Published in final edited form as:

IEEE Trans Med Imaging. 2015 February ; 34(2): 388–399. doi:10.1109/TMI.2014.2358499.

Fast X-Ray CT Image Reconstruction Using a Linearized Augmented Lagrangian Method with Ordered Subsets

Hung Nien [Student Member, IEEE] and Jeffrey A. Fessler [Fellow, IEEE]

Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48105, USA

Abstract

Augmented Lagrangian (AL) methods for solving convex optimization problems with linear constraints are attractive for imaging applications with composite cost functions due to the empirical fast convergence rate under weak conditions. However, for problems such as X-ray computed tomography (CT) image reconstruction, where the inner least-squares problem is challenging and requires iterations, AL methods can be slow. This paper focuses on solving regularized (weighted) least-squares problems using a linearized variant of AL methods that replaces the quadratic AL penalty term in the scaled augmented Lagrangian with its separable quadratic surrogate (SQS) function, leading to a simpler ordered-subsets (OS) accelerable splitting-based algorithm, OS-LALM. To further accelerate the proposed algorithm, we use a second-order recursive system analysis to design a deterministic downward continuation approach that avoids tedious parameter tuning and provides fast convergence. Experimental results show that the proposed algorithm significantly accelerates the convergence of X-ray CT image reconstruction with negligible overhead and can reduce OS artifacts when using many subsets.

Keywords

Statistical image reconstruction; computed tomography; ordered subsets; augmented Lagrangian

I. INTRODUCTION

STATISTICAL methods for image reconstruction have been explored extensively for computed tomography (CT) due to the potential of acquiring CT scans with lower X-ray dose while maintaining image quality. However, the much longer computation time of statistical methods still restrains their applicability in practice. To accelerate statistical methods, many optimization techniques have been investigated. Augmented Lagrangian (AL) methods (including the alternating direction variants) [1-4] are powerful techniques for solving regularized inverse problems using variable splitting. For example, in total-variation (TV) denoising and compressed sensing (CS) problems, AL methods can separate non-smooth l_1 regularization terms by introducing auxiliary variables, yielding simple penalized least-squares inner problems that are solved efficiently using the fast Fourier transform (FFT) algorithm and proximal mappings such as the soft-thresholding for the l_1 -norm [5, 6].

However, in applications like X-ray CT image reconstruction, the inner least-squares problem is challenging due to the highly shift-variant Hessian caused by the huge dynamic range of the statistical weighting. To solve this problem, Ramani *et al.* [7] introduced an additional variable that separates the shift-variant and approximately shift-invariant components of the statistically weighted quadratic data-fitting term, leading to a better-conditioned inner least-squares problem that was solved efficiently using the preconditioned conjugate gradient (PCG) method with an appropriate circulant preconditioner. Experimental results showed significant acceleration in 2D CT [7]; however, in 3D CT with cone-beam geometry, it is more difficult to construct a good preconditioner for the inner least-squares problem, and the method in [7] has yet to achieve the same acceleration as in 2D CT. Furthermore, even when a good preconditioner can be found, the iterative PCG solver requires several forward/back-projection operations per outer iteration, which is very time-consuming in 3D CT, significantly reducing the number of outer-loop image updates one can perform within a given reconstruction time.

The ordered-subsets (OS) algorithm [8] is a first-order method with a diagonal preconditioner that uses somewhat conservative step sizes but is easily applicable to 3D CT. By grouping the projections into M ordered subsets that satisfy the “subset balance condition” and updating the image incrementally using the M subset gradients, OS algorithms effectively perform M times as many image updates per outer iteration as the standard gradient descent method, leading to M times acceleration in early iterations. We can interpret the OS algorithm and its variants as incremental gradient methods [9]; when the subset is chosen randomly with some constraints so that the subset gradient is unbiased and with finite variance, they can also be called stochastic gradient methods [10]. Recently, OS variants [11, 12] of the fast gradient method [13–15] demonstrated dramatic acceleration (about M^2 times in early iterations) over their one-subset counterparts. However, when M increases, fast OS algorithms seem to have “larger” limit cycles and exhibit artifacts in the reconstructed images. This problem is also studied in the machine learning literature. Devolder showed that the error accumulation in fast gradient methods is inevitable when an inexact oracle is used, but it can be reduced by using relaxed momentum, i.e., a growing diagonal majorizer (or equivalently, a diminishing step size), at the cost of slower convergence rate [16]. Schmidt *et al.* also showed that an accelerated proximal gradient method is more sensitive to errors in the gradient and proximal mapping calculations [17].

OS-based algorithms, including the standard one and its fast variants, are not convergent in general (unless relaxation [18] or incremental majorization [19] is used, unsurprisingly, at the cost of slower convergence rate) and possibly introduce noise-like artifacts. Nevertheless, the effective M -times image updates using OS is still very promising for AL methods. As an example of combining OS (or stochastic gradients) with AL methods, Ouyang *et al.* [20] proposed a stochastic setting for the alternating direction method of multipliers (ADMM) [4, 6] that introduces an auxiliary variable for the regularization term and majorizes the smooth data-fitting term such as the logistic loss in the scaled augmented Lagrangian using a diagonal majorizer with stochastic gradients. For every stochastic ADMM iteration, only part of the data is visited (for evaluating the gradient of a subset of the data). This substantially reduces the cost per stochastic ADMM iteration, and one can

run more stochastic ADMM iterations in a given reconstruction time. However, stochastic ADMM simply combines the stochastic gradient method and ADMM, and it reverts to the stochastic gradient method when no variable splitting is considered. Therefore, the AL framework in stochastic ADMM extends the original stochastic gradient method so that it can use variable splitting for more complicated regularizations such as the non-smooth l_1 regularization, but it does not greatly accelerate convergence for problems with smooth regularizers (in which variable splitting is less compelling than in the non-smooth case) like those considered here for low-dose X-ray CT.

In this paper, we focus on solving regularized (weighted) least-squares problems using a linearized AL method (LALM) [21]. We majorize the quadratic AL penalty term, instead of the smooth data-fitting term, in the scaled augmented Lagrangian using a fixed diagonal majorizer, leading to a much simpler OS-accelerable splitting-based algorithm, OS-LALM. For further acceleration, we use a second-order recursive system analysis to design a deterministic downward continuation approach that avoids tedious parameter tuning and provides fast convergence. Experimental results show that the proposed algorithm significantly accelerates the convergence of X-ray CT image reconstruction in early iterations with negligible overhead and greatly reduces OS artifacts in the reconstructed image when using many subsets.

The paper is organized as follows. Section II reviews the linearized AL method in a general setting and shows new convergence properties of the linearized AL method with in-exact updates. Section III derives the proposed OS-accelerable splitting-based algorithm for solving regularized least-squares problems using the linearized AL method and develops a deterministic downward continuation approach for fast convergence without parameter tuning. Section IV considers solving X-ray CT image reconstruction problem with penalized weighted least-squares (PWLS) criterion using the proposed algorithm. Section V reports the experimental results of applying our proposed algorithm to X-ray CT image reconstruction. Further results are shown in the supplementary material. Finally, we draw conclusions in Section VI.

II. BACKGROUND

A. Linearized AL method (LALM)

Consider a general composite convex optimization problem:

$$\hat{\mathbf{x}} \in \underset{\mathbf{x}}{\operatorname{argmin}} \{g(\mathbf{Ax}) + h(\mathbf{x})\} \quad (1)$$

and its equivalent constrained minimization problem:

$$(\hat{\mathbf{x}}, \hat{\mathbf{u}}) \in \underset{\mathbf{x}, \mathbf{u}}{\operatorname{argmin}} \{g(\mathbf{u}) + h(\mathbf{x})\} \quad \text{s.t.} \quad \mathbf{u} = \mathbf{Ax}, \quad (2)$$

where both g and h are closed and proper convex functions. In CT, $\mathbf{A}$ denotes the system (projection) matrix, $\mathbf{x}$ denotes the image being reconstructed, g is a weighted quadratic data-fitting term, and h is an edge-preserving regularization term. One way to solve the

constrained minimization problem (2) is to use an (alternating direction) AL method that alternately minimizes the scaled augmented Lagrangian:

$$L_A(\mathbf{x}, \mathbf{u}, \mathbf{d}; \rho) \triangleq g(\mathbf{u}) + h(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{Ax} - \mathbf{u} - \mathbf{d}\|_2^2 \quad (3)$$

with respect to $\mathbf{x}$ and $\mathbf{u}$, followed by a gradient ascent of $\mathbf{d}$, yielding the following AL iterates [4, 6]:

$$\begin{cases} \mathbf{x}^{(k+1)} \in \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ h(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{Ax} - \mathbf{u}^{(k)} - \mathbf{d}^{(k)}\|_2^2 \right\} \\ \mathbf{u}^{(k+1)} \in \underset{\mathbf{u}}{\operatorname{argmin}} \left\{ g(\mathbf{u}) + \frac{\rho}{2} \|\mathbf{Ax}^{(k+1)} - \mathbf{u} - \mathbf{d}^{(k)}\|_2^2 \right\} \\ \mathbf{d}^{(k+1)} = \mathbf{d}^{(k)} - \mathbf{Ax}^{(k+1)} + \mathbf{u}^{(k+1)}, \end{cases} \quad (4)$$

where $\mathbf{d}$ is the scaled Lagrange multiplier of the split variable $\mathbf{u}$, and $\rho > 0$ is the corresponding AL penalty parameter.

In the linearized AL method (LALM) [21] (also known as the split inexact Uzawa method [22-24]), one replaces the quadratic AL penalty term in the $\mathbf{x}$ -update of (4):

$$\theta_k(\mathbf{x}) \triangleq \frac{\rho}{2} \|\mathbf{Ax} - \mathbf{u}^{(k)} - \mathbf{d}^{(k)}\|_2^2 \quad (5)$$

by its separable quadratic surrogate (SQS) function:

$$\begin{aligned} & \tilde{\theta}_k(\mathbf{x}; \mathbf{x}^{(k)}) \\ & \triangleq \theta_k(\mathbf{x}^{(k)}) + \langle \nabla \theta_k(\mathbf{x}^{(k)}), \mathbf{x} - \mathbf{x}^{(k)} \rangle + \frac{\rho L}{2} \|\mathbf{x} - \mathbf{x}^{(k)}\|_2^2 \\ & = \frac{\rho}{2t} \|\mathbf{x} - (\mathbf{x}^{(k)} - t\mathbf{A}'(\mathbf{Ax}^{(k)} - \mathbf{u}^{(k)} - \mathbf{d}^{(k)}))\|_2^2 + (\text{constant independent of } \mathbf{x}), \end{aligned} \quad (6)$$

where $L > \|\mathbf{A}\|_2^2 = \lambda_{\max}(\mathbf{A}'\mathbf{A})$ ensures $L\mathbf{I} - \mathbf{A}'\mathbf{A} > 0$, and $t \triangleq 1/L$. This function satisfies the “majorization” condition:

$$\begin{cases} \tilde{\theta}_k(\mathbf{x}; \bar{\mathbf{x}}) \geq \theta_k(\mathbf{x}), & \forall \mathbf{x}, \bar{\mathbf{x}} \in \operatorname{Dom} \theta_k \\ \tilde{\theta}_k(\bar{\mathbf{x}}; \bar{\mathbf{x}}) = \theta_k(\bar{\mathbf{x}}), & \forall \bar{\mathbf{x}} \in \operatorname{Dom} \theta_k. \end{cases} \quad (7)$$

It is trivial to generalize L to a symmetric positive semi-definite matrix $\mathbf{S}_+$, e.g., the diagonal matrix $\mathbf{D}_L$ used in OS-based algorithms [8, 25], and still ensure (7). When $\mathbf{S}_+ = \mathbf{A}'\mathbf{A}$, LALM just reverts to the standard AL method. Majorizing with a diagonal matrix leads to a simpler $\mathbf{x}$ -update. The corresponding LALM iterates are as follows [21]:

$$\begin{cases} \mathbf{x}^{(k+1)} \in \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ \phi_k(\mathbf{x}) \triangleq h(\mathbf{x}) + \check{\theta}_k(\mathbf{x}; \mathbf{x}^{(k)}) \right\} \\ \mathbf{u}^{(k+1)} \in \underset{\mathbf{u}}{\operatorname{argmin}} \left\{ g(\mathbf{u}) + \frac{\rho}{2} \|\mathbf{A}\mathbf{x}^{(k+1)} - \mathbf{u} - \mathbf{d}^{(k)}\|_2^2 \right\} \\ \mathbf{d}^{(k+1)} = \mathbf{d}^{(k)} - \mathbf{A}\mathbf{x}^{(k+1)} + \mathbf{u}^{(k+1)}. \end{cases} \quad (8)$$

The $\mathbf{x}$ -update can be written as the proximal mapping of h :

$$\begin{aligned} \mathbf{x}^{(k+1)} &\in \operatorname{prox}_{(\rho^{-1}t)h} \left(\mathbf{x}^{(k)} - t\mathbf{A}' \left(\mathbf{A}\mathbf{x}^{(k)} - \mathbf{u}^{(k)} - \mathbf{d}^{(k)} \right) \right) \\ &= \operatorname{prox}_{(\rho^{-1}t)h} \left(\mathbf{x}^{(k)} - (\rho^{-1}t) \mathbf{s}^{(k+1)} \right), \end{aligned} \quad (9)$$

where prox_f denotes the proximal mapping of f defined as:

$$\operatorname{prox}_f(\mathbf{z}) \triangleq \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ f(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 \right\}, \quad (10)$$

and

$$\mathbf{s}^{(k+1)} \triangleq \rho \mathbf{A}' \left(\mathbf{A}\mathbf{x}^{(k)} - \mathbf{u}^{(k)} - \mathbf{d}^{(k)} \right) \quad (11)$$

denotes the “search direction” of the proximal gradient-like $\mathbf{x}$ -update that consists of a descent step and a proximal mapping step using the same step size, for instance, $\rho^{-1}t$ in (9). Furthermore, θ_k can be written as:

$$\check{\theta}_k(\mathbf{x}; \mathbf{x}^{(k)}) = \theta_k(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{x}^{(k)}\|_G^2, \quad (12)$$

where $\mathbf{G} \triangleq L\mathbf{I} - \mathbf{A}'\mathbf{A} \succ 0$ by the definition of L . Hence, the LALM iterates (8) can be represented as a proximal-point variant [24] of the standard AL iterates (4) (also known as the preconditioned ADMM iterates [26] discussed later) by plugging (2) into (8):

$$\begin{cases} \mathbf{x}^{(k+1)} \in \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ h(\mathbf{x}) + \theta_k(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{x} - \mathbf{x}^{(k)}\|_G^2 \right\} \\ \mathbf{u}^{(k+1)} \in \underset{\mathbf{u}}{\operatorname{argmin}} \left\{ g(\mathbf{u}) + \frac{\rho}{2} \|\mathbf{A}\mathbf{x}^{(k+1)} - \mathbf{u} - \mathbf{d}^{(k)}\|_2^2 \right\} \\ \mathbf{d}^{(k+1)} = \mathbf{d}^{(k)} - \mathbf{A}\mathbf{x}^{(k+1)} + \mathbf{u}^{(k+1)}. \end{cases} \quad (13)$$

B. Convergence properties with inexact updates

The LALM iteration (8) is convergent for any fixed AL penalty parameter $\rho > 0$ and any $\mathbf{A}$ [21], while the standard AL method is convergent (in the primal) if $\mathbf{A}$ has full column rank [4, Theorem 8]. Furthermore, even if the AL penalty parameter varies every iteration, (8) is convergent when ρ is non-decreasing and bounded above [21]. However, existing convergence analyses of LALM assume that all updates are exact. In this paper, since some updates might not be solved exactly, we must consider the LALM iteration with inexact

updates. Specifically, instead of the exact LALM in (8), we focus on two closely related inexact LALM variants:

$$\begin{cases} \| \mathbf{x}^{(k+1)} - \underset{\mathbf{x}}{\operatorname{argmin}} \phi_k(\mathbf{x}) \| \leq \delta_k \\ \mathbf{u}^{(k+1)} \in \underset{\mathbf{u}}{\operatorname{argmin}} \left\{ g(\mathbf{u}) + \frac{\rho}{2} \| \mathbf{A}\mathbf{x}^{(k+1)} - \mathbf{u} - \mathbf{d}^{(k)} \|_2^2 \right\} \\ \mathbf{d}^{(k+1)} = \mathbf{d}^{(k)} - \mathbf{A}\mathbf{x}^{(k+1)} + \mathbf{u}^{(k+1)}, \end{cases} \quad (14)$$

where ϕ_k was defined in (8), and

$$\begin{cases} |\phi(\mathbf{x}^{(k+1)}) - \min_{\mathbf{x}} \phi_k(\mathbf{x})| \leq \epsilon_k \\ \mathbf{u}^{(k+1)} \in \underset{\mathbf{u}}{\operatorname{argmin}} \left\{ g(\mathbf{u}) + \frac{\rho}{2} \| \mathbf{A}\mathbf{x}^{(k+1)} - \mathbf{u} - \mathbf{d}^{(k)} \|_2^2 \right\} \\ \mathbf{d}^{(k+1)} = \mathbf{d}^{(k)} - \mathbf{A}\mathbf{x}^{(k+1)} + \mathbf{u}^{(k+1)}. \end{cases} \quad (15)$$

These inexact variants of LALM revert to the standard LALM when $\delta_k = 0$ and $\epsilon_k = 0$. The $\mathbf{u}$ -update could also be inexact; however, for simplicity, we focus on exact updates of $\mathbf{u}$. Considering inexact updates of $\mathbf{u}$ is a trivial extension.

Our convergence analysis of the inexact LALM is twofold. First, we show that the equivalent proximal-point variant of the standard AL iterates (3) can be interpreted as a convergent ADMM that solves another equivalent constrained minimization problem of the form (1) with a redundant split (the proof is in the supplementary material):

$$(\hat{\mathbf{x}}, \hat{\mathbf{u}}, \hat{\mathbf{v}}) \in \underset{\mathbf{x}, \mathbf{u}, \mathbf{v}}{\operatorname{argmin}} \{g(\mathbf{u}) + h(\mathbf{x})\} \quad \text{s.t. } \mathbf{u} = \mathbf{A}\mathbf{x} \quad \text{and} \quad \mathbf{v} = \mathbf{G}^{1/2}\mathbf{x}. \quad (16)$$

Therefore, the LALM iteration (8) is a convergent ADMM, and it inherits the nice properties of ADMM, including the tolerance of inexact updates [4, Theorem 8]. More formally, we have the following theorem:

Theorem 1. Consider a constrained composite convex optimization problem (2) where both g and h are closed and proper convex functions. Let $\rho > 0$ and $\{\delta_k\}_{k=0}^{\infty}$ denote a non-negative sequence such that

$$\sum_{k=0}^{\infty} \delta_k < \infty. \quad (17)$$

If (2) has a solution $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$, then the sequence of updates $\left\{ \left(\mathbf{x}^{(k)}, \mathbf{u}^{(k)} \right) \right\}_{k=0}^{\infty}$ generated by the inexact LALM in (4) converges to $(\hat{\mathbf{x}}, \hat{\mathbf{u}})$; otherwise, at least one of the sequences

$$\left\{ \left(\mathbf{x}^{(k)}, \mathbf{u}^{(k)} \right) \right\}_{k=0}^{\infty} \text{ or } \left\{ \mathbf{d}^{(k)} \right\}_{k=0}^{\infty} \text{ diverges.}$$

Theorem 1 shows that the inexact LALM in (14) converges if the error δ_k is absolutely summable. However, it does not describe how fast the iterates converge and more

importantly, how inexact updates affect the convergence rate. This leads to the second part of our convergence analysis.

In this part, we rely on the equivalence between LALM and the Chambolle-Pock first-order primal-dual algorithm (CP) [26]. Consider a minimax problem:

$$(\hat{\mathbf{z}}, \hat{\mathbf{x}}) \in \underset{\mathbf{z}}{\operatorname{argmin}} \underset{\mathbf{x}}{\operatorname{max}} \Omega(\mathbf{z}, \mathbf{x}), \quad (18)$$

where

$$\Omega(\mathbf{z}, \mathbf{x}) \triangleq \langle -\mathbf{A}' \mathbf{z}, \mathbf{x} \rangle + g^*(\mathbf{z}) - h(\mathbf{x}), \quad (19)$$

and f^* denotes the convex conjugate of a function f [27, p. 104]. Note that since both g and h are closed, proper, and convex, it follows $g^{**} = g$ and $h^{**} = h$. The sequence of updates

$\{(\mathbf{z}^{(k)}, \mathbf{x}^{(k)})\}_{k=0}^{\infty}$ generated by the CP iterates:

$$\begin{cases} \mathbf{x}^{(k+1)} \in \operatorname{prox}_{\sigma h} \left(\mathbf{x}^{(k)} - \sigma \mathbf{A}' \mathbf{z}^{(k)} \right) \\ \mathbf{z}^{(k+1)} \in \operatorname{prox}_{\tau g^*} \left(\mathbf{z}^{(k)} + \tau \mathbf{A} \mathbf{x}^{(k+1)} \right) \\ \mathbf{z}^{(k+1)} = \mathbf{z}^{(k+1)} + \left(\mathbf{z}^{(k+1)} - \mathbf{z}^{(k)} \right) \end{cases} \quad (20)$$

converges to a saddle-point $(\hat{\mathbf{z}}, \hat{\mathbf{x}})$ of (18), and the non-negative primal-dual gap $\Omega(\hat{\mathbf{z}}, \hat{\mathbf{x}}) - \Omega(\hat{\mathbf{z}}, \mathbf{x}_k)$ converges to zero with rate $O(1/k)$ [26, Theorem 1] provided that $\sigma\tau\|\mathbf{A}\|_2^2 < 1$ where $\mathbf{z}_k \triangleq \frac{1}{k} \sum_{j=1}^k \mathbf{z}^{(j)}$ and $\mathbf{x}_k \triangleq \frac{1}{k} \sum_{j=1}^k \mathbf{x}^{(j)}$ denote the arithmetic means of all previous $\mathbf{z}$ - and $\mathbf{x}$ -iterates. Since the CP iterates (20) solve the maximum problem (18), they also solve the primal problem:

$$\hat{\mathbf{z}} \in \underset{\mathbf{z}}{\operatorname{argmin}} \left\{ h^* \left(-\mathbf{A}' \mathbf{z} \right) + g^*(\mathbf{z}) \right\} \quad (21)$$

and the dual problem:

$$\hat{\mathbf{x}} \in \underset{\mathbf{x}}{\operatorname{argmax}} \left\{ -g(\mathbf{A}\mathbf{x}) - h(\mathbf{x}) \right\} \quad (22)$$

of (18), namely the composite convex optimization problem (1). Therefore, the CP iterates (20) solve (1) with rate $O(1/k)$ in an ergodic sense, i.e., with respect to $\mathbf{z}_k$ and $\mathbf{x}_k$ instead of $\mathbf{z}^{(k)}$ and $\mathbf{x}^{(k)}$. Furthermore, Chambolle *et al.* showed that their primal-dual algorithm is equivalent to a preconditioned ADMM. For example, the CP iteration (20) for solving (1) is a proximal-point variant of ADMM with a proximal term weighted by $\mathbf{M} \triangleq \sigma^{-1} \mathbf{I} - \tau \mathbf{A}' \mathbf{A}$ provided that $0 < \sigma\tau\|\mathbf{A}\|_2^2 < 1$ [26, Section 4.3]. Letting $\mathbf{z}^{(k)} = -\tau \mathbf{d}^{(k)}$ and $\mathbf{z}^{(0)} = \mathbf{z}^{(0)}$ and choosing $\sigma = \rho^{-1}t$ and $\tau = \rho$, the CP iteration (20) is just the proximal-point AL method (13)

and hence LALM (8) if we initialize $\mathbf{u}$ as $\mathbf{u}^{(0)} = \mathbf{A}\mathbf{x}^{(0)}$. This suggests that we can measure the convergence rate of LALM using the primal-dual gap that vanishes ergodically with rate $O(1/k)$. Finally, to consider inexact updates, we apply the error analysis technique developed in [17] to the convergence rate analysis of CP, leading to the following theorem (the proof is in the supplementary material):

Theorem 2. Consider a minimax problem (18) where both g and h are closed and proper convex functions. Suppose it has a saddle-point $(\hat{\mathbf{z}}, \hat{\mathbf{x}})$, where $\hat{\mathbf{z}}$ and $\hat{\mathbf{x}}$ are the solutions of the primal problem (21) and the dual problem (22) of (18), respectively. Let $\rho > 0$ and $\{\epsilon_k\}_{k=0}^{\infty}$ denote a non-negative sequence such that

$$\sum_{k=0}^{\infty} \sqrt{\epsilon_k} < \infty. \quad (23)$$

Then, the sequence of updates $\left\{ \left(-\rho \mathbf{d}^{(k)}, \mathbf{x}^{(k)} \right) \right\}_{k=0}^{\infty}$ generated by the inexact LALM in (15) is a bounded sequence that converges to $(\hat{\mathbf{z}}, \hat{\mathbf{x}})$, and the primal-dual gap of $(\mathbf{z}_k, \mathbf{x}_k)$ has the following bound:

$$\Omega(\mathbf{z}_k, \hat{\mathbf{x}}) - \Omega(\hat{\mathbf{z}}, \mathbf{x}_k) \leq \frac{(C + 2A_k + \sqrt{B_k})^2}{k}, \quad (24)$$

$$\text{where } \mathbf{z}_k \triangleq \frac{1}{k} \sum_{j=1}^k \left(-\rho \mathbf{d}^{(j)} \right), \quad \mathbf{x}_k \triangleq \frac{1}{k} \sum_{j=1}^k \mathbf{x}^{(j)},$$

$$C \triangleq \frac{\|\mathbf{x}^{(0)} - \hat{\mathbf{x}}\|_2}{\sqrt{2\rho^{-1}t}} + \frac{\|(-\rho \mathbf{d}^{(0)}) - \hat{\mathbf{z}}\|_2}{\sqrt{2\rho}}, \quad (25)$$

$$A_k \triangleq \sum_{j=1}^k \sqrt{\frac{\epsilon_{j-1}}{(1 - t\|\mathbf{A}\|_2^2)\rho^{-1}t}}, \quad (26)$$

and

$$B_k \triangleq \sum_{j=1}^k \epsilon_{j-1}. \quad (27)$$

Theorem 2 shows that the inexact LALM in (15) converges with rate $O(1/k)$ if the square root of the error ϵ_k is absolutely summable. In fact, even if $\{\sqrt{\epsilon_k}\}_{k=0}^{\infty}$ is not absolutely summable, say, $\sqrt{\epsilon_k}$ decreases as $O(1/k)$, A_k grows as $O(\log k)$ (note that B_k always grows slower than A_k), and the primal-dual gap converges to zero in $O(\log^2 k/k)$. To obtain convergence of the primal-dual gap, a necessary condition is that the partial sum of $\{\sqrt{\epsilon_k}\}_{k=0}^{\infty}$ grows no faster than $O(\sqrt{k})$.

The primal-dual gap convergence bound above is measured at the average point $(-\rho \mathbf{d}_k, \mathbf{x}_k)$ of the update trajectory. In practice, the primal-dual gap of $-\rho \mathbf{d}^{(k)}, \mathbf{x}^{(k)}$ converges much faster. Minimizing the constant in (24) need not provide the fastest convergence rate. However, the ρ -, t -, and ε_k -dependence in (24) suggests how these factors affect the convergence rate of LALM. Note that although we consider only one variable split in our derivation, it is easy to extend our proofs to support multiple variable splits by using the variable splitting scheme in [6]. Conventional LALM in (8) is not OS-accelerable because it needs one full forward projection for the $\mathbf{u}$ - and $\mathbf{d}$ -updates. We used LALM for analysis and to motivate the proposed algorithm in Section III, but it is not recommended for practical implementation in CT reconstruction. By restricting g to be a quadratic loss function, we show next that LALM becomes OS-accelerable and can further accelerate the conventional OS algorithms by decreasing or choosing a small AL penalty parameter.

III. PROPOSED ALGORITHM

A. OS-LALM: an OS-accelerable splitting-based algorithm

In this section, we restrict g to be a quadratic loss function, i.e., $g(\mathbf{u}) \triangleq \frac{1}{2} \|\mathbf{y} - \mathbf{u}\|_2^2$ and then the minimization problem (1) becomes a regularized least-squares problem:

$$\hat{\mathbf{x}} \in \underset{\mathbf{x}}{\operatorname{argmin}} \left\{ \Psi(\mathbf{x}) \triangleq \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + h(\mathbf{x}) \right\}. \quad (28)$$

Let $\mathbf{L}(\mathbf{x}) \triangleq g(\mathbf{A}\mathbf{x})$ denote the quadratic data-fitting term in (28). We assume that $\mathbf{L}$ is suitable for OS acceleration; i.e., $\mathbf{L}$ can be decomposed into M smaller quadratic functions $\mathbf{L}_1, \dots, \mathbf{L}_M$ satisfying the “subset balance condition” [8]:

$$\nabla \mathbf{L}(\mathbf{x}) \approx M \nabla \mathbf{L}_1(\mathbf{x}) \approx \dots \approx M \nabla \mathbf{L}_M(\mathbf{x}), \quad (29)$$

so that the subset gradients approximate the gradient of $\mathbf{L}$.

Since g is quadratic, its proximal mapping is linear. The $\mathbf{u}$ -update in LALM (8) has the following simple closed-form solution:

$$\mathbf{u}^{(k+1)} = \frac{\rho}{\rho+1} (\mathbf{A}\mathbf{x}^{k+1} - \mathbf{d}^{(k)}) + \frac{1}{\rho+1} \mathbf{y}. \quad (30)$$

Combining (30) with the $\mathbf{d}$ -update of (8) yields the identity

$$\mathbf{u}^{(k+1)} + \rho \mathbf{d}^{(k+1)} = \mathbf{y} \quad (31)$$

if we initialize $\mathbf{d}$ as $\mathbf{d}^{(0)} = \rho^{-1}(\mathbf{y} - \mathbf{u}^{(0)})$. Letting $\tilde{\mathbf{u}} \triangleq \mathbf{u} - \mathbf{y}$ denote the split residual and substituting (31) into (8) lead to the following simplified LALM iterates:

$$\begin{cases} \mathbf{s}^{(k+1)} = \mathbf{A}' \left(\rho \left(\mathbf{A} \mathbf{x}^{(k)} - \mathbf{y} \right) + (1 - \rho) \tilde{\mathbf{u}}^{(k)} \right) \\ \mathbf{x}^{k+1} \in \text{prox}_{(\rho^{-1}t)h} \left(\mathbf{x}^{(k)} - (\rho^{-1}t) \mathbf{s}^{(k+1)} \right) \\ \hat{\mathbf{u}}^{(k+1)} = \frac{\rho}{\rho+1} \left(\mathbf{A} \mathbf{x}^{(k+1)} - \mathbf{y} \right) + \frac{1}{\rho+1} \tilde{\mathbf{u}}^{(k)}. \end{cases} \quad (32)$$

The net computational complexity of (32) per iteration reduces to one multiplication by $\mathbf{A}$, one multiplication by $\mathbf{A}'$, and one proximal mapping of h that often can be solved non-iteratively or solved iteratively without using $\mathbf{A}$ or $\mathbf{A}'$. Since the gradient of L is $\mathbf{A}' (\mathbf{A} \mathbf{x} - \mathbf{y})$, letting $\mathbf{g} \triangleq \mathbf{A}' \tilde{\mathbf{u}}$ (a back-projection of the split residual) denote the split gradient, we rewrite (32) as:

$$\begin{cases} \mathbf{s}^{(k+1)} = \rho \nabla L(\mathbf{x}^{(k)}) + (1 - \rho) \mathbf{g}^{(k)} \\ \mathbf{x}^{(k+1)} \in \text{prox}_{(\rho^{-1}t)h} \left(\mathbf{x}^{(k)} - (\rho^{-1}t) \mathbf{s}^{(k+1)} \right) \\ \mathbf{g}^{(k+1)} = \frac{\rho}{\rho+1} \nabla L(\mathbf{x}^{(k+1)}) + \frac{1}{\rho+1} \mathbf{g}^{(k)}. \end{cases}$$

(33)

We call (33) the gradient-based LALM because only the gradients of L are used to perform the updates, and the net computational complexity of (33) per iteration becomes one gradient evaluation of L and one proximal mapping of h .

We interpret the gradient-based LALM (33) as a generalized proximal gradient descent of a regularized least-squares cost function ψ with step size $\rho^{-1}t$ and search direction $\mathbf{s}^{(k+1)}$ that is a weighted average of the gradient and split gradient of L . A smaller ρ can lead to a larger step size. When $\rho = 1$, (33) happens to be the proximal gradient method or the iterative shrinkage/thresholding algorithm (ISTA) [28]. In other words, by using LALM, we can arbitrarily increase the step size of the proximal gradient method by decreasing ρ , thanks to the simple ρ -dependent correction of the search direction in (33). To have a concrete example, suppose all updates are exact, i.e., $\varepsilon_k = 0$ for all k . From (31) and Theorem 2, we have $-\rho \mathbf{d}^{(k)} = \mathbf{u}^{(k)} - \mathbf{y} \rightarrow \mathbf{A} \hat{\mathbf{x}} - \mathbf{y} = \hat{\mathbf{z}}$ as $k \rightarrow \infty$. Furthermore, $(-\rho \mathbf{d}^{(0)}) - \hat{\mathbf{z}} = \mathbf{u}^{(0)} - \mathbf{A} \hat{\mathbf{x}}$. With a reasonable initialization, e.g., $\mathbf{u}^{(0)} = \mathbf{A} \mathbf{x}^{(0)}$ and consequently, $\mathbf{g}^{(0)} = \nabla L(\mathbf{x}^{(0)})$, the constant C in (25) can be rewritten as a function of ρ :

$$C(\rho) = \frac{\|\mathbf{x}^{(0)} - \hat{\mathbf{x}}\|_2}{\sqrt{2\rho^{-1}t}} + \frac{\|\mathbf{A}(\mathbf{x}^{(0)} - \hat{\mathbf{x}})\|_2}{\sqrt{2\rho}}. \quad (34)$$

This constant achieves its minimum at

$$\rho_{opt} = \frac{\|\mathbf{A}(\mathbf{x}^{(0)} - \hat{\mathbf{x}})\|_2}{\sqrt{L}\|\mathbf{x}^{(0)} - \hat{\mathbf{x}}\|_2} \leq 1, \quad (35)$$

suggesting that unity might be a reasonable upper bound on ρ for fast convergence. Note that the ratio of the first term to the second term in (34) is ρ/ρ_{opt} . When the majorization is loose, i.e., $L \gg \|\mathbf{A}\|_2^2 \rho_{\text{opt}} \ll 1$ and the first term in (34) dominates C for $\rho_{\text{opt}} < \rho \ll 1$ since $\rho/\rho_{\text{opt}} \gg 1$. The upper bound of the primal-dual gap becomes

$$\Omega(\mathbf{z}_k, \hat{\mathbf{x}}) - \Omega(\hat{\mathbf{z}}, \mathbf{x}_k) \leq \frac{C^2}{k} \approx \frac{\frac{L}{2} \|\mathbf{x}^{(0)} - \hat{\mathbf{x}}\|_2^2}{\rho^{-1}k}. \quad (36)$$

That is, comparing to the proximal gradient method ($\rho = 1$), the convergence rate (bound) of (33) is accelerated by a factor of ρ^{-1} for $\rho_{\text{opt}} < \rho \ll 1$.

Finally, since (33) requires only the gradients of L to perform updates, it is OS-accelerable. For OS acceleration, we simply replace ∇L in (33) with $M\nabla L_m$ using the approximation (29) and incrementally perform (33) for M times as a complete iteration, thus leading to the final proposed OS-accelerable LALM (OS-LALM):

$$\begin{cases} \mathbf{s}^{(k,m+1)} = \rho M \nabla L_m(\mathbf{x}^{(k,m)}) + (1 - \rho) \mathbf{g}^{(k,m)} \\ \mathbf{x}^{(k,m+1)} \in \text{prox}_{(\rho^{-1}t)h}(\mathbf{x}^{(k,m)} - (\rho^{-1}t) \mathbf{s}^{(k,m+1)}) \\ \mathbf{g}^{(k,m+1)} = \frac{\rho}{\rho+1} M \nabla L_{m+1}(\mathbf{x}^{(k,m+1)}) + \frac{1}{\rho+1} \mathbf{g}^{(k,m)} \end{cases}$$

(37)

with $\mathbf{c}^{(k,M+1)} = \mathbf{c}^{(k+1)} = \mathbf{c}^{(k+1,1)}$ for $\mathbf{c} \in \{\mathbf{s}, \mathbf{x}, \mathbf{g}\}$ and $L_{M+1} = L_1$. Like typical OS-based algorithms, this algorithm is convergent when $M = 1$, i.e., (33), but is not guaranteed to converge for $M > 1$. When $M > 1$, updates generated by OS-based algorithms approach a “limit cycle” in which updates stop nearing the optimum, and visible OS artifacts might be observed in the reconstructed image, depending on M .

B. Deterministic downward continuation

One drawback of conventional LALM is the difficulty of finding a fixed value for the penalty parameter ρ that provides the fastest convergence. The optimal penalty parameter ρ_{opt} in (35) minimizes the multiplicative constant but depends on the unknown solution $\hat{\mathbf{x}}$ of the problem. Intuitively, a smaller ρ is better because it leads to a larger step size. However, when the step size is too large, one can encounter overshoots and oscillations that slow down the convergence rate at first and when nearing the optimum. In fact, ρ_{opt} in (35) also suggests that ρ should not be arbitrarily small. Rather than estimating ρ_{opt} heuristically, we focus on using an iteration-dependent ρ , i.e., a continuation approach, for acceleration.

Classic continuation approaches increase ρ as iterations progress so that previous iterates can serve as a warm start for subsequent worse-conditioned but more penalized inner minimization problems [29, Proposition 4.2.1]. To implement this kind of approaches, one must specify an initial value of the penalty parameter ρ_0 and a rule for increasing ρ which are usually problem-dependent. For example, a small ρ_0 provides fast initial convergence but can cause overshoot (e.g., increasing the cost function) in early iterations. Choosing a

good ρ_0 that balances convergence rate and overshoot for a given problem is difficult. Furthermore, most classic continuation approaches adapt the penalty parameter dynamically by checking some conditions such as the primal and dual feasibilities and the decrease of cost function [30]. This can increase computational complexity per iteration, especially when involving expensive operations (e.g., $\mathbf{A}$ and $\mathbf{A}'$ in our case).

In this paper, unlike classic continuation approaches, we consider a downward continuation approach. It is inspired by Nesterov's second method [31] that starts as a proximal gradient method and gradually increases the step size (of the auxiliary sequence) deterministically. The intuition is that, for a fixed ρ , the step length $\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|$ is typically a decreasing sequence because the gradient norm vanishes as we approach the optimum, and an increasing sequence ρ_k (i.e., a diminishing step size) would aggravate the shrinkage of step length, slowing convergence. In contrast, a decreasing sequence ρ_k can compensate for step length shrinkage and accelerate convergence. Of course, ρ_k cannot decrease too fast; otherwise, the soaring step size might make the algorithm unstable or even divergent. To design a "good" decreasing sequence ρ_k for "effective" acceleration, we first analyze how LALM (the one-subset version (33) for simplicity) behaves for different values of ρ .

Consider a simple quadratic problem:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{A}\mathbf{x}\|_2^2, \quad (38)$$

corresponding to (28) with $h = 0$ and $\mathbf{y} = \mathbf{0}$. A trivial solution of (38) is $\hat{\mathbf{x}} = \mathbf{0}$. To ensure a unique solution, we assume that $\mathbf{A}'\mathbf{A}$ is positive definite (for this analysis only). Let $\mathbf{A}'\mathbf{A}$ have eigenvalue decomposition $\mathbf{V}\mathbf{\Lambda}\mathbf{V}^t$, where $\mathbf{\Lambda} \triangleq \operatorname{diag}\{\lambda_i\}$ and $0 < \lambda_1 \leq \dots \leq \lambda_n = L$. The updates generated by (33) that solve (38) can be written as

$$\begin{cases} \mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} - (1/L) \left(\mathbf{V}\mathbf{\Lambda}\mathbf{V}'\mathbf{x}^{(k)} + (\rho^{-1} - 1) \mathbf{g}^{(k)} \right) \\ \mathbf{g}^{(k+1)} = \frac{\rho}{\rho+1} \mathbf{V}\mathbf{\Lambda}\mathbf{V}'\mathbf{x}^{(k+1)} + \frac{1}{\rho+1} \mathbf{g}^{(k)}. \end{cases} \quad (39)$$

Furthermore, letting $\bar{\mathbf{x}} = \mathbf{V}'\mathbf{x}$ and $\bar{\mathbf{g}} = \mathbf{V}'\mathbf{g}$, the linear system can be further diagonalized, and we can represent the i th components of $\bar{\mathbf{x}}$ and $\bar{\mathbf{g}}$ as

$$\begin{cases} \bar{x}_i^{(k+1)} = \bar{x}_i^{(k)} - (1/L) \left(\lambda_i \bar{x}_i^{(k)} + (\rho^{-1} - 1) \bar{g}_i^{(k)} \right) \\ \bar{g}_i^{(k+1)} = \frac{\rho}{\rho+1} \lambda_i \bar{x}_i^{(k+1)} + \frac{1}{\rho+1} \bar{g}_i^{(k)}. \end{cases} \quad (40)$$

Solving this system of recurrence relations of $\bar{x}_i$ and $\bar{g}_i$, one can show that both $\bar{x}_i$ and $\bar{g}_i$ satisfy a second-order recursive system determined by the characteristic polynomial:

$$(1+\rho) r_i^2 - 2(1 - \lambda_i/L + \rho/2) r_i + (1 - \lambda_i/L). \quad (41)$$

The roots r_i of this polynomial determine the convergence rate of x_i and $\bar{g}_i$ in (40).

When $\rho = \rho_i^c$, where

$$\rho_i^c \triangleq 2 \sqrt{\frac{\lambda_i}{L} \left(1 - \frac{\lambda_i}{L}\right)} \in (0, 1], \quad (42)$$

the characteristic polynomial (41) has repeated roots. Hence, the system is critically damped, and x_i and $\bar{g}_i$ converge geometrically to zero with convergence rate

$$r_i^c = \frac{1 - \lambda_i/L + \rho_i^c/2}{1 + \rho_i^c} = \sqrt{\frac{1 - \lambda_i/L}{1 + \rho_i^c}}. \quad (43)$$

When $\rho > \rho_i^c$, the characteristic polynomial (41) has distinct real roots. Hence, the system is over-damped, and x_i and $\bar{g}_i$ converge geometrically to zero with convergence rate that is governed by the dominant root

$$r_i^o(\rho) = \frac{1 - \lambda_i/L + \rho/2 + \sqrt{\rho^2/4 - \lambda_i/L(1 - \lambda_i/L)}}{1 + \rho}. \quad (44)$$

It is easy to check that $r_i^o(\rho_i^c) = r_i^c$, and r_i^o is increasing. This suggests that the critically damped system always converges faster than the over-damped system. Finally, when $\rho < \rho_i^c$, the characteristic polynomial (41) has complex roots. In this case, the system is under-damped, and x_i and $\bar{g}_i$ converge geometrically to zero with convergence rate

$$r_i^u(\rho) = \frac{1 - \lambda_i/L + \rho/2}{1 + \rho}, \quad (45)$$

and oscillate at the damped frequency $\Psi_i/(2\pi)$, where

$$\cos \psi_i = \frac{1 - \lambda_i/L + \rho/2}{\sqrt{(1 + \rho)(1 - \lambda_i/L)}} \approx \sqrt{1 - \lambda_i/L} \quad (46)$$

when $\rho \approx 0$. Furthermore, by the small angle approximation:

$$\cos \sqrt{\theta} \approx 1 - \theta/2 \approx \sqrt{1 - \theta}, \quad \text{if } \lambda_i \ll L, \quad \psi_i \approx \sqrt{\lambda_i/L}. \text{ Again,}$$

$r_i^u(\rho_i^c) = r_i^c$, but r_i^u behaves differently from r_i^o . Specifically, r_i^u is decreasing if $\lambda_i/L < 1/2$, and it is increasing otherwise. This suggests that the critically damped system converges faster than the under-damped system if $\lambda_i/L < 1/2$, but it can be slower otherwise. In sum, the critically damped system is optimal for eigencomponents having smaller eigenvalues (i.e., $\lambda_i < L/2$), while for eigencomponents having larger eigenvalues (i.e., $\lambda_i > L/2$), the under-damped system is optimal.

The asymptotic convergence rate of the system is dominated by the smallest eigenvalue λ_1 . Eventually, only the component oscillating at frequency $\Psi_1/(2\pi)$ persists. Therefore, for the fastest asymptotic convergence rate, we would like to choose the AL penalty parameter ρ to be

$$\rho^* = \rho_1^c = 2 \sqrt{\frac{\lambda_1}{L} \left(1 - \frac{\lambda_1}{L}\right)} \in (0, 1]. \quad (47)$$

Figure 1 illustrates the dominant roots of (41) and the optimal asymptotic convergence rate that is the minimum of the largest dominant root, i.e., $\max_i r_i(\rho)$, of all possible ρ for a system with six distinct eigenvalues: $0.05L$, $0.1L$, $0.3L$, $0.7L$, $0.9L$, and L . For eigencomponents having smaller eigenvalues ($0.05L$, $0.1L$, and $0.3L$), the critically damped system has the fastest asymptotic convergence rate, while eigencomponents having larger eigenvalues ($0.7L$, $0.9L$, and L) attain the fastest asymptotic convergence rate in the under-damping regime. Moreover, when the smallest eigenvalue is less than $L/2$ (i.e., an ill-conditioned system), the eigencomponent with the smallest eigenvalue determines the optimal asymptotic convergence rate, i.e., $\min_\rho \{\max_i r_i(\rho)\}$, with ρ^* in (47). Unlike ρ_{opt} in (35), this choice of ρ does not depend on the initialization. It depends only on the geometry of the Hessian $\mathbf{A}'\mathbf{A}$. Furthermore, both ρ_{opt} and ρ^* fall in the interval $(0, 1]$. Hence, although LALM converges for any $\rho > 0$, we consider only $0 < \rho \leq 1$ in our downward continuation approach.

We can now interpret classic (upward) continuation approaches based on the second-order recursive system analysis. Classic continuation approaches usually start from a small ρ for better-conditioned inner minimization problem. Therefore, initially, the system is under-damped. Although the under-damped system has a slower asymptotic convergence rate, the oscillation can provide dramatic acceleration before the first zero-crossing of the oscillating components. We can think of the classic continuation approach as a greedy strategy that exploits the initial fast convergence rate of the under-damped system and carefully increases ρ to avoid oscillation and move toward the critical damping regime. However, this greedy strategy requires a “clever” update rule for increasing ρ . If ρ increases too fast, the acceleration ends prematurely; if ρ increases too slow, the system starts oscillating.

In contrast, we consider a more conservative strategy that starts from the over-damped regime, say, $\rho = 1$ as suggested in (47), and gradually reduces ρ to the optimal AL penalty parameter ρ^* . It sounds impractical at first because we do not know λ_1 beforehand. To solve this problem, we adopt the adaptive restart proposed in [32] and generate a decreasing sequence ρ_k that starts from $\rho = 1$ and reaches ρ^* every time the algorithm restarts. As mentioned before, the system oscillates at frequency $\Psi_1/(2\pi)$ when it is under-damped. This oscillating behavior can also be observed from the trajectory of updates. For example,

$$\xi(k) \triangleq \left(\mathbf{g}^{(k)} - \nabla L(\mathbf{x}^{(k+1)}) \right)' \left(\nabla L(\mathbf{x}^{(k+1)}) - \nabla L(\mathbf{x}^{(k)}) \right) \quad (48)$$

oscillates at the frequency Ψ_1/π [32]. Hence, we restart the algorithm (i.e., reset the decreasing penalty parameter ρ to be one and $\mathbf{g}$ to be the current gradient of L) every time $\xi > 0$, which should occur about every $(\pi/2) \sqrt{L/\lambda_1}$ iterations. Suppose we restart at the r th iteration, we have the approximation $\sqrt{\lambda_1/L} \approx \pi/(2r)$, and the ideal AL penalty parameter at the r th iteration should be

$$2 \sqrt{\left(\frac{\pi}{2r}\right)^2 \left(1 - \left(\frac{\pi}{2r}\right)^2\right)} = \frac{\pi}{r} \sqrt{1 - \left(\frac{\pi}{2r}\right)^2}. \quad (49)$$

The proposed downward continuation approach has the form (33), where we replace every ρ in (33) with

$$\rho_l = \begin{cases} 1, & \text{if } l=0 \\ \max \left\{ \frac{\pi}{l+1} \sqrt{1 - \left(\frac{\pi}{2l+2}\right)^2}, \rho_{\min} \right\}, & \text{otherwise,} \end{cases} \quad (50)$$

where l is a counter that starts from zero, increases by one everytime $\mathbf{g}$ is updated, and is reset to zero whenever $\xi(k) > 0$. The lower bound $\rho_{\min}$ is a small positive number for guaranteeing convergence. Note that ADMM is convergent if ρ is non-increasing and bounded below away from zero [33, Corollary 4.2]. As shown in Section II-B (Theorem 1), LALM is a convergent ADMM. Therefore, we can ensure convergence (of the one-subset version) of the proposed downward continuation approach if we set a non-zero lower bound for ρ_l , e.g., $\rho_{\min} = 10^{-3}$ in our experiments. Note that ρ_l in (50) is the same for any $\mathbf{A}$. The adaptive restart condition takes care of the dependence on $\mathbf{A}$. That is why we call this approach the deterministic downward continuation approach. When h is non-zero and/or $\mathbf{A}'\mathbf{A}$ is not positive definite, our analysis above does not hold. However, the deterministic downward continuation approach works well in practice for CT. One possible explanation is that the cost function can usually be well approximated by a quadratic near the optimum when the minimization problem is well-posed and h is locally quadratic.

Finally, in practice we do not restart our algorithm for X-ray CT image reconstruction (with ordered subsets, i.e., (37)) for two reasons. First, according to our analysis, the restart period is proportional to the square root of the local condition number, i.e., $\sqrt{L/\lambda_1}$. Since X-ray CT image reconstruction problems are usually very ill-conditioned, the algorithm usually terminates before a restart is needed. Second, since OS is used for acceleration, gradients used to compute $\xi^{(k)}$ are not accurate and might lead to premature restart. In our experimental results, we did not observe any problems without restart. However, restart may be useful in other applications.

IV. IMPLEMENTATION DETAILS

In this section, we consider solving the X-ray CT image reconstruction problem:

$$\hat{\mathbf{x}} \in \underset{\mathbf{x} \in \Omega}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_{\mathbf{W}}^2 + R(\mathbf{x}) \right\} \quad (51)$$

using the proposed OS-LALM algorithm (37), where $\mathbf{A}$ is the system matrix of a CT scan, $\mathbf{y}$ is the noisy sinogram, $\mathbf{W}$ is the statistical weighting matrix, R is an edge-preserving regularizer, and Ω denotes the convex set for a box constraint (usually the non-negativity constraint) on $\mathbf{x}$. We focus on the edge-preserving regularizer R defined as:

$$R(\mathbf{x}) \triangleq \sum_i \beta_i \sum_n k_n k_{n+s_i} \phi_i([\mathbf{C}_i \mathbf{x}]_n), \quad (52)$$

where β_i , s_i , ϕ_i , and $\mathbf{C}_i$ denote the regularization parameter, corresponding offset, potential function, and finite difference matrix in the i th direction, respectively, and k_n is a voxel-dependent weight for improving resolution uniformity [34]. In our experiments, we use 13 directions to include all neighbors in 3D CT.

A. OS-LALM for X-ray CT image reconstruction

The X-ray CT image reconstruction problem (51) is a constrained regularized weighted least-squares problem. To solve it using the proposed algorithm (33) and its OS variant (37), we use the following substitution:

$$\begin{cases} \mathbf{A} \leftarrow \mathbf{W}^{1/2} \mathbf{A} \\ \mathbf{y} \leftarrow \mathbf{W}^{1/2} \mathbf{y} \\ h \leftarrow R + \iota_{\Omega}, \end{cases} \quad (53)$$

where ι_C denotes the characteristic function of a convex set C . Thus, the inner minimization problem in (33) and its OS variant (37) becomes a constrained denoising problem. In our implementation, we solve this inner constrained denoising problem using n iterations of the fast iterative shrinkage/thresholding algorithm (FISTA) [15] starting from the previous update as a warm start. As discussed in Section II-B, inexact updates can slow down the convergence rate of the proposed algorithm. In general, the more FISTA iterations, the faster convergence rate of the proposed algorithm. However, the overhead of iterative inner updates is non-negligible for large n , especially when the number of subsets is large. Fortunately, in typical X-ray CT image reconstruction problems, the majorization is usually very loose (probably due to the huge dynamic range of the statistical weighting $\mathbf{W}$).

Therefore, $t \ll 1$ in most cases, greatly diminishing the regularization force in the constrained denoising problem. In practice, the constrained denoising problem can be solved up to some acceptable tolerance within just one or two iterations. For a fair comparison with the OS-based algorithms [8, 12] used in Section V that majorize the weighted quadratic data-fitting

term L using the SQS function with a diagonal Hessian $\mathbf{D}_L \triangleq \operatorname{diag} \{ \mathbf{A}' \mathbf{W} \mathbf{A} \mathbf{1} \}$ [8], in our experiments, we also majorize the quadratic penalty in the scaled augmented Lagrangian using the SQS function with Hessian $\mathbf{D}_L$ (in this case, the inner minimization problem in (33) and its OS variant (37) becomes a constrained weighted denoising problem that can also be solved by FISTA) and incrementally update the image using the subset gradients

with the bit-reversal order [35] that heuristically minimizes the subset gradient variance as in other OS-based algorithms. See the supplementary material for the outline of the proposed OS-LALM algorithm for solving (51).

The SQS function with Hessian $\mathbf{D}_L$ is a very loose majorizer. For fastest convergence, one might wish to use the tightest majorizer with Hessian $\mathbf{A}'\mathbf{W}\mathbf{A}$. However, this would revert to the standard AL method (4) with expensive $\mathbf{x}$ -updates. An alternative is the Barzilai-Borwein (spectral) method [36] that mimics the Hessian $\mathbf{A}'\mathbf{W}\mathbf{A}$ by $\mathbf{H}_k \triangleq \alpha_k \mathbf{D}_L$, where the scaling factor α_k is solved by fitting the secant equation in the (weighted) least-squares sense. Detailed derivation and additional experimental results can be found in the supplementary material.

B. Number of subsets

As mentioned in Section III-A, the number of subsets M can affect the stability of OS-based algorithms. When M is too large, OS algorithms typically become unstable, causing artifacts in the reconstructed image. Therefore, finding an appropriate number of subsets is very important. Since errors of OS-based algorithms come from the gradient approximation using subset gradients, artifacts might be suppressed using a better gradient approximation. Intuitively, to have a reasonable gradient approximation, each voxel in a subset should be sampled by a minimum number of views s . For simplicity, we consider the central voxel in the transaxial plane. In axial CT, the views are uniformly distributed in each subset, so we want

$$\frac{1}{M_{axial}} \cdot (\text{number of views}) \geq s_{axial}. \quad (54)$$

This leads to our maximum number of subsets for axial CT:

$$M_{axial} \leq (\text{number of views}) \cdot \frac{1}{s_{axial}}. \quad (55)$$

Helical CT is more complicated. Since the X-ray source moves in the z direction, a central voxel is only covered by $d_{so}/(p \cdot d_{sd})$ turns, where p is the pitch, d_{so} denotes the distance from the X-ray source to the isocenter, and d_{sd} denotes the distance from the X-ray source to the detector. Therefore, we want

$$\frac{1}{M_{helical}} \cdot (\text{number of views per turn}) \cdot \frac{d_{so}}{p \cdot d_{sd}} \geq s_{helical}. \quad (56)$$

This leads to our maximum number of subsets for helical CT:

$$M_{helical} \leq (\text{number of views per turn}) \cdot \frac{d_{so}}{p \cdot s_{helical} \cdot d_{sd}}. \quad (57)$$

Note that the maximum number of subsets for helical CT M_{helical} is inversely proportional to the pitch p . We set $s_{\text{axial}} \approx 40$ and $s_{\text{helical}} \approx 24$ for the proposed algorithm in our experiments.

V. EXPERIMENTAL RESULTS

This section reports numerical results for 3D X-ray CT image reconstruction from real CT scans with different geometries using various OS-based algorithms, including

- **OS-SQS- M** : the standard OS algorithm [8] with M subsets,
- **OS-Nes05- M** : the OS+momentum algorithm [12] based on Nesterov's fast gradient method [14] with M subsets,
- **OS-LALM- M - ρ - n** : the proposed algorithm using a fixed AL penalty parameter ρ with M subsets and n FISTA iterations for solving the inner constrained denoising problem, and
- **OS-LALM- M -c- n** : the proposed algorithm using the deterministic downward continuation approach described in Section III-B with M subsets and n FISTA iterations for solving the inner constrained denoising problem.

OS-SQS is a standard iterative method for tomographic reconstruction, and OS-Nes05 is a state-of-the-art method for fast X-ray CT image reconstruction using Nesterov's momentum technique. Unlike other OS-based algorithms, our proposed algorithm has additional overhead due to the iterative inner updates. However, when $n = 1$, i.e., with a single gradient descent for the constrained denoising problem, all algorithms listed above have the same computational complexity (one forward/back-projection pair and M regularizer gradient evaluations per iteration). When majorizing the regularizer, we use Huber's curvature [37, p. 185] for faster convergence in all algorithms. Therefore, comparing the convergence rate as a function of iteration is fair. We measured the convergence rate using the RMS difference (in the region of interest) between the reconstructed image $\mathbf{x}^{(k)}$ and the almost converged reference reconstruction $\mathbf{x}^*$ that we generated by running several iterations of the OS +momentum algorithm with a small M , followed by 2000 iterations of a convergent (i.e., one-subset) FISTA with adaptive restart [32]. We used a q -generalized Gaussian [38] potential function ϕ_i in (52), and β_i and κ_n were tuned to emulate GE's Veo method [39]. For reproducible results using an XCAT phantom, see the supplementary material.

A. Shoulder scan

In this experiment, we reconstructed a $512 \times 512 \times 109$ image from a shoulder region helical CT scan, where the sinogram has size $888 \times 32 \times 7146$ and pitch 0.5. The maximum number of subsets suggested by (57) is about 40. Figure 2 shows the cropped images from the central transaxial plane of the initial FBP image, the reference reconstruction, and the reconstructed image using the proposed algorithm (OS-LALM-40-c-1) at the 30th iteration (i.e., after 30 forward/back-projection pairs). In Figure 2, the reconstructed image using the proposed algorithm looks almost the same as the reference reconstruction in the display window from 800 to 1200 Hounsfield unit (HU, modified so that air is 0). The reconstructed image using

the OS+momentum algorithm (not shown here) also looks quite similar to the reference reconstruction.

Figure 3 shows the difference images, i.e., $\mathbf{x}^{(30)} - \mathbf{x}^*$, for different OS-based algorithms. The standard OS algorithm (with both 20 and 40 subsets) exhibits visible streak artifacts and structured high frequency noise in the difference image. When $M = 20$, the difference images look similar for the OS+momentum algorithm and our proposed algorithm, although that of the OS+momentum algorithm is slightly more structured and non-uniform. When $M = 40$, the difference image for our proposed algorithm remains uniform, whereas some noise-like OS artifacts appear in the OS+momentum algorithm's difference image. The OS artifacts for the OS+momentum algorithm worsen when M increases, e.g., $M = 80$ (not shown). Apparently OS-LALM has better gradient error tolerance than previous OS methods, probably due to the way we compute the search direction and the less aggressive acceleration to the regularization term. Additional experimental results (an XCAT phantom axial scan and a truncated abdomen scan) in the supplementary material demonstrate how different OS-based algorithms behave when M exceeds the suggested maximum number of subsets.

Figure 4 shows the convergence rate curves (RMS differences between the reconstructed image $\mathbf{x}^{(k)}$ and the reference reconstruction $\mathbf{x}^*$ as a function of iteration) using OS-based algorithms with 20 and 40 subsets. By exploiting the linearized AL method, the proposed algorithm accelerates the standard OS algorithm remarkably. As mentioned in Section III-A, a smaller ρ can provide greater acceleration due to the increased step size. Both plots show the acceleration of convergence as ρ decreases. Note that too large step sizes can cause overshoots in early iterations. For example, the proposed algorithm with $\rho = 0.05$ shows slower convergence rate in first few iterations but decreases more rapidly later. Our proposed deterministic downward continuation approach (50) overcomes this trade-off. In Figure 4, the proposed algorithm using deterministic downward continuation reaches the lowest RMSD (lower than 1 HU) within only 30 iterations. The slightly higher RMSD of the OS+momentum algorithm with 40 subsets is due to the OS artifacts seen in Figure 3.

Figure 5 illustrates the effectiveness of solving the inner constrained denoising problem using FISTA (for X-ray CT image reconstruction) mentioned in Section IV-A. In Figure 5, the convergence rate improves only slightly when using more than one FISTA iteration for solving the inner constrained denoising problem. In practice, one FISTA iteration, i.e., $n = 1$, per subset update suffices for fast and accurate X-ray CT image reconstruction.

B. GE performance phantom

In this experiment, we reconstructed a $1024 \times 1024 \times 90$ image from the GE performance phantom (GEPP) axial CT scan, where the sinogram has size $888 \times 64 \times 984$. The maximum number of subsets suggested by (55) is about 24. Figure 6 shows the cropped images from the central transaxial plane of the initial FBP image, the reference reconstruction, and the reconstructed image using the proposed algorithm (OS-LALM-24-c-1) at the 30th iteration. Again, the OS-LALM image at the 30th iteration is very similar to the reference reconstruction.

Figure 7 and Figure 8 show the difference images and convergence rate curves, respectively. Because of the lower view-redundancy in axial CT scans, the OS+momentum algorithm shows even more OS artifacts than the standard OS algorithm in the difference images, leading to a larger limit cycle in Figure 8. A relaxed OS+momentum algorithm [40] that uses a diminishing step size can address this problem but require careful tuning of the step sizes. In contrast, the proposed OS-LALM algorithm avoids the need for such parameter tuning; one only needs to choose the number of subsets M .

VI. CONCLUSION

Augmented Lagrangian (AL) methods and ordered subsets (OS) are two powerful techniques for accelerating optimization algorithms using decomposition and approximation, respectively. This paper combined these two techniques by considering a linearized variant of the AL method and proposed a fast OS-accelerated splitting-based algorithm, OS-LALM, for solving regularized (weighted) least-squares problems. We also proposed a novel deterministic downward continuation approach based on a second-order damping system that simplifies parameter selection; only the number of subsets M needs to be selected, and we provided heuristics for that based on sampling considerations in (55) and (57). We applied OS-LALM to X-ray computed tomography (CT) image reconstruction problems and compared with some state-of-the-art OS methods using real CT scans with different geometries. Experimental results showed that OS-LALM exhibits fast convergence rate and excellent gradient error tolerance.

In (33), the search direction $\mathbf{s}$ is a weighted average of the current gradient and the split gradient of L corresponding to a low-pass infinite-impulse-response filter (across iterations). The gradient error might be suppressed by this low-pass filter, improving stability. A similar averaging technique (with a low-pass finite-impulse-response filter) is used in the stochastic average gradient (SAG) method [41, 42]. In contrast, the OS+momentum algorithm computes the search direction using only the current gradient (of the auxiliary image), so the gradient error can accumulate when OS is used, providing a less stable reconstruction. When the inner constrained denoising problem is more difficult to solve, one could run more FISTA iterations or introduce an additional split variable for the regularizer as in [7] at the cost of higher memory burden, thus leading to a “high-memory” version of OS-LALM [43]. A recent alternative without variable splitting would be to use a grouped coordinate descent (GCD) denoising with a GPU implementation [44, 45].

As future work, we are interested in the convergence rate analysis of the proposed algorithm with the deterministic downward continuation approach and a more rigorous convergence analysis of OS-LALM for M that is greater than one.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

The authors thank GE Healthcare for providing sinogram data in our experiments.

This work is supported in part by National Institutes of Health (NIH) grants R01-HL-098686 and U01-EB-18753 and by an equipment donation from Intel Corporation.

REFERENCES

- [1]. Hestenes MR. Multiplier and gradient methods. *J. Optim. Theory Appl.* Nov.1969 4:303–20.
- [2]. Glowinski R, Marrocco A. Sur l'approximation par éléments nis d'ordre un, et la résolution par pénalisation-dualité d'une classe de problèmes de dirichlet nonlinéaires, *rev. française d'auto. Inf. Rech. Oper.* 1975; R-2:41–76.
- [3]. Gabay D, Mercier B. A dual algorithm for the solution of nonlinear variational problems via finite-element approximations. *Comput. Math. Appl.* 1976; 2(1):17–40.
- [4]. Eckstein J, Bertsekas DP. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming.* Apr.1992 55:293–318.
- [5]. Goldstein T, Osher S. The split Bregman method for L1-regularized problems. *SIAM J. Imaging Sci.* 2009; 2(2):323–43.
- [6]. Afonso MV, Bioucas-Dias JM, Figueiredo MAT. An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems. *IEEE Trans. Im. Proc.* Mar. 2011 20:681–95.
- [7]. Ramani S, Fessler JA. A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction. *IEEE Trans. Med. Imag.* Mar.2012 31:677–88.
- [8]. Erdem H, Fessler JA. Ordered subsets algorithms for transmission tomography. *Phys. Med. Biol.* Nov.1999 44:2835–51. [PubMed: 10588288]
- [9]. Bertsekas, DP. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. Aug. 2010 2010(revised December 2010) Report LIDS 2848
- [10]. Robbins H, Monro S. A stochastic approximation method. *Ann. Math. Stat.* Sep.1951 22:400–7.
- [11]. Kim, D.; Ramani, S.; Fessler, JA. Ordered subsets with momentum for accelerated X-ray CT image reconstruction. *Proc. IEEE Conf. Acoust. Speech Sig. Proc.*; 2013. p. 920-3.
- [12]. Kim, D.; Ramani, S.; Fessler, JA. Accelerating X-ray CT ordered subsets image reconstruction with Nesterov's first-order methods. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.*; 2013. p. 22-5.
- [13]. Nesterov Y. A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$. *Dokl. Akad. Nauk. USSR.* 1983; 269(3):543–7.
- [14]. Nesterov Y. Smooth minimization of non-smooth functions. *Mathematical Programming.* May. 2005 103:127–52.
- [15]. Beck A, Teboulle M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.* 2009; 2(1):183–202.
- [16]. Devolder, O. Stochastic first order methods in smooth convex optimization. 2011.
- [17]. Schmidt, M.; Le Roux, N.; Bach, F. Convergence rates of inexact proximal-gradient methods for convex optimization. *Adv. in Neural Info. Proc. Sys.*; 2011. p. 1458-66.
- [18]. Ahn S, Fessler JA. Globally convergent image reconstruction for emission tomography using relaxed ordered subsets algorithms. *IEEE Trans. Med. Imag.* May.2003 22:613–26.
- [19]. Ahn S, Fessler JA, Blatt D, Hero AO. Convergent incremental optimization transfer algorithms: Application to tomography. *IEEE Trans. Med. Imag.* Mar.2006 25:283–96.
- [20]. Ouyang, H.; He, N.; Tran, L.; Gray, AG. Stochastic alternating direction method of multipliers. In: Dasgupta, S.; Mcallester, D., editors. *Proc. Intl. Conf. Machine Learning; JMLR Workshop and Conference Proceedings*; 2013. p. 80-8.
- [21]. Lin, Z.; Liu, R.; Su, Z. Linearized alternating direction method with adaptive penalty for low-rank representation. *Adv. in Neural Info. Proc. Sys.*; 2011. p. 612-20.
- [22]. Esser E, Zhang X, Chan T. A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science. *SIAM J. Imaging Sci.* 2010; 3(4):1015–46.
- [23]. Zhang X, Burger M, Bresson X, Osher S. Bregmanized nonlocal regularization for deconvolution and sparse reconstruction. *SIAM J. Imaging Sci.* 2010; 3(3):253–76.

- [24]. Zhang X, Burger M, Osher S. A unified primal-dual algorithm framework based on Bregman iteration. *Journal of Scientific Computing*. 2011; 46(1):20–46.
- [25]. Kim D, Pal D, Thibault J-B, Fessler JA. Accelerating ordered subsets image reconstruction for X-ray CT using spatially non-uniform optimization transfer. *IEEE Trans. Med. Imag.* Nov.2013 32:1965–78.
- [26]. Chambolle A, Pock T. A first-order primal-dual algorithm for convex problems with applications to imaging. *J. Math. Im. Vision*. 2011; 40(1):120–45.
- [27]. Rockafellar, RT. *Convex analysis*. Princeton University Press; Princeton: 1970.
- [28]. Daubechies I, Defrise M, Mol CD. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.* Nov.2004 57:1413–57.
- [29]. Bertsekas, DP. *Nonlinear programming*. 2. Athena Scientific; Belmont: 1999.
- [30]. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. & Trends in Machine Learning*. 2010; 3(1):1–122.
- [31]. Nesterov Y. On an approach to the construction of optimal methods of minimization of smooth convex functions. *Ekonomika i Mateaticheskie Metody*. 1988; 24:509–17. In Russian.
- [32]. O'Donoghue B, Candès E. Adaptive restart for accelerated gradient schemes. *Found. Comput. Math.* Jul.2013 13
- [33]. Kontogiorgis S, Meyer RR. A variable-penalty alternating directions method for convex optimization. *Mathematical Programming*. 1998; 83(1–3):29–53.
- [34]. Fessler JA, Rogers WL. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Trans. Im. Proc.* Sep.1996 5:1346–58.
- [35]. Herman GT, Meyer LB. Algebraic reconstruction techniques can be made computationally efficient. *IEEE Trans. Med. Imag.* Sep.1993 12:600–9.
- [36]. Barzilai J, Borwein J. Two-point step size gradient methods. *IMA J. Numerical Analysis*. 1988; 8(1):141–8.
- [37]. Huber, PJ. *Robust statistics*. Wiley; New York: 1981.
- [38]. Thibault J-B, Sauer K, Bouman C, Hsieh J. A three-dimensional statistical approach to improved image quality for multi-slice helical CT. *Med. Phys.* Nov.2007 34:4526–44. [PubMed: 18072519]
- [39]. Shuman WP, Green DE, Busey JM, Kolokythas O, Mitsumori LM, Koprowicz KM, Thibault J-B, Hsieh J, Alessio AM, Choi E, Kinahan PE. Model based iterative reconstruction versus adaptive statistical iterative reconstruction and filtered back projection in 64-MDCT: Focal lesion detection, lesion conspicuity, and image noise. *Am. J. Roentgenol.* May.2013 200:1071–6. [PubMed: 23617492]
- [40]. Kim, D.; Fessler, JA. Ordered subsets acceleration using relaxed momentum for X-ray CT image reconstruction. *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*; 2013. p. 1-5.
- [41]. Le Roux, N.; Schmidt, M.; Bach, F. A stochastic gradient method with an exponential convergence rate for strongly-convex optimization with finite training sets. *Adv. in Neural Info. Proc. Sys.*; 2012. p. 2672-80.
- [42]. Schmidt, M.; Le Roux, N.; Bach, F. Minimizing finite sums with the stochastic average gradient. 2013.
- [43]. Nien, H.; Fessler, JA. Fast splitting-based ordered-subsets X-ray CT image reconstruction. *Proc. 3rd Intl. Mtg. on image formation in X-ray CT*; 2014. p. 291-4.
- [44]. McGaffin, MG.; Fessler, JA. Fast edge-preserving image denoising via group coordinate descent on the GPU. *Proc. SPIE 9020 Computational Imaging XII*; 2014. p. 90200P
- [45]. McGaffin MG, Fessler JA. Edge-preserving image denoising via group coordinate descent on the GPU. *IEEE Trans. Im. Proc.* 2014 Submitted.

Fig. 1.

The dominant roots $r_i(\rho)$ of (41) and the optimal asymptotic convergence rate

$r^* \triangleq \min_{\rho} \{ \max_i r_i(\rho) \}$ for a system with six distinct eigenvalues: $0.05L$, $0.1L$, $0.3L$, $0.7L$, $0.9L$, and L .

Fig. 2. Shoulder scan: cropped images (displayed from 800 to 1200 HU) from the central transaxial plane of the initial FBP image $\mathbf{x}^{(0)}$ (left), the reference reconstruction $\mathbf{x}^*$ (center), and the reconstructed image using the proposed algorithm (OS-LALM-40-c-1) at the 30th iteration $\mathbf{x}^{(30)}$ (right).

Fig. 3. Shoulder scan: cropped difference images (displayed from -30 to 30 HU) from the central transaxial plane of $\mathbf{x}^{(30)} - \mathbf{x}^*$ using OS-based algorithms.

Fig. 4.

Shoulder scan: RMS differences between the reconstructed image $\mathbf{x}^{(k)}$ and the reference reconstruction $\mathbf{x}^*$ as a function of iteration using OS-based algorithms with (a) 20 subsets and (b) 40 subsets, respectively. The dotted lines show the RMS differences using the standard OS algorithm with one subset.

Fig. 5. Shoulder scan: RMS differences between the reconstructed image $\mathbf{x}^{(k)}$ and the reference reconstruction $\mathbf{x}^*$ as a function of iteration using the proposed algorithm with different number of FISTA iterations n (1, 2, and 5) for solving the inner constrained denoising problem.

Fig. 6.

GE performance phantom: cropped images (displayed from 800 to 1200 HU) from the central transaxial plane of the initial FBP image $\mathbf{x}^{(0)}$ (left), the reference reconstruction $\mathbf{x}^*$ (center), and the reconstructed image using the proposed algorithm (OS-LALM-24-c-1) at the 30th iteration $\mathbf{x}^{(30)}$ (right).

Fig. 7. GE performance phantom: cropped difference images (displayed from -30 to 30 HU) from the central transaxial plane of $\mathbf{x}^{(30)} - \mathbf{x}^*$ using OS-based algorithms.

Fig. 8.

GE performance phantom: RMS differences between the reconstructed image $\mathbf{x}^{(k)}$ and the reference reconstruction $\mathbf{x}^*$ as a function of iteration using OS-based algorithms with 24 subsets. The dotted line shows the RMS differences using the standard OS algorithm with one subset.