right]$

    3) Partial sketched Hessian:

        $\mathbf{H}^t(\mathbf{x}^t) = [\mathbf{G}^t(\mathbf{x}^t)]^T \mathbf{G}^T(\mathbf{x}^t) + \frac{1}{\nu} \left(\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)] - \mathbf{I} \right)$

    —- Inner iterations —-

    4) Solve the liner system using CG:

        $\mathbf{H}^t(\mathbf{x}^t)\mathbf{p}^t = -\nabla g(\mathbf{x}^t)$

end for

return $\mathbf{x}^t$

5 Denoiser directional derivative estimators

The major computational cost of the inexact Newton solver relies on the calculation at each outer iteration t of the Jacobian matrix, of dimensions $n \times n$, of the denoising function $D_{\boldsymbol{\sigma}}$ in Eq. (20) whose computation is unfeasible in high dimensions. On a closed look, CG requires instead to calculate explicitly the direction derivative $\mathbf{J}\left[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)\right] \mathbf{p}^t$ that for the n-dimensional vector $\mathbf{p}^t$,can be computed using finite differences at the cost of a single denoising evaluation as in Eq. (13) through the following expression

$$\mathbf{J}\left[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)\right] \mathbf{p}^t = \lim_{\epsilon \to 0} \frac{D_{\boldsymbol{\sigma}}(\mathbf{x}^t + \epsilon\mathbf{p}^t) - D_{\boldsymbol{\sigma}}(\mathbf{x}^t)}{\epsilon} \tag{26}$$

Although the finite differences approach suffers from numerical instabilities and also requires the computationally expensive evaluation of non-linear functions, there is an efficient procedure for computing exactly the Jacobian-vector product for neural networks.

5.1 Automatic differentiation for CNN denoisers

The procedure to calculate $\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]\mathbf{p}^t$ is based on deriving the right multiplication of Jacobian matrices. We denote the denoiser $D_{\boldsymbol{\sigma}}(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}^m$ and let $\mathbf{p} \in \mathbb{R}^n$ be the right multiplication vector and $\mathbf{u} \in \mathbb{R}^n$ be a dummy variable. The left multiplication of the Jacobian matrix, i.e. $\mathbf{u}^T\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]$, can be calculated by back propagation, which is common in reverse mode automatic differentiation packages

$$\mathbf{u}^T \frac{\partial D}{\partial \mathbf{x}^T} = \mathbf{u}^T\mathbf{J}\left[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)\right] \tag{27}$$

We exploit (27) to compute $\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]\mathbf{p}^t$ by defining $g(\mathbf{u}) = \mathbf{u}^T\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]$. Since $\mathbf{u}^T\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]$ is a vector in $\mathbb{R}^n$, the mapping of $\mathbf{p}^t$ can be defined as the function $g(\mathbf{u}) : \mathbb{R}^n \to \mathbb{R}^n$. We can take derivative of $g(\mathbf{u})$ with respect to $\mathbf{u}$, while providing the left multiplying vector $\mathbf{p}^t$:

$$\mathbf{p}^t \frac{\partial g}{\partial \mathbf{u}} = \mathbf{p}^t \frac{\partial \left(\mathbf{u}^T\mathbf{J}\left[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)\right]\right)}{\partial \mathbf{u}} = \mathbf{p}^t\mathbf{J}\left[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)\right]^T = \left(\mathbf{J}\left[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)\right] \mathbf{p}^t\right)^T \tag{28}$$

6 Block Leverage Scores Sketching

We focus on sketches based on non-uniform random block rows sampling which can be easily and efficiently implemented on GPU-based imaging operators such as CT [49]. Given a probability distribution $\{p_i\}_{i=1}^m$, we randomly sample the rows of the matrix $\mathbf{A}$ a total of s times with replacement from the given probability distribution. The rows of $\mathbf{S}$ are independent and take the values $\mathbf{s}^T = \frac{\mathbf{e}_i}{\sqrt{p_i}}$ with probability p_i for $i = 1, \ldots, m$, where $\mathbf{e}_i \in \mathbb{R}^m$ is the i-th canonical basis vector. The discrete probability distribution is chosen according to the ridge leverage scores.

Definition 6.1. *(Ridge Leverage Scores) Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a scalar $\lambda \in \mathbb{R}_+$, then for $i = 1, \ldots, n$, the i-th leverage scores of A is defined as $l_i^\lambda(\mathbf{A}) = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A} + \lambda \mathbf{I})^\dagger \mathbf{a}_i$.*

We consider to estimate the ridge leverage scores of the square root matrix of the Hessian matrix of the loss function, i.e. $\mathbf{B} = [\nabla^2 f(\mathbf{x}^t)]^{\frac{1}{2}} = \mathbf{\Sigma}^{-\frac{1}{2}} \mathbf{A}$ as defined in (10). The idea is to adaptively penalize the leverage scores of the Hessian-related matrix $\mathbf{B}$ according to the noise level estimated through the trace of the Jacobian of the denoiser based regularizer $\mathbf{H}_\rho^d$. Defining $\mathbf{H}_\rho = \frac{1}{\nu}\Big(\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)] - \mathbf{I}\Big)$, we estimate the trace

$$H_\rho^d = \text{Trace}[\mathbf{H}_\rho] = \frac{1}{\nu}\Big(\text{Trace}\Big[\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]\Big] - 1\Big) \tag{29}$$

We exploit the property that $H_\rho^d = \text{Trace}[\mathbf{H}_\rho] = \mathbb{E}[\mathbf{n}^T \mathbf{H}_\rho \mathbf{n}]$ if $\mathbf{n} \sim \mathcal{N}(0, \mathbf{I})$ is independent of $\mathbf{J}$ (where the expectation is taken over $\mathbf{n}$). Therefore, given a noise realization normally distributed $\mathbf{n}$, we propose to compute the vector $\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]\mathbf{n}$ instead of the matrix $\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]$. The expectation over $\mathbf{n}$ is calculated using a Monte-Carlo method, i.e. generate K i.i.d. $\mathcal{N}(0, \mathbf{I})$ samples vectors, estimate the divergence for each vector and obtain the global divergence by averaging:

$$\text{Trace}\Big[\mathbf{J}[D_{\boldsymbol{\sigma}}(\mathbf{x}^t)]\Big] = \frac{1}{m} \sum_{j=1}^K \mathbf{n}_j^T \left(\frac{D_{\boldsymbol{\sigma}}(\mathbf{x}^t + \epsilon \mathbf{n}_j) - D_{\boldsymbol{\sigma}}(\mathbf{x}^t)}{\epsilon}\right) \tag{30}$$

Given the vectorized image lying in a high dimensional space, it has been empirically observed [50] that we can accurately approximate the expected value using only a single random sample, i.e. $K = 1$. In all the simulations we have used the Monte Carlo method with $K = 1$.

Definition 6.2. *(Block partial leverage scores) Given $\mathbf{B} \in \mathbb{R}^{N_d N_p \times N_v}$ and the scalar $H_\rho^d \in \mathbb{R}_+$ defined in (29), let $\{l_i\}_{i=1}^{N_d N_p}$ the leverage scores of the matrix $\mathbf{B}$. The block partial leverage score for the i-th block is defined as*

$$l_i^{H_\rho^d}(\mathbf{B}) = \sum_{j=(i-1)N_d+1}^{iN_d} l_j, \quad i = 1, \ldots, N_p \tag{31}$$

and the non-uniform sampling distribution is $p_i = \dfrac{l_i^{H_\rho^d}(\mathbf{B})}{\sum_{j=1}^{N_p} l_i^{H_\rho^d}(\mathbf{B})}, \quad i = 1, \ldots, N_p$

Theorem 6.1 ([37]). *Given $\mathbf{B} \in \mathbb{R}^{N_d N_p \times N_v}$, $H_\rho^d \in \mathbb{R}_+$ and $\epsilon \in (0, 1)$, let $\{l_i^{H_\rho^d}(\mathbf{B})\}_{i=1}^{N_p}$ be the block partial leverage scores and $p_i \dfrac{l_i^{H_\rho^d}(\mathbf{B})}{\sum_{j=1}^{N_p} l_i^{H_\rho^d}(\mathbf{B})}$ Let construct the sketch $\mathbf{G} = \mathbf{SB}$, as in (17), by sampling*

the i-th block of $\mathbf{G}$ with probability p_i and rescaling it by $\frac{1}{\sqrt{p_i}}$. If $s \geq 4 \left(\sum_{i=1}^{N_p} l_i^{H_\rho^d}(\mathbf{B}) \right) \log \frac{4N_v}{\delta\epsilon^2}$ with probability at least $1 - \delta$, then the following condition holds

$$\left\| \left(\mathbf{G}^T \mathbf{S}^T \mathbf{S} \mathbf{G} + \mathbf{H}_\rho^d \right) - \left(\mathbf{G}^T \mathbf{G} + \mathbf{H}_\rho^d \right) \right\|^2 \leq \epsilon \left\| \mathbf{G}^T \mathbf{G} + \mathbf{H}_\rho^d \right\|^2 \tag{32}$$

Corollary 6.1. *Let focus on the case of one-row sampling, $d = 1$ and $\mathbf{H}_\rho$ a scaled identity matrix, $\mathbf{H}_\rho = m\gamma\mathbf{I}$. Let define $\mathbf{G} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_{i=1}^m \sigma_i \mathbf{u}_i \mathbf{v}_i^T$ as the singular value decomposition. Then, the leverage score is*

$$l_i^\gamma = \mathbf{g}_i (\mathbf{G}^T \mathbf{G} + m\gamma\mathbf{I})^\dagger \mathbf{g}_i = \sum_{j=1}^n \frac{\sigma_j^2}{\sigma_j^2 + m\gamma} u_{ij}^2, \quad i = 1, \ldots, m \tag{33}$$

and therefore, the effective n-dimension, n_{eff}, is much smaller than n

$$n_{eff}(\mathbf{G}) = \sum_{i=1}^m l_i^\gamma = \sum_{j=1}^n \frac{\sigma_j^2}{\sigma_j^2 + m\gamma} \ll n \tag{34}$$

if $m\gamma$ is larger than the majority of singular values of $\mathbf{G}^T\mathbf{G}$.

It is worth noting we have used the notion of SVD to show the property of the dimensionality reduction but using the SVD is not practical for calculating the leverage scores. Instead, we will show how to use fast matrix decomposition through the FFT.

6.1 Convolutional operators

In principle, computing the leverage scores for a generic matrix $\mathbf{A}$ has a time complexity $\mathcal{O}(mn^2)$ equivalent to doing an SVD which is computationally demanding in high dimensions. We focus our attention on convolutional operator which can be used to describe the structure of the physical CT operator. We show that the convolutional structure allows to compute efficiently the leverage scores which determines the non-uniform random sub-sampling of the Hessian.

Definition 6.3. *(CT Convolutional operator) The continuous 2D X-ray Transform is a convolutional linear operator $\mathcal{R} : L_2(\mathbb{R}^2) \to L_2([0, \pi) \times \mathbb{R})$ which computes the line integral of a function in the 2D input space. The Fourier central slice Theorem states that $\mathcal{R} = \mathcal{F}_\gamma^{-1}\Omega_{\omega^{-1}}\mathcal{F}$ where $\mathcal{F}$ is the 2D Fourier transform (FT), Ω is the coordinate transform operator from Cartesian to polar coordinates, $\omega^{-1} = (\gamma \cos\delta, \gamma \sin\delta)$ and $\mathcal{F}_\gamma^{-1}$ is the inverse 1D FT with respect to γ.*

The output of the linear operator is the sinogram that is a function of δ and the polar space variable ρ. Both $\mathcal{R}$ and $\mathcal{R}^T\mathcal{R}$ are normal-convolutional operators since in the frequency domain

$$\mathcal{R}^T\mathcal{R} = \mathcal{F}^H (\Omega_{\omega^{-1}})^T (\mathcal{F}_\gamma^{-1})^H \mathcal{F}_\gamma^{-1} \Omega_{\omega^{-1}} \mathcal{F} \overset{(a)}{=} \mathcal{F}^H |\det J_\omega| \mathcal{F} \overset{(b)}{=} \mathcal{F}^H \mathcal{D} \left(\frac{1}{|\rho|} \right) \mathcal{F} \tag{35}$$

where (a) follows from $\mathcal{F}_\gamma$ being an unitary operator and (b) derives from the back-projection CT filter formulation where J_ω defines the Jacobian of ω and $\mathcal{D} \left(\frac{1}{|\rho|} \right)$ is the diagonal polar Fourier space operator. From the continuous to the discrete domain, the Fourier-based Radon transform can be written as [51, 52]:

$$\mathbf{R} = \mathbf{F}_\gamma^{-1} \mathbf{\Omega}_{\boldsymbol{\omega}^{-1}} \mathbf{F} \tag{36}$$

where $\mathbf{F}$ is the 2D unitary discrete Fourier transform operator which takes as input the image $\boldsymbol{\mu}$ of dimensions $I \times I$, with $I = \sqrt{N_v}$. The operator $\boldsymbol{\Omega}_{\boldsymbol{\omega}^{-1}}$ performs a discretized version of the continuous coordinate transform in (6.3) which outputs a matrix of polar coordinate samples that are equally-spaced along ρ at the discrete locations $\{i\Delta_\rho\}$ for $i = -\frac{I}{2}, \ldots, \frac{I}{2} - 1$. The non-uniform FFT operator $\boldsymbol{\Omega}_{\boldsymbol{\omega}^{-1}}\mathbf{F}$ takes as input the $I \times I$ input image matrix and output a matrix of dimensions $N_d \times N_p$; $\mathbf{F}_\gamma$ applies the 1D unitary discrete Fourier transform (DFT) matrix separately to each of the radial lines vectors of dimension N_{dec} and it is defined as the Kronecker product between a 1D DFT matrix $\mathbf{F}_1$ and the identity matrix, i.e. $\mathbf{F}_\gamma = \mathbf{F}_1 \otimes \mathbf{I}_{N_p}$. Therefore, the final output is a vector of dimensions $J = N_d \cdot N_p$ where N_d is the number of detectors and N_p is the number of angles (or number of projections) in agreement with (3).

7 Numerical simulations

We have preliminary tested the Denoising-IHS algorithm using the following parameters and geometry; the X-ray spectrum is obtained by using the Spektr [53] with tube potential $E = 80$ keV and initial average photon flux of $I_0 = 2 \cdot 10^3$. Figure 1 shows the spectrum in the energy interval $[0, 150]$ keV of the five discrete selected energy bins. In Figure 1(a), the mass attenuation coefficients of the materials to be estimated are shown; it is interesting to note that the two distinct K-edges associated with the iodine and gadolinium appear at energies corresponding to two distinct energy bins as shown in Figure 1(b).
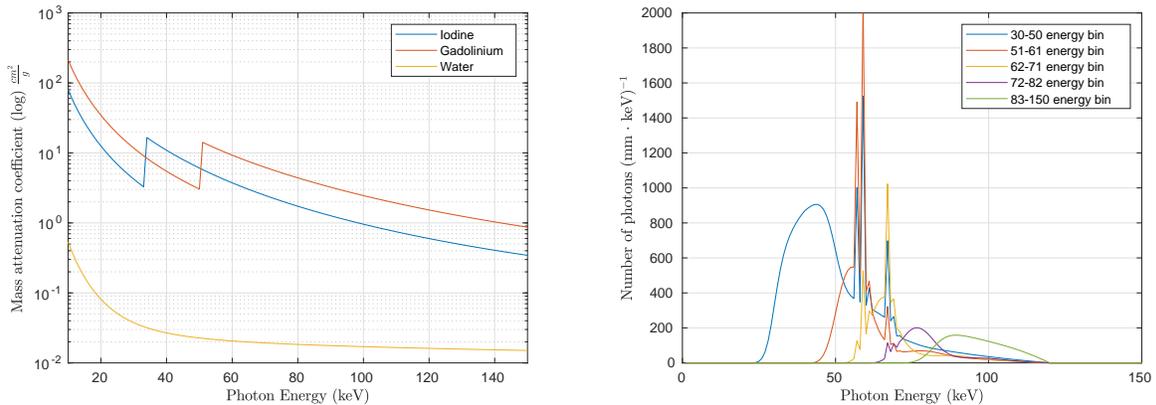
Figure 1: (a) Mass attenuation coefficients of iodine, gadolinium and water as a function of energy and (b) energy spectra of x-ray photons interacting in detector for different energy bins.

To generate the multichannel projections, we use the ASTRA-toolbox [49] to implement the forward operator associated with a fan-beam CT geometry; the distance from the source to the detector is 45 cm, the source to object distance is 31 cm and the detector cell size is 50 μm. The pixel resolution of the reconstructed image is 256×256 and the data is simulated on a double grid of 512×512 to avoid the inverse crime, the number of projection angles is set to 360 and the Poisson distributed noise has been added to the projection data. Figure 2(a) shows the plot of the probabilities of sampling each projection according to the leverage score metric. This plot is obtained by calculating exactly the leverage score of the ASTRA operator offline using the power method. Many projections have a very low score and this shows that a random one-by-one sampling would lose an important quantity of information or it would require to run the algorithm on a number of samples of the order of the projection's dimension.
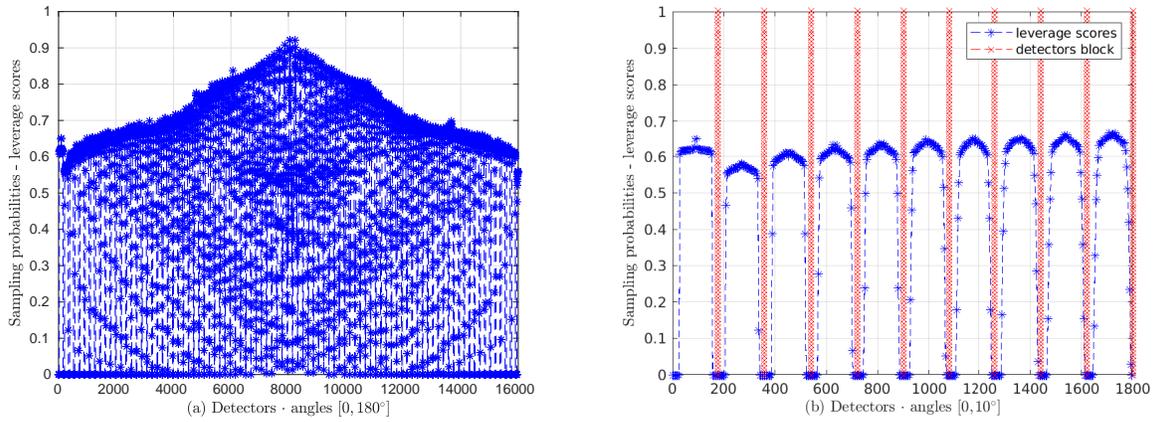
Figure 2: (a) Leverage scores probabilities on the projection up to 180° angles and (b) zoom on the structure of the detectors block leverage scores.
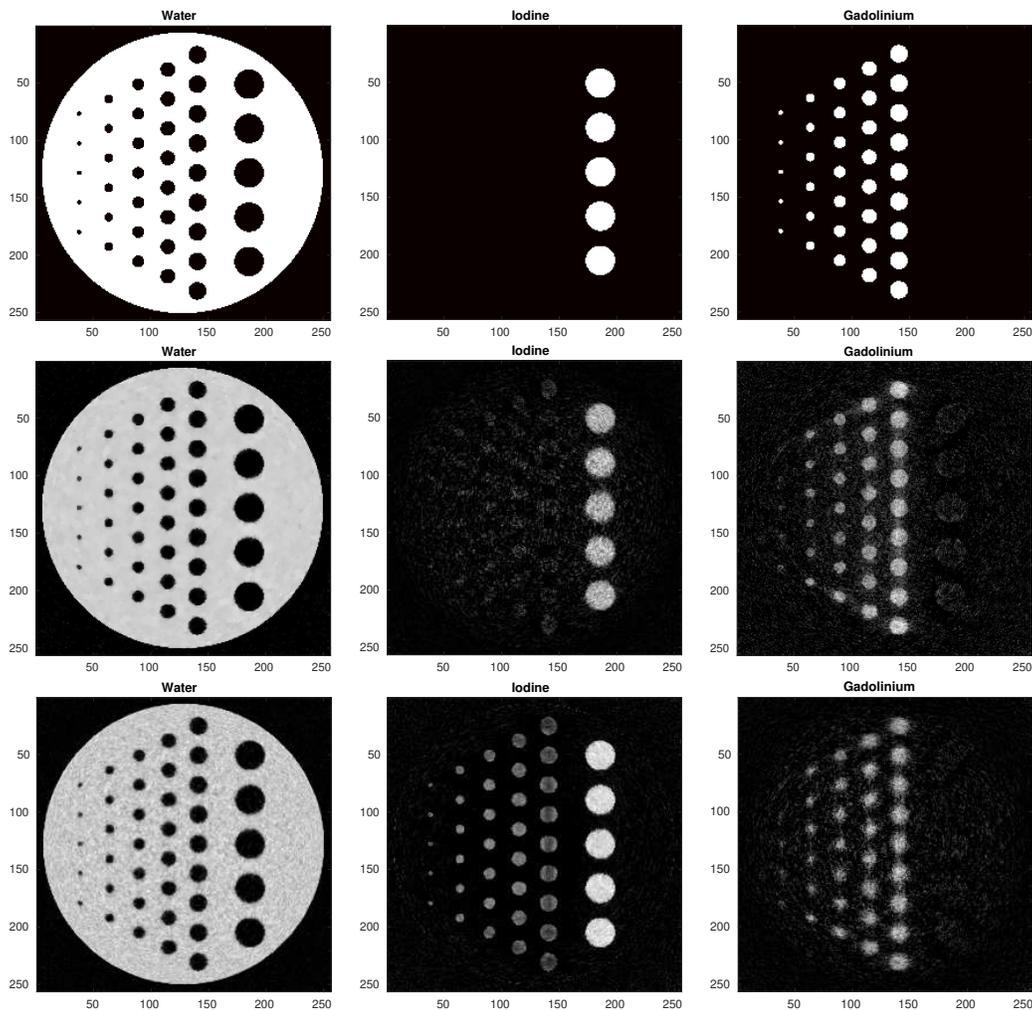
Figure 3: (a) Ground truth phantom with 3 materials: water, iodine, gadolinium, (b) qualitative materials estimation with the proposed Denoising-IHS algorithm, (c) qualitative materials estimation with algorithm [14] with non-adaptive edge-preserving hyperbola regularizer.

Instead, the block sampling procedure takes advantage, by averaging, of this block structure of the leverage scores induced by the CT operator as depicted in Figure 2(b) where each vertical red line corresponds to an adjacent angle and the inner point coincides with the detector array elements. Although this plot represents the exact sampling probabilities, for computational constraint the algorithm estimates these probabilities by exploiting the spectral decomposition of the ASTRA operator through (36).

We have tested the proposed algorithm on a synthetic phantom composed of three basis materials, water, iodine and gadolinium. The geometric shape of the 2D phantom is build used the model number 4 of the TomoPhantom Matlab software [54] containing three materials, water, iodine and gadolinium. The phantom is constituted of different circles with increasing diameter as shown in Figure 3(a). As denoiser-based regularizer D_{σ}, we have used the implementation of the U-Net [55] trained on a large set of 1000 images obtained by rotating and placing in different position the circles that are included in the test image. Furthermore, we have used the Autograd software in pytorch to compute the Jacobian-vector product of the U-Net denoiser [56]. The Hessian block sub-sampling for the Denoising-IHS was set to $s = \frac{m}{3}$. Figure 3(b) shows the quantitative reconstruction of the three materials while Figure 3(c) reports the results obtained using the multi-material framework (OS-PWSQS) proposed in [14] which considers a model-based optimization with an edge-preserving hyperbola regularizer. By comparing the images for each materials, it is possible to clearly note how the the Denoising-IHS approach is able to separate the each material and to reduces the noise in the reconstruction, as can be seen in particular for the water and iodine. The iodine image in figure 3(c) indicates that algorithm [14] does not manage to fully separate the iodine and gadolinium materials.
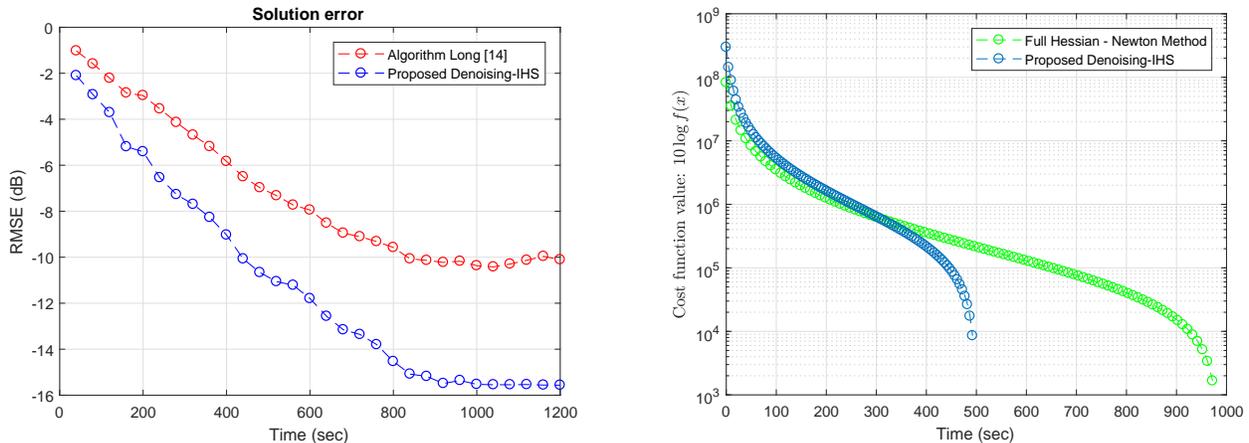
Figure 4: (a) Relative solution error comparison between the algorithm [14] and the proposed Denoising-IHS algorithm and (b) relative solution error comparison between full Hessian and the proposed Denoising-IHS algorithm.

Figure 4(a) shows the comparison between the root-mean-square error (RMSE), i.e. $\sqrt{\frac{1}{N}(\mathbf{x} - \mathbf{x}^*)}$ of the vector $\mathbf{x}$ containing the concatenation of the estimated materials respect to the ground truth $\mathbf{x}^*$ between the OS-PWSQS algorithm and the proposed algorithm. The solution achieved by the two algorithms is different because the cost function, both data fidelity term and regularizer, that OS-PWSQS [14] aims at minimizing is different from (11). By analyzing Figure 4(a), it is possible to highlight that the Denoising-IHS algorithm achieves a more accurate solution. Looking in detail at the blue curve, it is possible to note some local plateaus with sudden reduction of the error and

this is in agreement with the fact that in the outer loop the divergence of the denoiser is calculated to update the ridge leverage scores (29) which has a cost equivalent of applying the denoiser as stated in (30). Furthermore, Figure 4(b) shows the comparison in objective function between the convergence plot (semi-log) in dB of the Denoising-IHS algorithm and the solver which consider the full Hessian of the data loss at each iteration. This plot highlights how the strategy for block sub-sampling of the Hessian leads to a noticeable reduction in computation (almost 2.2 times) while both reach a similar minimum value in the cost function. We can analyze that at early iterations when the estimate is inaccurate both algorithms have similar accuracy/computation trade-off, while going further the Newton solver with full Hessian requires higher computation per iteration and overall it converges close to the Denoising-IHS solution but with increased computational time. From Figure 4(b) the time reduction is slightly less compared to the measurement sub-sampling rate and this could be anticipated since the ASTRA operator is not exactly linear in the number of measurements and the Denoising-IHS requires an additional denoising step at each outer iteration for the sampling probabilities update.

The following numerical test considers a more structured dataset, derived from the Adult Reference Computational Phantom (ICRP Publication 110, 2009), which is a segmented image of defined chemical composition to represent real tissues [57]. Figure 5(a) shows two slices selected for testing through the pelvis with image resolution 137×299 with a display window of $[0, 2]$.
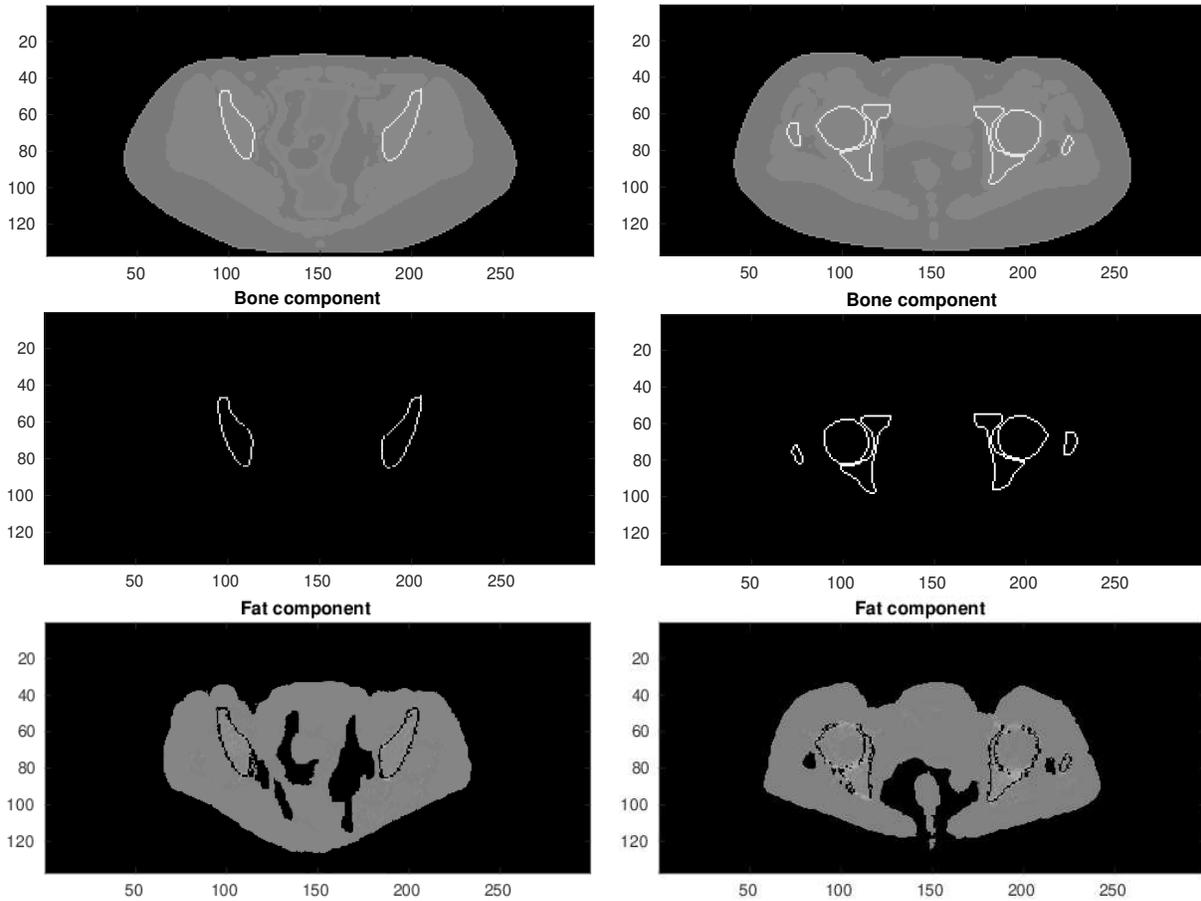
Figure 5: (a) Testing slices of the pelvis both with a display window $[0, 2]$, (b) reconstruction of the bone material of the two testing images with Denoising-IHS, (c) estimation of the fat material for the two testing images with Denoising-IHS with display window $[0, 0.15]$.

Figure 5(b) reports the qualitative estimation of the bone component in the testing images; it is possible to note how the Denoising-IHS algorithm is able to tackle the contours of the bone precisely in both images. Figure 5(c) shows the qualitative estimation of the fat component in the testing images; in this case the Denoising-IHS algorithm manages to estimate correctly most of the complicated structure in the pelvis with some spurious pixels missing.

To evaluate the quantitative estimation, Figure 6 reports the normalized oracle mass attenuation of the fat and the estimated one with Denoising-IHS related to longitudinal section at pixel $y = 80$ of the first test image.
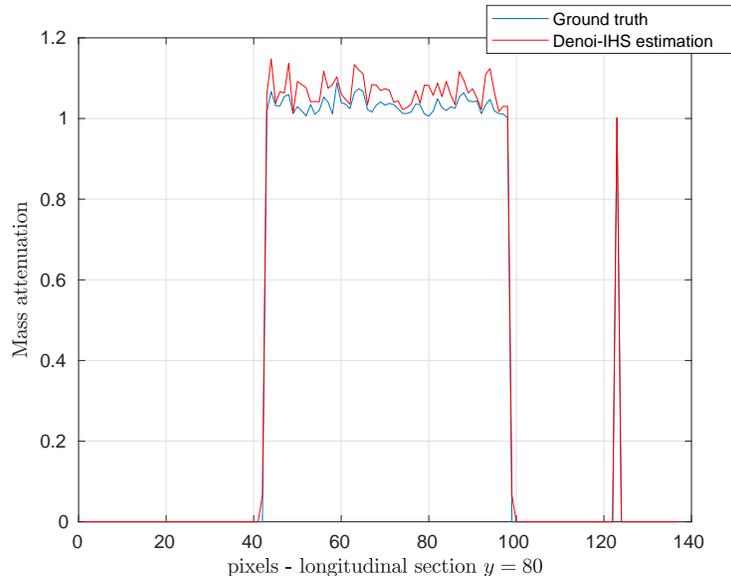
Figure 6: Ground truth values of the fat component for the longitudinal section of pixel 80 and the Denoising-IHS estimate.

This plot confirms that the proposed algorithm is able to accurately estimate the mass attenuation with and error which is less than 3%. In particular, Table 1 reports the root mean square error (RMSE) for different materials compared to the oracle ground truth.

Algorithms	Bone	Fat	Air
Denoising-IHS	0.031%	0.053%	0.025%
OS-PWSQS	0.041%	0.082%	0.035%

Table 1: RMSE comparison of the material images using Denoising-IHS and OS-PWSQS.

8   Preliminary Experimental Results

We have assessed the algorithm with a real acquisition using a photon counting detector CT scanner depicted in Figure 7(a) with a specimen in (b) containing cylinders of iodine and PMMA at different concentrations. The geometry is cone beam with source to axis of rotation distance of 8.1 cm and source to origin distance of 11.4 cm.

Figure 8(a) show the decrease of normalized cost function for the Denoising-IHS subsamplig case and full Hessian; it is worth noting the reduction in computation although is slower compared

17

Figure 7: (a) Scanner setup and (b) support specimen for the cylinders of different materials.

to the results achieved in the numerical simulations. Table (2) and Figures 8(b-c) show the quantitative and qualitative error in concentration of the iodine.

Table 2: Iodine concentration error

Ground truth [mg/ml]	estimated
8	8.15
16	15.72
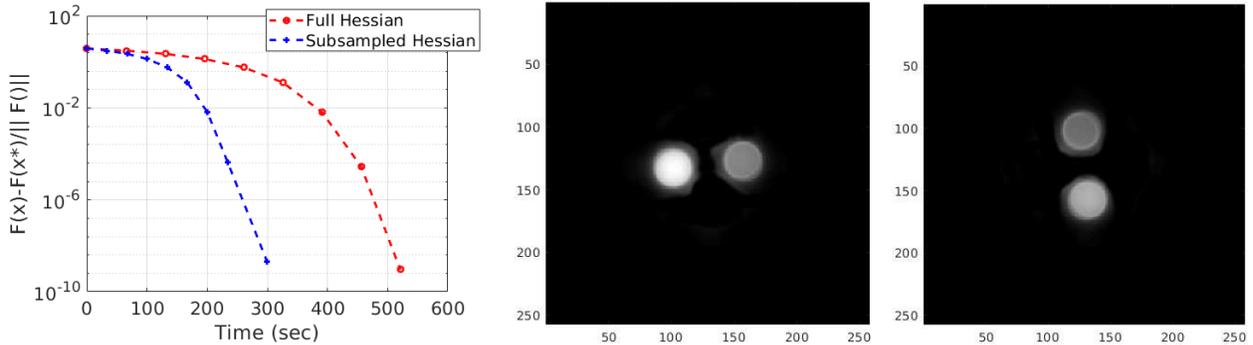
Figure 8: Error in objective function and reconstruction of the cylinders phantoms.

9   Conclusion

We propose a preliminary study for an efficient second order optimization method, Denoising-IHS, for multi-material decomposition in spectral CT which combines dimensionality reduction through Hessian matrix non-uniform random sub-sampling and regularization by denoising. A practical procedure for efficiently calculating the non-uniform sampling probabilities is shown which involves Fourier spectral approximation for estimating the ridge leverage scores. Furthermore, we analyze the CG solver which does not require the computation of the full Jacobian of the denoiser but only directional derivatives with a substantial reduction in computation. This problem is tested on numerical spectral CT simulations with Poisson noise model. We evaluate the convergence of the algorithm and we are able to show improvements both in terms of qualitative and quantitative accuracy compared to state-of-the-art algorithms. Furthermore, we show reduction in the

18

computational complexity respect to the Newton method involving the full Hessian matrix of the loss function. Further analysis to assess the proposed framework on more complex phantom with experimental acquisitions will be conducted together with a tighter analysis on the lower bounds for the dimensional reduction that it is possible to theoretically achieve and an analysis of the robustness of the algorithm to training sample dimension.

Acknowledgments

A   Convergence Analysis

Recalling the definition of the function $g(\mathbf{x}) = f(\mathbf{x}) + \rho_\nu(\mathbf{x})$ that we aim at minimizing as in (11) where we hide for simplicity the dependency over $\mathbf{y}$, we use the denotation $\bar{\mathbf{H}}(\mathbf{x}) = \nabla^2 g(\mathbf{x})$ to define the full Hessian, i.e. the one that would be obtained without block sub-sampling, and $\mathbf{H}^t(\mathbf{x}^t)$ the block sub-sampled Hessian defined in (19). We assume the following conditions

Assumption A.1. *The Hessian of the cost function $g(\mathbf{x})$ is Lipschitz continuous, therefore there is there exists a constant $L \geq 0$ such that*

$$\|\bar{\mathbf{H}}(\mathbf{w}) - \bar{\mathbf{H}}(\mathbf{z})\| \leq L\|\mathbf{w} - \mathbf{z}\|, \quad \forall \mathbf{w}, \mathbf{z} \in \mathbb{R}^n \tag{37}$$

Assumption A.2. *For any set S_t with $|S_t| = s$, there is a positive constant such that*

$$\lambda_{min}(\mathbf{H}^t(\mathbf{x}^t))\mathbf{I} \preceq \mathbf{H}^t(\mathbf{x}^t) \preceq \lambda_{max}(\mathbf{H}^t(\mathbf{x}^t))\mathbf{I} \tag{38}$$

Assuming $\bar{\mathbf{H}}(\mathbf{x}^t)$ the Hessian that would be obtained without block sub-sampling, the error compared to the Hessian obtained by data-loss term sketching is bounded according to the following condition

$$\|\mathbf{H}^t(\mathbf{x}^t) - \bar{\mathbf{H}}(\mathbf{x}^t)\| \leq \epsilon\|\bar{\mathbf{H}}(\mathbf{x}^t)\| \tag{39}$$

Given (39) we have

$$\begin{aligned}
\lambda_{min}(\mathbf{H}^t(\mathbf{x}^t)) &\geq \lambda_{min}(\bar{\mathbf{H}}(\mathbf{x}^t)) - \epsilon\lambda_{max}(\bar{\mathbf{H}}(\mathbf{x}^t)) = \left(1 - \epsilon\kappa(\bar{\mathbf{H}}(\mathbf{x}^t))\right)\lambda_{min}(\bar{\mathbf{H}}(\mathbf{x}^t)) \\
&\geq (1 - 2\epsilon\kappa)\lambda_{min}(\bar{\mathbf{H}}(\mathbf{x}^t)) = \lambda_b(\bar{\mathbf{H}}(\mathbf{x}^t))
\end{aligned} \tag{40}$$

where κ is the condition number of the problem.

We consider the expected error between the oracle solution and the estimate at iteration t where the expectation is taken over the random block selection at iteration t

$$\begin{aligned}
\mathbb{E}_t \left\|\mathbf{x}^{t+1} - \mathbf{x}_*\right\| &= \mathbb{E}_t \left\|\mathbf{x}^t + \mathbf{p}_r^t - \mathbf{x}_*\right\| \\
&\leq \underbrace{\mathbb{E}_t \left\|\mathbf{x}^t - \left[\mathbf{H}^t(\mathbf{x}^t)\right]^{-1}\nabla g(\mathbf{x}^t) - \mathbf{x}_*\right\|}_{1^{st} \text{ term}} + \underbrace{\mathbb{E}_t \left\|\mathbf{p}_r^t + \left[\mathbf{H}^t(\mathbf{x}^t)\right]^{-1}\nabla g(\mathbf{x}^t)\right\|}_{2^{nd} \text{ term}}
\end{aligned} \tag{41}$$

Regarding the 1^{st} term in (41), we have that

$$\mathbb{E}_t \left\| \mathbf{x}^t - \left[\mathbf{H}^t(\mathbf{x}^t) \right]^{-1} \nabla g(\mathbf{x}^t) - \mathbf{x}_* \right\| = \mathbb{E}_t \left\| \left[\mathbf{H}^t(\mathbf{x}^t) \right]^{-1} \left(\mathbf{H}^t(\mathbf{x}^t)(\mathbf{x}^t - \mathbf{x}_*) - \nabla g(\mathbf{x}^t) \right) \right\|$$

$$\leq \frac{1}{\lambda_b(\bar{\mathbf{H}}(\mathbf{x}^t))} \mathbb{E}_t \left\| \left(\mathbf{H}^t(\mathbf{x}^t) - \bar{\mathbf{H}}(\mathbf{x}^t) \right)(\mathbf{x}^t - \mathbf{x}_*) + \bar{\mathbf{H}}(\mathbf{x}^t)(\mathbf{x}^t - \mathbf{x}_*) - \nabla g(\mathbf{x}^t) \right\|$$

$$\leq \frac{1}{\lambda_b(\bar{\mathbf{H}}(\mathbf{x}^t))} \left[\underbrace{\mathbb{E}_t \left\| \left(\mathbf{H}^t(\mathbf{x}^t) - \bar{\mathbf{H}}(\mathbf{x}^t) \right)(\mathbf{x}^t - \mathbf{x}_*) \right\|}_{\text{term (a)}} + \underbrace{\left\| \nabla^2 g(\mathbf{x}^t)(\mathbf{x}^t - \mathbf{x}_*) - \nabla g(\mathbf{x}^t) \right\|}_{\text{term (b)}} \right]$$

The term (a) represents the error due to the Hessian sketching and using condition (39) we have

$$\mathbb{E}_t \left\| \left(\mathbf{H}^t(\mathbf{x}^t) - \bar{\mathbf{H}}(\mathbf{x}^t) \right)(\mathbf{x}^t - \mathbf{x}_*) \right\| \leq \epsilon \|\bar{\mathbf{H}}(\mathbf{x}^t)\| \|\mathbf{x}^t - \mathbf{x}_*\| = \epsilon \lambda_{max}\left(\bar{\mathbf{H}}(\mathbf{x}^t) \right) \|\mathbf{x}^t - \mathbf{x}_*\|$$

For the term (b), using the Lipschitz continuity of the Hessian (37) we have

$$\left\| \nabla^2 g(\mathbf{x}^t)(\mathbf{x}^t - \mathbf{x}_*) - \nabla g(\mathbf{x}^t) \right\| \leq \|\mathbf{x}^t - \mathbf{x}_*\| \int_{t=0}^1 \|\nabla^2 g(\mathbf{x}^t) - \nabla^2 g(\mathbf{x}^t + t(\mathbf{x}_* - \mathbf{x}^t)) dt\|$$

$$\leq \|\mathbf{x}^t - \mathbf{x}_*\|^2 \int_{t=0}^1 Lt dt = \frac{L}{2}\|\mathbf{x}^t - \mathbf{x}_*\|^2 \tag{42}$$

Therefore, the overall 1^{st} term is bounded by

$$\mathbb{E}_t \left\| \mathbf{x}^t - \left[\mathbf{H}^t(\mathbf{x}^t) \right]^{-1} \nabla g(\mathbf{x}^t) - \mathbf{x}_* \right\| \leq \frac{1}{(1 - 2\epsilon\kappa)\lambda_{min}(\bar{\mathbf{H}}(\mathbf{x}^t))} \left(\frac{L}{2}\|\mathbf{x}^t - \mathbf{x}_*\|^2 + \right.$$

$$\left. + \epsilon\lambda_{max}(\bar{\mathbf{H}}(\mathbf{x}^t))\|\mathbf{x}^t - \mathbf{x}_*\| \right) \tag{43}$$

The 2^{nd} term in (41) represents the residual error after r CG iterations. The convergence rate of the CG method varies at every iteration depending on the spectrum $\{\lambda_1 \leq \lambda_2 \leq, \ldots, \leq \lambda_n\}$ of the positive definite matrix $\mathbf{H}^t(\mathbf{x}^t)$. After r steps of the CG method applied to the linear system in (21), the iterate p_r^t satisfies

$$\left\| \mathbf{p}_r^t - \mathbf{p}_*^t \right\|_{\mathbf{C}}^2 \leq 2\sqrt{\kappa(\mathbf{C})} \left(\frac{\sqrt{\kappa(\mathbf{C})} - 1}{\sqrt{\kappa(\mathbf{C})} + 1} \right)^r \left\| \mathbf{p}_0^t - \mathbf{p}_*^t \right\|_{\mathbf{C}}^2, \quad \mathbf{C} = \mathbf{H}^t(\mathbf{x}^t) \tag{44}$$

where $\mathbf{p}_*^t$ indicates the exact solution and $\|\mathbf{x}\|_{\mathbf{C}}^2 \triangleq \mathbf{x}^T\mathbf{C}\mathbf{x}$. Using the CG convergence rate (44) and assuming $\mathbf{p}_0^t = 0$ and $\mathbf{p}_*^t = - \left[\bar{\mathbf{H}}(\mathbf{x}^t) \right]^{-1} \nabla g(\mathbf{x}^t)$, we obtain

$$\left\| \mathbf{p}_r^t + \left[\mathbf{H}^t(\mathbf{x}^t) \right]^{-1} \nabla g(\mathbf{x}^t) \right\|_{\mathbf{H}^t(\mathbf{x}^t)} \leq 2\sqrt{\kappa(\mathbf{C})} \left(\frac{\sqrt{\kappa(\mathbf{C})} - 1}{\sqrt{\kappa(\mathbf{C})} + 1} \right)^r \left\| \left[\mathbf{H}^t(\mathbf{x}^t) \right]^{-1} \nabla g(\mathbf{x}^t) \right\|_{\mathbf{H}^t(\mathbf{x}^t)} \tag{45}$$

Given $\|\mathbf{a}\|_{\mathbf{C}}^2 \leq \|\mathbf{b}\|_{\mathbf{C}}^2$, we can obtain the following relationship in terms of unweighted norm

$$\lambda_1\|\mathbf{a}\|^2 \leq \mathbf{a}^T\mathbf{C}\mathbf{a} \leq \mathbf{b}^T\mathbf{C}\mathbf{b} \leq \lambda_n\|\mathbf{b}\|^2 \Rightarrow \|\mathbf{a}\| \leq \sqrt{\kappa(\mathbf{C})}\|\mathbf{b}\| \leq \sqrt{\kappa}\|\mathbf{b}\| \tag{46}$$

By substituting the (45) in the 2^{nd} term of (41), it results

$$\mathbb{E}_t \left\| \mathbf{p}_r^t + \left[\mathbf{H}^t(\mathbf{x}^t)\right]^{-1} \nabla g(\mathbf{x}^t) \right\| \leq 2\sqrt{\kappa(\mathbf{C})} \left(\frac{\sqrt{\kappa(\mathbf{C})}-1}{\sqrt{\kappa(\mathbf{C})}+1} \right)^r \mathbb{E}_t \left\| \left[\mathbf{H}^t(\mathbf{x}^t)\right]^{-1} \nabla g(\mathbf{x}^t) \right\|$$

$$\leq 2\sqrt{\kappa(\mathbf{C})} \left(\frac{\sqrt{\kappa(\mathbf{C})}-1}{\sqrt{\kappa(\mathbf{C})}+1} \right)^r \nabla g(\mathbf{x}^t) \mathbb{E}_t \left\| \left[\mathbf{H}^t(\mathbf{x}^t)\right]^{-1} \right\|$$

$$\leq \frac{2L\sqrt{\kappa(\mathbf{C})}}{\mu} \left(\frac{\sqrt{\kappa(\mathbf{C})}-1}{\sqrt{\kappa(\mathbf{C})}+1} \right)^r \left\| \mathbf{x}^t - \mathbf{x}_* \right\| \qquad (47)$$

By substituting the expressions (43) and (47) in (41), we have

$$\mathbb{E}_t \left\| \mathbf{x}^{t+1} - \mathbf{x}_* \right\| \leq \frac{1}{(1-2\epsilon\kappa)\,\lambda_{min}(\bar{\mathbf{H}}(\mathbf{x}^t))} \left(\frac{L}{2} \|\mathbf{x}^t - \mathbf{x}_*\|^2 + \right.$$

$$\left. + \epsilon\lambda_{max}(\bar{\mathbf{H}}(\mathbf{x}^t))\|\mathbf{x}^t - \mathbf{x}_*\| \right) + 2\sqrt{\kappa(\mathbf{C})} \left(\frac{\sqrt{\kappa(\mathbf{C})}-1}{\sqrt{\kappa(\mathbf{C})}+1} \right) \|\mathbf{x}^t - \mathbf{x}_*\|$$

$$\leq \frac{L}{2\,(1-2\epsilon\kappa)\,\lambda_{min}(\bar{\mathbf{H}}(\mathbf{x}^t))} \left\| \mathbf{x}^t - \mathbf{x}_* \right\|^2 +$$

$$+ \frac{\epsilon\lambda_{max}(\bar{\mathbf{H}}(\mathbf{x}^t))}{2\,(1-2\epsilon\kappa)\,\lambda_{min}(\bar{\mathbf{H}}(\mathbf{x}^t))} \left\| \mathbf{x}^t - \mathbf{x}_* \right\| + 2\sqrt{\kappa(\mathbf{C})} \left(\frac{\sqrt{\kappa(\mathbf{C})}-1}{\sqrt{\kappa(\mathbf{C})}+1} \right) \left\| \mathbf{x}^t - \mathbf{x}_* \right\|$$

References

[1] R. E. Alvarez and A. Macovski, "Energy-selective reconstructions in X-ray computerised tomography," *Physics in Medicine & Biology*, vol. 21, no. 5, p. 733, 1976.

[2] A. M. Alessio and L. R. MacDonald, "Quantitative material characterization from multi-energy photon counting CT," *Medical Physics*, vol. 40, no. 3, p. 031108, 2013.

[3] L. Yu, J. A. Christner, S. Leng, J. Wang, J. G. Fletcher, and C. H. McCollough, "Virtual monochromatic imaging in dual-source dual-energy CT: Radiation dose and image quality," *Medical Physics*, vol. 38, no. 12, pp. 6371–6379, 2011.

[4] E. Roessl and R. Proksa, "K-edge imaging in X-ray computed tomography using multi-bin photon counting detectors," *Physics in Medicine & Biology*, vol. 52, p. 4679, 2007.

[5] S. Kappler, D. Niederlöhner, K. Stierstorfer, and T. Flohr, "Contrast-enhancement, image noise, and dual-energy simulations for quantum-counting clinical CT," in *Medical Imaging: Physics of Medical Imaging*. Proceedings of SPIE, 2010.

[6] R. E. Alvarez, "Estimator for photon counting energy selective X-ray imaging with multibin pulse height analysis," *Medical Physics*, vol. 38, no. 5, pp. 2324–2334, 2011.

[7] G. Vilches-Freixas, J. M. Létang, N. Ducros, and S. Rit, "Optimization of dual-energy CT acquisitions for proton therapy using projection-based decomposition," *Medical Physics*, vol. 44, no. 9, pp. 4548–4558, 2017.

[8] C. Maaß, M. Baer, and M. Kachelrieß, "Image-based dual energy CT using optimized pre-correction functions: A practical new approach of material decomposition in image domain," *Medical Physics*, vol. 36, no. 8, pp. 3818–3829, 2009.

[9] H. Gao, H. Yu, S. Osher, and G. Wang, "Multi-energy CT based on a prior rank, intensity and sparsity model (PRISM)," *Inverse Problems*, vol. 27, p. 115012, 2011.

[10] K. Kim, J. C. Ye, W. Worstell, J. Ouyang, Y. Rakvongthai, G. Fakhri, and Q. Li, "Sparse-view spectral CT reconstruction using spectral patch-based low-rank penalty," *IEEE Trans. on Medical Imaging*, vol. 34, no. 3, 2015.

[11] D. Kazantsev, J. S. Jørgensen, M. S. Andersen, W. Lionheart, P. D. Lee, and P. J. Withers, "Joint image reconstruction method with correlative multi-channel prior for X-ray spectral computed tomography," *Inverse Problems*, vol. 34, p. 064001, 2018.

[12] W. Wu, D. Hu, K. An, S. Wang, and F. Luo, "A high-quality photon-counting CT technique based on weight adaptive total-variation and image-spectral tensor factorization for small animals imaging," *IEEE Trans. on Instrumentation and Measurement*, vol. 70, pp. 1–14, 2020.

[13] W. Wu, P. Chen, S. Wang, V. Vardhanabhuti, F. Liu, and H. Yu, "Image-domain material decomposition for spectral CT using a generalized dictionary learning," *IEEE Trans. on Radiation and Plasma Medical Sciences*, 2020.

[14] Y. Long and J. A. Fessler, "Multi-material decomposition using statistical image reconstruction for spectral CT," *IEEE Trans. on Medical Imaging*, vol. 33, no. 8, pp. 1614–1626, 2014.

[15] R. Barber, E. Y. Sidky, T. G. Schmidt, and X. Pan, "An algorithm for constrained one-step inversion of spectral CT data," *Phys. Med. Biol*, vol. 61, pp. 3784–3818, 2016.

[16] T. G. Schmidt, R. F. Barber, and E. Y. Sidky, "A spectral CT method to directly estimate basis material maps from experimental photon-counting data," *IEEE Trans. on Medical Imaging*, vol. 36, no. 9, pp. 1808–1819, 2017.

[17] T. Weidinger, T. M. Buzug, T. Flohr, S. Kappler, and K. Stierstorfer, "Polychromatic iterative statistical material image reconstruction for photon-counting computed tomography," *International Journal of Biomedical Imaging*, vol. 2016, pp. 1–15, 2016.

[18] K. Mechlem, S. Ehn, T. Sellerer, E. Braig, D. Münzel, F. Pfeiffer, and P. B. Noël, "Joint statistical iterative material image reconstruction for spectral computed tomography using a semi-empirical forward model," *IEEE Trans. on Medical Imaging*, vol. 37, no. 1, pp. 68–80, 2018.

[19] D. Needell, R. Ward, and N. Srebro, "Stochastic gradient descent, weighted sampling, and the randomized kaczmarz algorithm," in *Advances in neural information processing systems*, 2014, pp. 1017–1025.

[20] J. Konečný, J. Liu, P. Richtárik, and M. Takáč, "Mini-batch semi-stochastic gradient descent in the proximal setting," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 2, pp. 242–255, 2015.

[21] D. Kim, S. Ramani, and J. A. Fessler, "Combining ordered subsets and momentum for accelerated x-ray CT image reconstruction," *IEEE Trans. on Medical Imaging*, vol. 34, no. 1, pp. 167–178, 2014.

[22] A. Sawatzky, Q. Xu, C. O. Schirra, and A. M. A., "Proximal ADMM for multi-channel image reconstruction in spectral X-ray CT," *IEEE Trans. on Medical Imaging*, vol. 33, pp. 1637–368, 2014.

[23] S. Ramani and J. A. Fessler, "A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction," *IEEE Trans. on Medical Imaging*, vol. 31, no. 3, pp. 677–688, 2011.

[24] R. M. Gower, N. Loizou, X. Qian, A. Sailanbayev, E. Shulgin, and P. Richtárik, "SGD: General analysis and improved rates," in *International Conference on Machine Learning*, 2019, pp. 5200–5209.

[25] J. He, Y. Yang, Y. Wang, D. Zeng, Z. Bian, H. Zhang, J. Sun, Z. Xu, and J. Ma, "Optimizing a parameterized plug-and-play ADMM for iterative low-dose CT reconstruction," *IEEE Trans. on Medical Imaging*, vol. 38, no. 2, pp. 371–382, 2018.

[26] M. Schmidt, D. Kim, and S. Sra, "Projected newton-type methods in machine learning," in *Optimization for Machine Learning*. MIT Press, 2011, pp. 305–330.

[27] J. D. Lee, Y. Sun, and M. A. Saunders, "Proximal newton-type methods for minimizing composite functions," *SIAM Journal on Optimization*, vol. 24, no. 3, pp. 1420–1443, 2014.

[28] J. Nocedal and S. Wright, "Numerical optimization," *Springer Science & Business Media*, 2006.

[29] H. G. Gene and F. Charles, "Matrix computations," *Johns Hopkins Universtiy Press, 3rd edtion*, 1996.

[30] D. P. Woodruff, "Sketching as a tool for numerical linear algebra," *Foundations and Trends® in Theoretical Computer Science*, vol. 10, no. 1–2, pp. 1–157, 2014.

[31] R. Bollapragada, R. H. Byrd, and J. Nocedal, "Exact and inexact subsampled newton methods for optimization," *IMA Journal of Numerical Analysis*, vol. 39, no. 2, pp. 545–578, 2019.

[32] A. S. Berahas, R. Bollapragada, and J. Nocedal, "An investigation of newton-sketch and subsampled newton methods," *Optimization Methods and Software*, pp. 1–20, 2020.

[33] M. Pilanci and M. J. Wainwright, "Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence," *SIAM Journal on Optimization*, vol. 27, no. 1, pp. 205–245, 2017.

[34] R. H. Byrd, G. M. Chin, W. Neveitt, and J. Nocedal, "On the use of stochastic hessian information in optimization methods for machine learning," *SIAM Journal on Optimization*, vol. 21, no. 3, pp. 977–995, 2011.

[35] M. A. Erdogdu and A. Montanari, "Convergence rates of sub-sampled newton methods," *arXiv preprint arXiv:1508.02810*, 2015.

[36] F. Roosta-Khorasani and M. W. Mahoney, "Sub-sampled newton methods i: globally convergent algorithms," *arXiv preprint arXiv:1601.04737*, 2016.

[37] P. Xu, J. Yang, F. Roosta, C. Ré, and M. W. Mahoney, "Sub-sampled newton methods with non-uniform sampling," in *Advances in Neural Information Processing Systems*, 2016, pp. 3000–3008.

[38] F. Roosta-Khorasani and M. W. Mahoney, "Sub-sampled newton methods ii: Local convergence rates," *arXiv preprint arXiv:1601.04738*, 2016.

[39] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *IEEE Global Conference on Signal and Information Processing*, 2013, pp. 945–948.

[40] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM Journal on Imaging Sciences*, vol. 10, no. 4, pp. 1804–1844, 2017.

[41] E. T. Reehorst and P. Schniter, "Regularization by denoising: Clarifications and new interpretations," *IEEE Trans. on Computational Imaging*, vol. 5, no. 1, pp. 52–67, 2019.

[42] Y. Sun, B. Wohlberg, and U. S. Kamilov, "An online plug-and-play algorithm for regularized image reconstruction," *IEEE Trans. on Computational Imaging*, vol. 5, no. 3, pp. 395–408, 2019.

[43] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *IEEE Trans. on image processing*, vol. 16, no. 8, pp. 2080–2095, 2007.

[44] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.

[45] Y. Sun, J. Liu, and U. Kamilov, "Block coordinate regularization by denoising," in *Advances in Neural Information Processing Systems*, 2019, pp. 382–392.

[46] Y. Hu, M. Andersen, and J. G. Nagy, "Spectral computed tomography with linearization and preconditioning," *SIAM Journal on Scientific Computing*, vol. 41, no. 5, pp. S370–S389, 2019.

[47] A. Hyvärinen, "Estimation of non-normalized statistical models by score matching," *Journal of Machine Learning Research*, vol. 6, no. Apr, pp. 695–709, 2005.

[48] P. Vincent, "A connection between score matching and denoising autoencoders," *Neural computation*, vol. 23, no. 7, pp. 1661–1674, 2011.

[49] W. Van Aarle, W. J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabravolski, J. De Beenhouwer, K. J. Batenburg, and J. Sijbers, "Fast and flexible x-ray tomography using the astra toolbox," *Optics express*, vol. 24, no. 22, pp. 25 129–25 147, 2016.

[50] S. Ramani, T. Blu, and M. Unser, "Monte-carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms," *IEEE Trans. on Image Processing*, vol. 17, no. 9, pp. 1540–1554, 2008.

[51] S. Matej, J. A. Fessler, and I. Kazantsev, "Iterative tomographic image reconstruction using fourier-based forward and back-projectors," *IEEE Trans. on Medical Imaging*, vol. 23, no. 4, pp. 401–412, 2004.

[52] Y. O'Connor and J. A. Fessler, "Fourier-based forward and back-projectors in iterative fan-beam tomographic image reconstruction," *IEEE Trans. on Medical Imaging*, vol. 25, no. 5, pp. 582–589, 2006.

[53] J. H. Siewerdsen, A. M. Waese, D. J. Moseley, S. Richard, and D. A. Jaffray, "Spektr: A computational tool for x-ray spectral analysis and imaging system optimization," *Medical physics*, vol. 31, no. 11, pp. 3057–3067, 2004.

[54] D. Kazantsev, V. Pickalov, S. Nagella, E. Pasca, and P. J. Withers, "Tomophantom, a software package to generate 2d–4d analytical phantoms for CT image reconstruction algorithm benchmarks," *SoftwareX*, vol. 7, pp. 150–155, 2018.

[55] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention.* Springer, 2015, pp. 234–241.

[56] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017.

[57] J. H. Mason, A. Perelli, W. H. Nailon, and M. E. Davies, "Polyquant ct: direct electron and mass density reconstruction from a single polyenergetic source," *Physics in Medicine & Biology*, vol. 62, no. 22, p. 8739, 2017.

[58] A. Perelli and M. S. Andersen, "Data support for "Regularization by denoising sub-sampled newton method for spectral CT multi-material decomposition," https://doi.org/10.5281/zenodo.4482071, 2021.