



HAL
open science

Comparison of five one-step reconstruction algorithms for spectral CT

Cyril Mory, Bruno Sixou, Salim Si-Mohamed, Loic Boussel, Simon Rit

► **To cite this version:**

Cyril Mory, Bruno Sixou, Salim Si-Mohamed, Loic Boussel, Simon Rit. Comparison of five one-step reconstruction algorithms for spectral CT. 2018. hal-01760845v1

HAL Id: hal-01760845

<https://hal.science/hal-01760845v1>

Preprint submitted on 6 Apr 2018 (v1), last revised 10 Feb 2019 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparison of five one-step reconstruction algorithms for spectral CT

Cyril Mory, Bruno Sixou, Salim Si-mohamed, Loïc Boussel and Simon Rit

Abstract—Over the last decade, dual-energy CT scanners have gone from prototypes to clinically available machines, and spectral photon counting CT scanners are following. They require a specific reconstruction process, consisting of two steps: material decomposition and tomographic reconstruction. The two steps can be done separately in either order, but in both cases, some information is lost along the way. As an alternative, “one-step inversion” methods have been proposed, which perform decomposition and reconstruction simultaneously. For most CT applications, reconstruction time is critical for practical usability, and one-step methods are typically slower than their two-step counterparts. The goal of this paper is to provide an independent comparison of five one-step inversion algorithms, focused mainly on convergence speed, but also on memory footprint, stability, and ease of use. We adapted and implemented a Bayesian method which uses non-linear conjugate gradient for minimization [1], three methods based on quadratic surrogates [2, 3, 4], and a primal-dual method based on MOCCA, a modified Chambolle-Pock algorithm [5]. Experiments were performed on both simulated and real data. Some of these methods can be accelerated by using μ -preconditioning, i.e. by performing all internal computations not with the actual materials the object is made of, but with carefully chosen linear combinations of those. In this paper, we also evaluate the impact of three different μ -preconditioners on convergence speed. Our results show that the method of Mechlem *et al.* [4] is much faster than the others, while being only slightly less stable and more complex: it requires less than 100 iterations, versus several thousands for other methods. It seems to be the only viable candidate for implementation into a real multi-energy scanner.

I. INTRODUCTION

Dual energy computed tomography (CT) systems are now commercially available, and spectral photon counting CT systems already exist as research prototypes. These new types of scanners provide information on the energy distribution of the X-ray photons that reach the detector: dual energy systems acquire two sets of projections (for low and high energies), while spectral photon counting systems group the incoming photons into so-called “energy bins”, which form a partition of the incident spectrum’s bandwidth. For each projection’s pixel, the “raw” measurements acquired by a spectral photon counting CT scanner are therefore the number of photons in each energy bin, which are commonly referred to as “photon counts”. Assuming that the object’s attenuation can be obtained by linear combination of the attenuations of only a few materials, this energy-resolved information allows to reconstruct multiple volumes,

C. Mory & S. Rit are with the Université de Lyon, CREATIS ; CNRS UMR5220 ; Inserm U1206 ; INSA-Lyon ; Université Lyon 1 ; Centre Léon Bérard, France.

Bruno Sixou is with the Université de Lyon, CREATIS ; CNRS UMR5220 ; Inserm U1206 ; INSA-Lyon ; Université Lyon 1, France.

Loïc Boussel and Salim Si-mohamed are with the Université de Lyon, CREATIS ; CNRS UMR5220 ; Inserm U1206 ; INSA-Lyon ; Université Lyon 1, France and with the Radiology Department, Lyon University Hospital, Lyon, France

each one representing a different material’s concentration map. Alternatively, one can assume that the attenuation stems from only a few physical phenomena, e.g. Compton scattering, photoelectric effect and K-edge [6, 7], and reconstruct maps of the characteristic quantities of each of these phenomena. For the sake of readability, focus on spectral photon counting CT and a basis of materials, but almost everything in the present paper also holds for dual energy CT and phenomena, and the other combinations of device and decomposition method. It is assumed that the attenuation coefficient of the imaged object μ , which depends on the spatial position in the object and on the energy of the incident photons, can be separated into a function of space and a function of energy, i.e.

$$\mu(\vec{s}, E) = \sum_{m=1}^{N_m} \mu_m(E) x_m(\vec{s}) \quad (1)$$

where \vec{s} is the position in space, E the energy of an incident photon, N_m the number of materials, $\mu_m(E)$ the mass-attenuation coefficient of material m at energy E , and $x_m(\vec{s})$ the concentration of material m at position \vec{s} . The attenuation profiles μ_m are known, and the aim is to find the concentration maps x_m . The most widespread methods to reconstruct material-specific volumes from photon counts can be divided into two categories. First, two-step image-based methods, in which each energy bin of the photon counts is log-transformed and reconstructed, just like the projections dataset of a standard CT. This process yields one volume per energy bin and these intermediate volumes are then decomposed into material-specific volumes [8, 9]. Unfortunately, the intermediate volumes are corrupted by beam hardening artefacts, because the energy dependence of the attenuation coefficients is still averaged in each energy bin, and also because a vast proportion of the primary photons reaching the detector lose energy during interactions inside the detector, and are therefore detected in a lower energy bin, which is not accounted for during the reconstruction. Second, two-step projection-based methods, in which multi-channel projections are first decomposed into material-specific projections, and then reconstructed, independently [6, 10] or jointly [11]. It turns out that for typical choices of materials (e.g. water, bone, and a high-Z contrast agent), the normalized attenuation profiles of the materials do not differ enough to yield a robust decomposition: the inevitable statistical noise on photon counts is often sufficient to cause aberrant material line integrals, which result in strong streak artefacts in the reconstructed volumes. Some methods were proposed to regularize the decomposition process [12, 13], but projections regularized using standard priors (e.g. total variation) usually violate data consistency conditions [14, 15]. In addition, projection-based methods cannot be used easily with fast kV switching [16] or dual source dual-energy systems [17].

Two-step methods also have a structural drawback: unless the first step is a bijection, it implies a loss of information, for which the second step cannot compensate.

Recently, several methods have been proposed which reconstruct material-specific volumes directly from the photon counts [1, 2, 3, 5, 4, 18]. They are commonly referred to as “one-step inversion”, or simply “one-step” methods. All of these methods are iterative: there is currently no analytical inversion formula for the material decomposition problem, let alone for one-step inversion. They consist in combining the forward models of the tomographic reconstruction and the material separation inverse problems, yielding a single (but more complex) forward model, which takes as input a set of material volumes, and yields photon counts as output. In discrete form, the forward model reads

$$y_{ib} = \sum_E S_{b,E} \exp \left(- \sum_j a_{ij} \sum_m \mu_m(E) x_{jm} \right) \quad (2)$$

where y_{ib} is the photon count in pixel i and bin b , $S_{b,E}$ is the effective spectrum, i.e. the number of photons of actual energy E that are expected to be detected in bin b if there was no attenuating object, a_{ij} is the coefficient of the forward projection matrix A at row i and column j , and x_{jm} is the concentration of material m in voxel j . One-step inversion then means finding the volumes x_m for which the forward model yields photon counts y as similar as possible to those measured by the scanner. By construction, the drawbacks of two-step methods are circumvented, and there is no risk of losing information. One-step methods therefore have the potential to yield higher image quality than their two-step counterparts, but they are also slower to converge. This is likely to prevent their use in many applications, as it happened for regularized iterative reconstruction techniques on standard CT [19]. Since in the multi-energy CT case the amount of data contained in a single acquisition is multiplied by the number of energy bins, and the one-step problem is more complex than the reconstruction problem alone, the need for fast methods is even more acute than for standard CT.

In this paper, we present a comparison of five one-step methods [1, 2, 3, 4, 5]. Among these, only Cai *et al.* [1] reported a comparison with another one-step method [20] and an advanced image-based two-step method [21]. Barber *et al.* [5] showed many results, but only with the proposed method, and the other ones [2, 3, 4] are compared to two-step methods, either projection-based or image-based [22], with the exception that Mechlem *et al.* [4] contains a comparison with Weidinger *et al.* [3] in terms of convergence speed. An independent comparison between many one-step methods is therefore still missing. In the present work, experiments were conducted on both simulated and real data on a single slice. A strong emphasis was put on convergence speed, but other aspects like memory footprint, stability, necessity to tune additional parameters and mathematical rigour were also taken into account. Adequately comparing the image quality of the reconstruction results of each method is notoriously difficult [23], and would have required many choices (on datasets, method to determine the regularization parameters, convergence criteria and image quality metrics). It is therefore beyond the scope of the present paper.

In [5, 24], it is proposed to perform internal computations not

with the actual materials into which one wishes to decompose the object, but with carefully chosen linear combinations of these materials. Since the materials are characterized only by their attenuation coefficient μ as a function of the incident energy, Barber *et al.* call this step μ -preconditioning. In this paper, we present three different ways to perform μ -preconditioning and evaluate their impact on convergence speed. Since most methods were not explicitly designed for μ -preconditioning, a small adaptation is required (and described in the present paper) for the regularization step.

II. MATERIAL AND METHODS

A. Compared methods

This paper is a comparison of five one-step inversion methods. We have tried to stick as much as possible to the original algorithm description but some modifications were required. In this section, we briefly introduce each method, and describe the modifications we had to perform. Though we have tried to make the present paper as self-contained as possible, it is impossible to explain all the methods in details without ending up duplicating large passages of the original papers. For the same reason, we sometimes use the notations of the original papers without always re-defining them, but only for very specific comments, which anyway do not make sense without reading the original papers, and are not necessary to understand the message of the present paper. We use a contraction of the first author’s name and publication year (e.g. “Cai2013”) to refer to each method. When the distinction is relevant, it should be clear from the context whether we mean the method implemented and modified by us, or the one described in the original papers.

1) *Cai2013*: Cai2013 [1] was originally published for dual energy CT. The cost function is built with a variable y that is the ratios between photon counts with object and photon counts without object (instead of just the photon counts with object, in the other methods). The variance of these ratios y is assumed to be proportional to y , with a factor k_d that has to be estimated. The regularization term is the Huber function of the spatial gradient of the volumes. The cost function is minimized by a non-linear conjugate gradient, with a heuristically determined step size. Given the descent direction, the step size is computed as the one that would minimize the cost function’s second-order Taylor expansion. Then the cost function is computed at the candidate destination point, and the step size divided by 2 if it causes an increase of the cost instead of a decrease. According to the authors, the case where the step length had to be divided by 2 never occurred in the experiments they conducted (which were limited to dual-energy CT). In ours, it did occur, and we had to implement a reset of the descent direction to the opposite of the gradient (instead of the one provided by non-linear CG, which includes the descent direction at previous iteration) when the step size has been divided by 2 too many times (in our implementation, 10 times) and still yields an ascent step. Cai2013 is the generalization of [1] to 3 materials, which is straightforward, and k_d was determined by simulations. Note that Cai2013 does not minimize the same data-attachment term as the other four methods.

2) *Long2014*: Long2014 [2] was originally published as a method for reconstructing three or more materials from dual

energy CT scans. To that end, it adds a lot of *a priori* information, in the form of equality and inequality constraints on the reconstructed volumes. It uses surrogates of the cost function to minimize it. Surrogates are defined as follows: let $\Psi : \mathbb{R}^N \rightarrow \mathbb{R}$ and $x_0 \in \mathbb{R}^N$. The function $\Phi_{x_0} : \mathbb{R}^N \rightarrow \mathbb{R}$ is a surrogate of Ψ at x_0 if and only if

$$\begin{cases} \forall x \in \mathbb{R}^N, \Phi_{x_0}(x) \geq \Psi(x) \\ \Phi_{x_0}(x_0) = \Psi(x_0) \end{cases} \quad (3)$$

i.e. if Φ_{x_0} is above Ψ on \mathbb{R}^N , and tangent to Ψ at x_0 . Note that we have used \mathbb{R}^N as domain for simplicity, but in Long2014 the domain actually used is \mathbb{R}^{+N} . The algorithms based on surrogates consist in finding a surrogate Φ_{x_0} to the cost function Ψ at the current iterate x_0 , minimizing the surrogate instead of the original cost function, adopting that minimizer as the new iterate x_1 , and starting over. Obviously the surrogates must be chosen carefully, in order to be easier to minimize than the original cost function. In [2], surrogates are derived one after the other: first a convex one, then a quadratic one, and finally a separable quadratic one, each new one being a surrogate for all the previous ones, and therefore for the original cost function. Performing minimization while complying with all constraints requires an inner loop of quadratic programming using the Generalized Sequential Minimal Optimization (GSMO) algorithm [25]. Since in our case, there are more bins than the number of materials, we have removed all constraints in our implementation of Long2014, which allowed to replace the quadratic programming inner loop by a single iteration of Newton's method. Long2014 uses Ordered Subsets (OS) to speed up convergence.

3) *Weidinger2016*: Weidinger2016 [3] also minimizes the Poisson negative log-likelihood of the data by Separable Quadratic Surrogates (SQS). As in [2], surrogates are derived sequentially. However, the quadratic "surrogate" is actually not a surrogate, but an approximation, of the convex surrogate. It simplifies the calculations, but the method is no longer guaranteed to converge.

4) *Mechlem2017*: Mechlem2017 [4] builds upon [3], adding three features: calibration, Ordered Subsets and Nesterov acceleration. The calibration part allows to estimate the product of the incident spectrum and detector response, in cases where it is either not known at all, or known with insufficient accuracy. In this study, we have assumed perfect knowledge of these system characteristics, and therefore dropped the calibration part. In addition to that, Mechlem2017 [4] suggests to compute the Hessian (and its inverse) only once, thus updating only the gradient at each iteration. This approach is only valid if the Hessian is near-constant, i.e. if the estimated volumes undergo little change over the course of the iterations. Therefore, the paper proposes to start the iterations from a previously known approximate result, like a set of two-step reconstructed volumes. Since we start from zero-filled estimates, we do not use this acceleration.

5) *Barber2016*: Barber2016 [5] minimizes the transmission Poisson likelihood (or a least squares term) plus a Total Variation (TV) regularization term, using the mirrored convex/concave (MOCCA) algorithm [26]. MOCCA is a primal-dual algorithm designed to minimize non-convex functions that can be expressed as the difference of two convex functions. As

explained in section 3.1 of [5], the algorithm should have an outer and an inner loop, but to speed up convergence, the inner loop is set to perform only a single iteration. Unfortunately, that modification can cause the algorithm to diverge, so an additional parameter λ has been introduced and tuned to avoid that. Roughly, λ is tuned to the largest value that does not cause divergence. We also adapted the parameter θ of the Chambolle-Pock algorithm (see the last line of equation 7 in [27]), which must lie in $[0, 1]$, and in [5] is implicitly set to 1. With $\theta = 1$, examining the sequence of iterates, we observed fast and ample oscillations, and artifacts constantly appearing and disappearing. These oscillations and artifacts did not significantly decrease in intensity over the course of the iterations. Setting $\theta = 0.5$ got rid of them, yielding a more stable sequence of iterates. Computing the primal and dual step sizes Σ and T involves the element-wise absolute value of products of matrices, which is not trivial if both matrices have positive and negative values and are too large to be actually stored. The computation of the step sizes proved particularly tedious to implement, so we restricted our study to a μ -preconditioning method (see section II-B) that does not generate negative attenuation values (i.e. simply normalizing the attenuation coefficient vectors), which makes simplifications possible. This μ -preconditioning is different from the one presented in the original paper [5] (i.e. orthonormalizing the attenuation coefficient vectors), but our study on Cai2013 in section III-F shows a better convergence speed for normalization than for orthonormalization, so it is unlikely that switching to orthonormalization for Barber2016 would greatly improve its convergence speed.

B. μ -preconditioning

We form the matrix M as follows:

$$M = \begin{pmatrix} \mu_I(1\text{keV}) & \mu_{\text{Gd}}(1\text{keV}) & \mu_{\text{H}_2\text{O}}(1\text{keV}) \\ \vdots & \vdots & \vdots \\ \mu_I(150\text{keV}) & \mu_{\text{Gd}}(150\text{keV}) & \mu_{\text{H}_2\text{O}}(150\text{keV}) \end{pmatrix} \quad (4)$$

where $\mu_m(E)$ is the mass-attenuation coefficient of material m for incident X-ray photons of energy E , I stands for iodine, Gd for gadolinium and H₂O for water. For $E \leq 150$ keV, $\mu_{\text{H}_2\text{O}}(E)$ and $\mu_{\text{Gd}}(E)$ typically differ by 2 orders of magnitude, which causes conditioning issues in the resolution of the one-step inversion problem, and slows down convergence. Section 3.4 of [5] proposes to modify the material attenuation matrix M to speed up and stabilize the internal computations. They call this step " μ -preconditioning". It can be interpreted as creating synthetic materials, which are linear combinations of the original ones, possibly with negative coefficients. We present three different μ -preconditioning methods. Each one yields a synthetic material attenuation matrix \tilde{M} and a matrix P such that $MP = \tilde{M}$:

- Normalization: \tilde{M} is obtained by dividing each column of M by its norm. This has the advantage of not introducing negative attenuations. The added complexity associated with the use of negative attenuation coefficients, like the use of absolute values in the calculation of $\Sigma^{(n)}$ and $T^{(n)}$ in [5], can be avoided with this synthetic materials basis.
- Orthonormalization: \tilde{M} is obtained either using the Gram-Schmidt algorithm on M , or diagonalizing $M^T M$ (as

proposed in [5]). This will generate negative attenuation coefficients (for the dot products between columns to be zero).

- Fessler’s method:

$$Q = \frac{SM}{S\mathbb{1}} \quad \text{i.e.} \quad (Q)_{b,m} = \frac{\sum_E S_{b,E} \mu_m(E)}{\sum_E S_{b,E}} \quad (5)$$

with $S_{b,E}$ the effective spectrum at energy E in bin b , and $\mu_m(E)$ the attenuation of material m at energy E . Each value $(Q)_{b,m}$ is a weighted mean of the attenuation coefficients of material m over the whole energy spectrum, the weights being the number of photons seen by the detector in bin b . Roughly, it evaluates how much the material m attenuates the photons that will end up being detected in bin b . Now, the problem of determining which linear combination of materials causes the observed attenuation would be much easier if each material attenuated the photons of only one bin, i.e. if Q was diagonal. Therefore, we define $\tilde{M} = MQ^{-1}$, i.e. $P = Q^{-1}$. When Q is not square, e.g. when there are more bins than materials, Q cannot be inverted, but P can be defined as the Moore-Penrose pseudo inverse of Q , i.e. $P = (Q^T Q)^{-1} Q^T$. This approach was proposed in [24] for dual energy CT, but its extension to more materials and more bins is straightforward.

When the algorithm is fed a matrix $\tilde{M} = MP$ of the attenuations of synthetic materials, it reconstructs synthetic material volumes \tilde{x} . No change is required for the data-attachment term, but the regularization must be performed on the original materials volumes $x = P\tilde{x}$, since we want the original material volumes to be regular, not the synthetic ones.

For Cai2013, the regularization term J_2 , its gradient g_2 and its Hessian H_2 must be replaced by their counterparts adapted to synthetic materials \tilde{J}_2 , \tilde{g}_2 and \tilde{H}_2 defined as follows:

$$\begin{aligned} \tilde{J}_2(\tilde{x}) &= J_2(P\tilde{x}) = 2\Phi(Dx, \theta_2) \\ \tilde{g}_2(\tilde{x}) &= P^T g_2(P\tilde{x}) = 2P^T D^T \Phi'(Dx, \theta_2) \\ \tilde{H}_2(\tilde{x}) &= P^T H_2(P\tilde{x}) P = 2P^T D^T \Phi''(Dx, \theta_2) DP \end{aligned}$$

The notations are those of the original paper [1], except for the matrix P , which is the one defined above. Similar changes are required for Long2014, Weidinger2016 and Mechlem2017: the gradient of the regularization term must be left-multiplied by P^T , and its Hessian left-multiplied by P^T and right-multiplied by P . For Barber2016, the required changes are described in the original paper.

C. Simulation setup

We have designed a simple two-dimensional 3-materials phantom (see Figure 1), consisting of a large square of water at 1 g/ml, a small square of iodine at 10 mg/ml, and a small square of gadolinium at 10 mg/ml. The iodine and gadolinium squares are inside the water square, but do not overlap. The phantom has 256^2 voxels. Through this phantom, 725 parallel projections were simulated, with 362 rays per projection, using the sparse forward projection matrix generated by the AIR toolbox [28]. The line integrals obtained were then converted to photon counts, following the classical polychromatic Beer-Lambert attenuation law. For the incident spectrum, we used the

same one as in Ducros *et al.* 2017 [13] which, quoting the paper, “was simulated using the SpekCalc software [29], considering a tube voltage of 120 kV, a 12° anode angle, and a filter of 1.2 mm of Al”. The detector response was simulated according to the model presented in appendix A.2 of [10]. In the end, the photon counts were corrupted with Poisson noise. This model therefore neglects pile-up, scatter, charge sharing and probably many other complex effects. For all methods, the initial guess is a set of zero-filled volumes.




Fig. 1: Ground truth of the phantom

D. Real data

A rabbit underwent an injection of both iodine and gadolinium in the kidney. The contrast agents were injected at different times, allowing to visualize a cortical enhancement with the iodinated contrast agent and a medullar phase with the gadolinated contrast agent. This injection protocol highlights the spectral capabilities of SPCCT for multiphase imaging within one single acquisition. The data was acquired on a Philips spectral photon counting CT prototype in Lyon. The scan consists of 2400 projections, each one of size 643×9 pixels, each pixel containing 5 energy bins. Only the central row of the projections was kept, and only one projection out of 5, ending up with a downsampled dataset of $643 \times 1 \times 480$ pixels. Each material volume was reconstructed as a 380×380 grid of 2D voxels. The detector response and incident spectrum descriptions were provided by the manufacturer of the spectral photon counting CT scanner prototype. The Philips prototype used a two-step projection-based reconstruction method: material decomposition is performed with an adaptation of Schlomka *et al.* [10], then filtered to remove aberrant pixels, and finally reconstructed by a standard filtered backprojection.

E. Comparison criteria

Comparison between iterative methods must be performed at convergence. How one determines that a method has converged, though, is a difficult question. Finding mathematical criteria that match with simple visual examination is a research topic in itself, far beyond the scope of the present paper. In one-step spectral CT, there are two potentially competing objectives: tomographic reconstruction, and material separation. It turns out that material separation is typically achieved much slower than tomographic reconstruction: with all studied methods, early iterates clearly show the shape of the imaged object, but exhibit important levels of cross-talk, which is unacceptable for clinical practice. Therefore, in the reconstruction of simulated data, tracking the concentration of a material over a region of interest (ROI) with known target concentration is a simple but efficient measure of convergence. We have run all methods with many

Method	δ_{contrast}	δ_{water}	Regularization contrast	Regularization water
Cai2013	0.001	0.1	100000 / 10000	30 / 10
Long2014	0.001	0.1	300000 / 10000	30 / 10
Weidinger2016	N/A	N/A	30000 / 3000	3 / 1
Barber2016	N/A	N/A	100 / 200	10000 / 20000
Mechlem2017	0.001	0.1	30000 / 10000	10 / 3

TABLE I: Regularization parameters of the five methods used to obtain the reconstructions on simulated data / real data. “Contrast” means iodine and gadolinium, which share the same set of parameters. For Barber2016, the “Regularization contrast” and “Regularization water” are actually the TV limits for contrast and water. The δ parameter is used in the Huber function or its equivalent, in the regularization terms of the methods.

iterations (the exact number varies with the method), and for each method, we provide:

- the concentrations of iodine and gadolinium in their respective ROIs over the course of the iterations
- the time per iteration with our implementation. Measurements were performed on a machine equipped with an Intel Xeon E5-2620 v4 CPU running at 2.10GHz and an NVidia GeForce GTX Titan X, and running OpenSuse Leap 42.2 and Matlab R2017b
- the number of iterations required for all materials to reach their target concentrations in their respective ROIs, within a certain tolerance
- the mean and standard deviation in each ROI at the last iteration, only to prove that the reconstructed volumes are reasonably close to the solution
- the theoretical memory footprint, since the actual memory footprint depends too much on the implementation

Since storing all iterates is impractical, for all methods except Mechlem2017 we stored one iterate every ten, which is why Table III contains mostly multiples of ten.

F. Choice of the regularization parameters

Automatically finding an appropriate regularization parameter is also a research topic in itself, and not in the scope of the present paper. However, the problem gets even more acute in spectral CT, since regularization typically creates cross-talk artifacts: most regularizations blur sharp borders, which causes incorrect concentration values around the borders of an object. This in turn causes a mismatch between the measured photon counts and those simulated through the reconstructed volume, and the algorithm fixes it by introducing some amount (positive or negative) of another material to obtain a better match. In this study, we have determined the regularization parameters empirically, choosing the largest possible regularization that did not cause significant amounts of cross-talk. This constrains the regularization parameters to remain small, which is why the results presented do not look patchy, and the noise level looks similar between analytical and iterative reconstruction. Table I summarizes the regularization parameters chosen for simulated data and real data, separated by a ‘/’.

G. Software

For this study, we re-implemented all of the five methods compared here in a lightweight Matlab framework. That code is free and open source, available at <https://github.com/cyrlmory/>

OneStepSpectralCT, and we appreciate any relevant contribution to it. In particular, we encourage the authors of the five methods we have studied to check that code against their own implementation. Together with the code, we provide detailed notes on how to re-do the calculations of some of these methods. In some of the papers ([1, 3, 5]), while re-doing the calculations, we found mistakes in one or more equations. There is no doubt that the authors of the papers have actually implemented and tested the algorithms they present with software that performs the right calculations, and that these mistakes are only “typos”. We provide the corrected equations in our notes.

III. RESULTS

A. Simulated data

Figure 2 shows the last iterate computed for each of the methods presented here. For Cai2013, we used Fessler’s method for μ -preconditioning. A video of the iterates, on which it is easy to evaluate the relative convergence rates of the various methods, is available as supplementary material. It has the same grayscale and is arranged the same way as Figure 2. The remaining cross-talk between gadolinium and water in the reconstruction by Barber2016 is most probably due to the fact that the algorithm has not converged: on very small-scale experiments (32x32 pixels), we were able to push the iterations further, and Barber2016 did yield results free from cross-talk.

B. Qualitative criteria and various comments

In addition to the quantitative criteria, we mention other important aspects of the algorithms: mathematical rigour, stability issues and need for additional parameters. Though we admit it is highly subjective, we also describe the “difficulty” of each method, i.e. how much mathematics one has to master to be able to follow the calculations (and ideally, re-do them), and how much time it took to implement them. All methods require parameters for regularization, but some need additional ones for other purposes: Cai2013 needs a parameter k_d characterizing the detector, and a threshold on how many times one tries to halve the step size before resetting the descent direction (see section II-A1). Barber2016 has a parameter $\lambda = 10^{-4}$, and a parameter $\theta = 0.5$ (see section II-A5). Long2014 splits the projections into 20 subsets, and Mechlem2017 into 4 subsets, in order to speed up convergence. Mechlem2017 also uses Nesterov’s momentum technique to speed it up even more. This does reduce the number of iterations required to obtain a sensible result by one to two orders of magnitude with respect to Weidinger2016, but makes Mechlem2017 unstable: when splitting the projections into too many subsets, Mechlem2017 sometimes diverges: the iterates accumulate strong streak artefacts, then they reach a point where the Hessian of the SQS is not invertible, which stops the algorithm.

In terms of mathematical difficulty, Cai2013 is probably the simplest method. It only requires to derive the gradient and Hessian of the cost function, and use them in a non-linear conjugate gradient framework. Weidinger2016 follows closely: the calculations only require a basic understanding of optimization transfer principles (see chapter 1 of [30]), and a few simple tricks based on the convexity inequality. The main theoretical difficulty (determining the optimal curvature for the

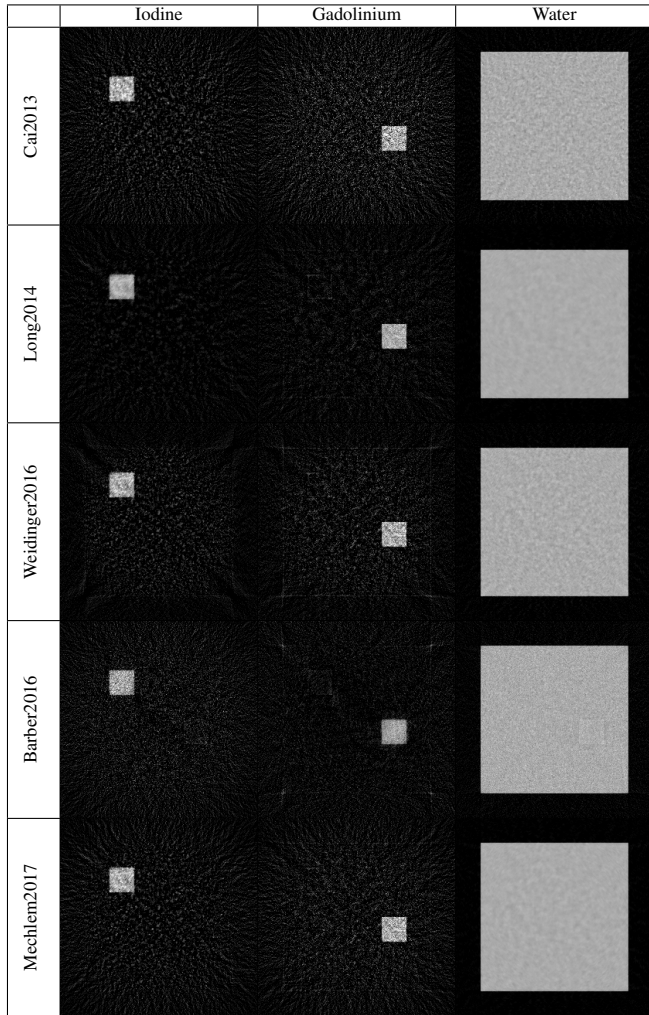


Fig. 2: Reconstructed volumes with the five methods, with grayscale ranges equal to $[0; 15]$ mg/ml for iodine and gadolinium, and $[0; 1.5]$ g/ml for water

quadratic surrogate) is avoided by using an approximation. Then comes Mechlem2017, which adds Nesterov's momentum technique and ordered subsets. Both can be implemented without fully understanding the mathematics from which they stem. They make the final implementation slightly more complex than that of Weidinger2016. Long2014 comes next, and solves the optimal curvature problem, with the assumption that the reconstructed volumes only contain positive attenuation values. With respect to the approximation used in Weidinger2016 and Mechlem2017, this rigorous solution requires additional calculations and intermediate variables, and its complete proof occupies several pages of the appendix. Lastly, Barber2016 is far more complex than the other four methods, and requires an in-depth understanding of proximal algorithms.

C. Quantitative criteria

We show in Figures 3 and 4 the evolution of the concentrations of iodine and gadolinium over the course of iterations, for all five methods, with a logarithmic scale only on the x-axis. The concentrations are measured in the square where iodine or gadolinium is expected, eroded by 2 pixels in order to avoid partial volume effects. In all methods, the pixels at the border of the region of interest (ROI) have significantly lower

concentrations. The full line is the ground truth concentration value in the ROI. Note that from these two graphs, Barber2016 seems to have reached convergence after 100 iterations, which is wrong: the concentration of water after 100 iterations is completely off target, and a quick look at the video of the iterates reveals that all three volumes are still contain a lot of cross talk. Table II contains the mean and standard deviation of




Fig. 3: Concentration of iodine over the course of the iterations for the five methods, in the ROI where it is expected eroded by 2 pixels

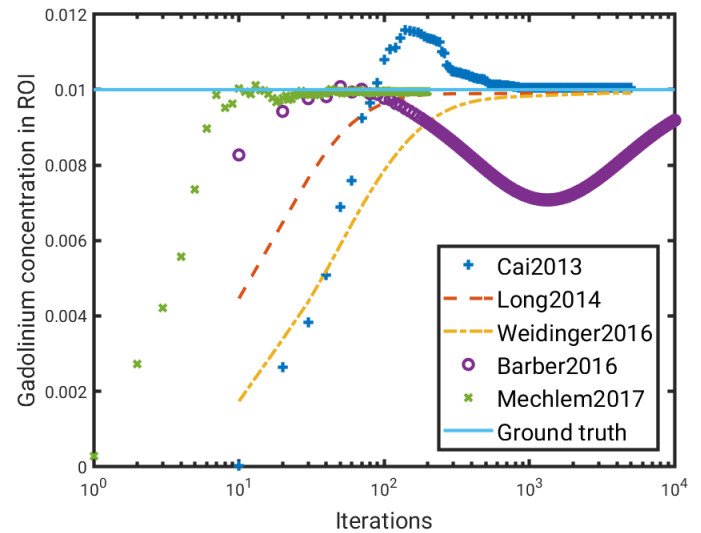


Fig. 4: Concentration of gadolinium over the course of the iterations for the five methods, in the ROI where it is expected eroded by 2 pixels

iodine, gadolinium and water concentrations in the ROI where they are expected eroded by 2 pixels, on the final iterate.

Table III shows the total number of iterations performed, the time per iteration, and the number of iterations it took for all materials to reach target concentration in their respective ROIs, within 20% or 10% tolerance. From these results, Mechlem2017 clearly appears as the best candidate for the processing of large real datasets: it converges 1 to 3 orders of magnitude faster than its competitors, and yields similar results.

	Iodine (mg/ml)		Gadolinium (mg/ml)		Water (g/ml)	
	mean	std	mean	std	mean	std
	Cai2013	10.4	2.14	10.1	3.75	0.999
Long2014	9.96	0.952	9.91	1.46	1	0.0191
Weidinger2016	9.92	1.86	9.92	2.44	1	0.0425
Barber2016	9.61	2.43	9.19	1.39	1	0.072
Mechlem2016	9.97	1.93	9.94	2.67	1	0.0227

TABLE II: Mean and standard deviation of iodine, gadolinium and water concentrations in the ROI where they are expected eroded by 2 pixels on the final iterate

Method	Iterations performed	Time per iter (s)	Within 20% tolerance	Within 10% tolerance
Cai2013	5000	5.2	270	430
Long2014	5000	1.33	150	290
Weidinger2016	5000	0.48	190	390
Barber2016	10000	1.02	3900	8360
Mechlem2017	200	0.62	5	10

TABLE III: Number of iterations performed and required to reach target concentrations within 20% and 10% tolerance.

D. Memory footprint of each method

We provide here a simple analysis of the memory footprint of each method, by listing the variables that have to be stored and recalling their size. Note that the implementation we provide keeps many intermediate variables in memory, for convenience and speed considerations, and therefore uses much more memory than the minimum required. The measured photon counts, which constitute the main input of these reconstruction algorithms, are assumed to be available on disk, so they never have to be fully loaded in memory and are not mentioned in this section. Obviously, re-reading them from disk at each iteration is highly inefficient, and one should load them in memory if possible. We use the notations of the original papers for the names of the variables, and the following notations for various quantities: N_v is the number of voxels of one reconstructed material volume, N_m the number of materials, N_p the number of pixels of all projections for one bin, N_b the number of energy bins in photon counts measurements. Despite the fact that we performed 2D simulations, we assume for these calculations that the volumes are 3D (which matters for the computation of the spatial gradient).

1) *Cai2013*: The following variables must be stored and updated at each iteration: $x^{(k)}$, $d^{(k)}$, $g^{(k)}$ and $g^{(k-1)}$, which each contain a set of material volumes. Since all computations can be performed pixel-wise, the expected photon counts \bar{y} do not have to be fully stored in memory, and neither does any of the intermediate variables in the calculation of the gradient and Hessian of the data-attachement term, so we do not count them. Lastly, the spatial gradient of $x^{(k)}$, which is 3 times as large as $x^{(k)}$, is used in the calculation of g and of $d^T Hd$. More precisely, it is the divergence of the first and second derivative of the Huber function of the gradient that is used. These calculations can be computed piecewise, with overlapping pieces since the gradient and divergence require neighboring voxels. In total, at least $4N_v N_m$ floating point values must be kept in memory, and storing $3N_v N_m$ additional values for the gradient of $x^{(k)}$ avoids redundant computations, and the problem of having to divide the volumes into overlapping pieces.

Variable	Stored ?	Size in memory (in floating point values)
f_0	yes	$N_v N_m$
$f^{(n-1)}$ $f^{(n)}$ $f^{(n+1)}$	only $f^{(n-1)}$	$N_v N_m$
Σ_{sino}	yes	$N_p N_b$
Σ_{grad} w	only one of both	$3N_v N_m$
$T^{(n)}$	yes: used only once, but computed as the sum of two very different calculations, unlikely to occur voxelwise in parallel	$N_v N_m$
$z_0^{(n+1)}$	no: used only once, pixelwise	0
$y_{\text{sino}}^{(n-1)}$ $y_{\text{sino}}^{(n)}$ $y_{\text{sino}}^{(n+1)}$ y_{sino}	not $y_{\text{sino}}^{(n+1)}$, because it replaces $y_{\text{sino}}^{(n)}$	$2N_p N_b$
$D_1(f_0)$ $b_1(f_0)$ $E_1(f_0)$	no: used only to compute $y_{\text{sino}}^{(n+1)}$, pixelwise	0
y_{grad}^+	yes	$3N_v N_m$
$y_{\text{grad}}^{(n)}$ $y_{\text{grad}}^{(n+1)}$ y_{grad}	only one of both	$3N_v N_m$
g^+	yes: to be projected onto L1 ball	$3N_v N_m$
\hat{g}^+	no: used only once, voxelwise	0
$f^{(n)}$ $f^{(n+1)}$	yes	$2N_v N_m$

TABLE IV: Memory footprint of Barber2016, for each variable used in appendix C.6 of [5]. N_v , N_m , N_p and N_b are defined in section III-D of the present paper. The words ‘‘pixelwise’’ and ‘‘voxelwise’’ mean that the variable can be computed one pixel (or voxel) at a time, thus its computation requires negligible amounts of memory

2) *Long2014*, *Weidinger2016* and *Mechlem2017*: In all three SQS-based methods we have implemented, the following variables must be stored and updated at each iteration: the current iterate ($x^{(n)}$ as in [2]), of size $N_v N_m$, and the gradient and Hessian of the surrogate at $x^{(n)}$, of size $N_v N_m$ and $N_v N_m^2$, respectively. The gradient and the Hessian of the surrogate of the data-attachment term both result from one forward and one back projection. By performing these one angle at a time, they can be computed with negligible amounts of memory. The gradient and Hessian of the surrogate of the regularization term both require a memory size of $N_v N_m$ (the Hessian has $N_v N_m^2$ elements, but is known to be diagonal), and internal computations can be done in place. Since typically they are calculated separately, and then added to those of the regularization terms, their requirements in memory add up. In total, the memory footprint is at least $(4 + N_m)N_v N_m$ floating point values for Long2014 and Weidinger2016. For Mechlem2017, two additional Nesterov variables must be stored ($z^{(k)}$ and $v^{(k)}$), taking the memory footprint to a minimum of $(6 + N_m)N_v N_m$ floating point values.

3) *Barber2016*: This method makes use of more intermediate variables than the other four. Table IV lists the variables used in the pseudo code given in appendix C.6 of [5], whether they must be stored or not, why, and how many floating point values are therefore required. In total, it uses at least $17N_v N_m + 3N_p N_b$ floating point values.

E. Real data

Figure 5 shows reconstructions of the rabbit acquisition described in section II-D. For each method presented here, it shows the last iterate computed (the number of iterations is the same as for the simulated data). For Cai2013, we used Fessler’s μ -preconditioning. A video of the iterates is available as supplementary material. The circular artefacts are due to imperfect calibration of the detector’s response and/or physical phenomena that are not taken into account in the forward model. Again, cross-talk can be observed in the reconstruction by Barber2016. It is unclear whether it is due to the fact that the algorithm has not converged, or to excessive regularization because of a too restrictive TV constraint. For Cai2013, we set the parameter k_d to the same value as in the simulations, which is probably not optimal, and may explain why Cai2013 results are noisier than the others. It illustrates why requiring extra parameters is a serious drawback for a reconstruction method. The spatial distribution of the gadolinium is different in the prototype’s reconstruction, where the gadolinium seems mainly located in a small region of the kidney, and in the one step reconstructions, where also the cortex of the kidney seems to contain gadolinium. Without a ground truth, it is difficult to determine which distribution is the correct one, if any.

F. Impact of mu-preconditioning

Surprisingly, none of the SQS-based methods is impacted by μ -preconditioning: once mapped back to the original materials, the iterates are almost identical to the ones obtained without μ -preconditioning, with difference maps two orders of magnitude less intense than the reconstructed volumes. The only impact is with Fessler’s method [24]: since in our case (three materials, five energy bins) it artificially increases the number of materials to that of the number of bins, the Hessians of the SQS are no longer invertible, so this μ -preconditioning cannot be used at all with SQS-based methods. On Cai2013, however, μ -preconditioning has a strong impact. Figure 6 shows the evolution of the cost function over the course of the iterations minus its minimum, using Cai2013 without μ -preconditioning, with normalization, orthonormalization and Fessler’s method. Since all four of these experiments aim to minimize the same cost function, the minimum we subtracted is the minimum over all iterations and all four experiments (reached, as can be seen on the graph, with μ -preconditioning by normalization). It clearly shows that Fessler’s method has the fastest initial convergence, which is confirmed by looking at the iterates (not shown here). After more than 1000 iterations, the experiment using normalization reached a lower value of the cost function than the one using Fessler’s method.

IV. DISCUSSION

We have presented a comparison of five one-step inversion methods for spectral CT, focused on aspects important for implementation into a real scanner: convergence speed, memory footprint and required mathematical background for the programmer. This choice of criteria is in part due to the initial aim of our study, which was to determine which algorithm should be implemented to process large real datasets in reasonable time. But more importantly, it was dictated by the behavior of the




Fig. 5: Reconstructions of a rabbit injected with iodine and gadolinium into the kidney, with the five one-step methods presented, and reconstruction from the Philips prototype. Grayscale ranges are [0;20] mg/ml for iodine, [0;7] mg/ml for gadolinium and [0;1.5] g/ml for water.

algorithms themselves: while all methods seemed to perform similarly in terms of image quality, provided one takes the time to adjust the regularization parameters, they differed a lot by practical aspects like convergence speed and ease of implementation. We provide the complete Matlab code we have used for this study, so that readers interested in conducting an image quality comparison study have the means to do so.

There are other one-step inversion methods in the literature than the ones we selected, e.g. [18, 20], and it is likely that more will be proposed in the future. We provide our code (see section II-G), hoping that others will be interested in implementing their method in the same framework, to extend the comparison.

A limitation of the present study is the small size of the data used in the experiments, which was dictated by our implementation using a sparse system matrix stored in computer memory. It is therefore possible that on larger datasets, e.g. on




Fig. 6: Cost function of Cai2013 over the course of the iterations minus its minimum, depending on the μ -preconditioning method used

a full real helical acquisition, the relative convergence speeds of the studied methods would be different. However, the other observations made in the paper would remain valid.

Another limitation is that since one-step inversion is a non-convex problem, the results may strongly depend on the initialization. Mechlem *et al.* recommend to initialize the algorithms with reconstruction results from a two-step method. We have not tried that approach, and have instead initialized all reconstructions with zero-filled volumes. Different initialization strategies may lead to different conclusions on relative convergence speeds and on image quality.

Since the same sparse system matrix was used both for simulation and reconstruction, volumes reconstructed from simulated data result from an inverse-crime setup [31]. However, this holds for all methods, so the comparison remains valid. On the real data, no inverse-crime is possible.

Sidky *et al.* have recently proposed a more efficient μ -preconditioning method [32], better-suited to three-materials reconstructions than the one described in Barber2016. In the experiment they present, this new μ -preconditioning method speeds up convergence by a factor 10. The slow convergence of Barber2016 in our experiments must therefore be mitigated by the fact that we did not implement that method. However, a factor 10 would only make its convergence speed similar to Cai2013, Long2014 and Weidinger2016, and our other observations on Barber2016 (memory footprint, mathematical complexity) remain valid.

Our implementation of Mechlem2017 differs from that of Weidinger2016 only by the use of Nesterov’s momentum technique and a small number of ordered subsets, and yet greatly outperforms Weidinger2016. Therefore we confirm the observation made in the discussion of Mechlem2017, and in several other papers [33, 34]: Nesterov’s momentum technique combined with ordered subsets highly accelerates the convergence speed per iteration. We have also tried modifying Cai2013 to include Nesterov’s momentum technique (it is available in the code we provide), and it does accelerate convergence there too.

The five methods we have compared all use the same forward

model, which does not simulate advanced effects like charge sharing or pulse pile-up. A way to improve the reconstruction results of one-step methods on real data would be to take these effects into account in the forward model. How to minimize a cost function with such a modified forward model is a topic for future research on one-step methods.

V. CONCLUSION

We have compared five one-step reconstruction methods: Cai2013 [1], Long2014 [2], Weidinger2016 [3], Barber2016 [5], Mechlem2017 [4], leaving aside image quality considerations to focus on convergence speed and other practical aspects. From these observations, it is clear that among the five methods studied, the only one suitable for implementation in a real spectral CT scanner is Mechlem2017.

ACKNOWLEDGEMENTS

This work was performed within the framework of the EU’s H2020 research and innovation program under the grant agreement No. 633937, the SIRIC LYric Grant INCa-DGOS-4664, and the LABEX PRIMES (ANR-11-LABX-0063) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-0007) operated by the ANR.

REFERENCES

- [1] C. Cai, T. Rodet, S. Legoupil, and A. Mohammad-Djafari, “A full-spectral Bayesian reconstruction approach based on the material decomposition model applied in dual-energy computed tomography,” *Medical Physics*, vol. 40, pp. 111916–111931, Nov. 2013.
- [2] Y. Long and J. A. Fessler, “Multi-Material Decomposition Using Statistical Image Reconstruction for Spectral CT,” *IEEE Transactions on Medical Imaging*, vol. 33, pp. 1614–1626, Aug. 2014.
- [3] 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.
- [4] K. Mechlem, S. Ehn, T. SELLERER, E. Braig, D. Munzel, F. Pfeiffer, and P. B. Noel, “Joint statistical iterative material image reconstruction for spectral computed tomography using a semi-empirical forward model,” *IEEE Transactions on Medical Imaging*, pp. 1–1, 2017.
- [5] R. F. Barber, E. Y. Sidky, T. G. Schmidt, and X. Pan, “An algorithm for constrained one-step inversion of spectral CT data,” *Physics in Medicine and Biology*, vol. 61, no. 10, p. 3784, 2016.
- [6] R. E. Alvarez and A. Macovski, “Energy-selective reconstructions in X-ray computerized tomography,” *Physics in Medicine and Biology*, vol. 21, pp. 733–744, Sept. 1976.
- [7] D. Pan, E. Roessl, J.-P. Schlomka, S. D. Caruthers, A. Sempin, M. J. Scott, J. S. Allen, H. Zhang, G. Hu, P. J. Gaffney, E. T. Choi, V. Rasche, S. A. Wickline, R. Proksa, and G. M. Lanza, “Computed Tomography in Color: NanoK-Enhanced Spectral CT Molecular Imaging,” *Angewandte Chemie*, vol. 122, pp. 9829–9833, Dec. 2010.

- [8] C. Maass, M. Baer, and M. Kachelriess, "Image-based dual energy CT using optimized precorrection functions: A practical new approach of material decomposition in image domain," *Medical Physics*, vol. 36, pp. 3818–3829, Aug. 2009.
- [9] T. Niu, X. Dong, M. Petrongolo, and L. Zhu, "Iterative image-domain decomposition for dual-energy CT," *Medical Physics*, vol. 41, pp. n/a–n/a, Apr. 2014.
- [10] J. P. Schlomka, E. Roessl, R. Dorscheid, S. Dill, G. Martens, T. Stel, C. Bäumer, C. Herrmann, R. Steadman, G. Zeitler, A. Livne, and R. Proksa, "Experimental feasibility of multi-energy photon-counting K-edge imaging in pre-clinical computed tomography," *Physics in Medicine and Biology*, vol. 53, no. 15, p. 4031, 2008.
- [11] A. Sawatzky, Q. Xu, C. O. Schirra, and M. A. Anastasio, "Proximal ADMM for Multi-Channel Image Reconstruction in Spectral X-ray CT," *IEEE Transactions on Medical Imaging*, vol. 33, pp. 1657–1668, Aug. 2014.
- [12] B. Brendel, F. Bergner, K. Brown, and T. Koehler, "Penalized Likelihood Decomposition for Dual Layer Spectral CT," in *Proceeding of the Fourth International Meeting in X-Ray Computed Tomography*, (Bamberg, Germany), pp. 41–44, July 2016.
- [13] N. Ducros, J. F. P. J. Abascal, B. Sixou, S. Rit, and F. Peyrin, "Regularization of Nonlinear Decomposition of Spectral X-ray Projection Images," *Medical Physics*, 2017.
- [14] R. Clackdoyle, L. Desbat, J. Lesaint, and S. Rit, "Data Consistency Conditions for Cone-Beam Projections on a Circular Trajectory," *IEEE Signal Processing Letters*, vol. 23, pp. 1746–1750, Dec. 2016.
- [15] J. Lesaint, S. Rit, R. Clackdoyle, and L. Desbat, "Calibration for Circular Cone-Beam CT Based on Consistency Conditions," *IEEE Transactions on Radiation and Plasma Medical Sciences*, vol. 1, no. 6, pp. 517–526, 2017.
- [16] A. So and T.-Y. Lee, "Rapid kV Switching Dual-Energy CT Imaging," in *Dual-Energy CT in Cardiovascular Imaging*, pp. 45–60, Springer, Cham, 2015.
- [17] T. G. Flohr, C. H. McCollough, H. Bruder, M. Petersilka, K. Gruber, C. Süß, M. Grasruck, K. Stierstorfer, B. Krauss, R. Raupach, A. N. Primak, A. Küttner, S. Achenbach, C. Becker, A. Kopp, and B. M. Ohnesorge, "First performance evaluation of a dual-source CT (DSCT) system," *European Radiology*, vol. 16, pp. 256–268, Feb. 2006.
- [18] B. Chen, Z. Zhang, E. Y. Sidky, D. Xia, and X. Pan, "Image reconstruction and scan configurations enabled by optimization-based algorithms in multispectral CT," *Physics in Medicine & Biology*, vol. 62, no. 22, p. 8763, 2017.
- [19] X. Pan, E. Y. Sidky, and M. Vannier, "Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction?," *Inverse Problems*, vol. 25, p. 123009, Dec. 2009.
- [20] P. Sukovic and N. H. Clinthorne, "Penalized weighted least-squares image reconstruction for dual energy X-ray transmission tomography," *IEEE Transactions on Medical Imaging*, vol. 19, pp. 1075–1081, Nov. 2000.
- [21] N. Maass, S. Sawall, M. Knaup, and M. Kachelriess, "Empirical multiple energy calibration (EMEC) for material-selective CT," pp. 4222–4229, IEEE, Oct. 2011.
- [22] P. R. S. Mendonca, P. Lamb, and D. V. Sahani, "A Flexible Method for Multi-Material Decomposition of Dual-Energy CT Images," *IEEE transactions on medical imaging*, vol. 33, pp. 99–116, Jan. 2014.
- [23] C. H. McCollough, G. H. Chen, W. Kalender, S. Leng, E. Samei, K. Taguchi, G. Wang, L. Yu, and R. I. Pettigrew, "Achieving Routine Submillisievert CT Scanning: Report from the Summit on Management of Radiation Dose in CT," *Radiology*, vol. 264, pp. 567–580, Aug. 2012.
- [24] J. A. Fessler, "Method for statistically reconstructing images from a plurality of transmission measurements having energy diversity and image reconstructor apparatus utilizing the method," Aug. 2003. International Classification G06T5/00; Cooperative Classification G06T2207/10081, G06T5/002, G06T11/006, A61B6/482, A61B6/4241, A61B6/405, A61B6/032, A61B6/4035, G06T2211/408, G06T2211/424; European Classification A61B6/48D, G06T11/00T3, A61B6/03B, A61B6/40H, A61B6/40F, A61B6/42B8, G06T5/00D.
- [25] S. S. Keerthi and E. G. Gilbert, "Convergence of a Generalized SMO Algorithm for SVM Classifier Design," *Machine Learning*, vol. 46, pp. 351–360, Jan. 2002.
- [26] R. F. Barber and E. Y. Sidky, "MOCCA: Mirrored convex/concave optimization for nonconvex composite functions," *arXiv:1510.08842 [math]*, Oct. 2015.
- [27] A. Chambolle and T. Pock, "A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging," *Journal of Mathematical Imaging and Vision*, vol. 40, pp. 120–145, May 2011.
- [28] P. C. Hansen and M. Saxild-Hansen, "AIR Tools — A MATLAB package of algebraic iterative reconstruction methods," *Journal of Computational and Applied Mathematics*, vol. 236, pp. 2167–2178, Feb. 2012.
- [29] G. Poludniowski, G. Landry, F. DeBlois, P. M. Evans, and F. Verhaegen, "SpekCalc: A program to calculate photon spectra from tungsten anode x-ray tubes," *Physics in Medicine and Biology*, vol. 54, pp. N433–438, Oct. 2009.
- [30] J. Beutel, *Handbook of Medical Imaging: Medical Image Processing and Analysis*. SPIE Press, 2000.
- [31] A. Wirgin, "The inverse crime," *arXiv:math-ph/0401050*, Jan. 2004.
- [32] E. Y. Sidky, R. F. Barber, T. Gilat-Schmidt, and X. Pan, "Three material decomposition for spectral computed tomography enabled by block-diagonal step-preconditioning," *arXiv:1801.06263 [physics]*, Jan. 2018.
- [33] D. Kim, S. Ramani, and J. A. Fessler, "Combining Ordered Subsets and Momentum for Accelerated X-Ray CT Image Reconstruction," *IEEE Transactions on Medical Imaging*, vol. 34, pp. 167–178, Jan. 2015.
- [34] D. C. Hansen and T. S. Sørensen, "Fast 4D cone-beam CT from 60 s acquisitions," *Physics and Imaging in Radiation Oncology*, vol. 5, pp. 69–75, Jan. 2018.