



Published in final edited form as:

IEEE Trans Med Imaging. 2012 March ; 31(3): 677–688. doi:10.1109/TMI.2011.2175233.

A Splitting-Based Iterative Algorithm for Accelerated Statistical X-Ray CT Reconstruction

Sathish Ramani [Member, IEEE] and

Department of Electrical Engineering and Computer Science, University of Michigan, 1301 Beal Ave., Ann Arbor, MI 48109-2122, U.S.A

Jeffrey A. Fessler [Fellow, IEEE]

Department of Electrical Engineering and Computer Science, University of Michigan, 1301 Beal Ave., Ann Arbor, MI 48109-2122, U.S.A

Sathish Ramani: sramani@umich.edu; Jeffrey A. Fessler: fessler@umich.edu

Abstract

Statistical image reconstruction using penalized weighted least-squares (PWLS) criteria can improve image-quality in X-ray CT. However, the huge dynamic range of the statistical weights leads to a highly shift-variant inverse problem making it difficult to precondition and accelerate existing iterative algorithms that attack the statistical model directly. We propose to alleviate the problem by using a variable-splitting scheme that separates the shift-variant and (“nearly”) invariant components of the statistical data model and also decouples the regularization term. This leads to an equivalent constrained problem that we tackle using the classical method-of-multipliers framework with alternating minimization. The specific form of our splitting yields an alternating direction method of multipliers (ADMM) algorithm with an inner-step involving a “nearly” shift-invariant linear system that is suitable for FFT-based preconditioning using cone-type filters. The proposed method can efficiently handle a variety of convex regularization criteria including smooth edge-preserving regularizers and nonsmooth sparsity-promoting ones based on the ℓ_1 -norm and total variation. Numerical experiments with synthetic and real *in vivo* human data illustrate that cone-filter preconditioners accelerate the proposed ADMM resulting in fast convergence of ADMM compared to conventional (nonlinear conjugate gradient, ordered subsets) and state-of-the-art (MFISTA, split-Bregman) algorithms that are applicable for CT.

Index Terms

Statistical Image Reconstruction; Regularization; Iterative Algorithm; Method of Multipliers; Alternating Minimization

I. Introduction

Statistical image reconstruction methods in X-ray CT minimize a cost function consisting of a data-fidelity term that accommodates the measurement statistics and the geometry of the data-acquisition process, and a regularization term that reduces noise. For example, PWLS cost functions for X-ray CT use a (statistically) weighted quadratic data-fidelity term [1], [2] and can provide improved image-quality compared to filtered back-projection (FBP) [1], [2]. However, computation-intensive iterative methods are needed to minimize such cost

functions. This paper describes a new minimization algorithm that uses variable splitting to provide accelerated convergence.

Several types of iterative algorithms have been proposed for statistical image reconstruction in X-ray CT, including iterative coordinate descent (ICD) methods [1], block-based coordinate descent [3], ordered subsets (OS) algorithms based on separable quadratic surrogates (SQS) [4], [5] and (preconditioned) nonlinear conjugate gradient (NCG) methods [6]. For fast computation on multiprocessor computers, (P)NCG-type methods appear to be particularly amenable to efficient parallelization because they update all voxels simultaneously using all measurements.

Developing suitable preconditioners for (P)NCG is challenging for X-ray CT because the enormous dynamic range of the transmission data causes the Hessian of the statistical data-fidelity term to be highly shift-variant [6]. Clinthorne *et al.* [7] showed that for unweighted least-squares reconstruction, one can precondition the problem effectively using FFTs with a kind of cone filter. This cone filter amplifies high spatial frequencies, helping to accelerate convergence. But that cone filter is ineffective for (P)NCG in the PWLS case [6]. Delaney *et al.* [8] considered a very special type of shift-invariant weighting and also demonstrated accelerated convergence, but for low-dose X-ray CT the appropriate statistical weighting does not satisfy the assumptions in [8]. Shift-variant pre-conditioners based on multiple FFTs were proposed in [6] for 2D transmission tomography, but never became popular due to their complexity and were never investigated for 3D problems. Another way to introduce a cone filter is the iterative FBP approach [9], [10]. Initially these algorithms “converge” rapidly compared to (P)NCG methods, but typically they do not have any theoretical convergence properties and “too many” iterations lead to undesirably noisy images. Furthermore, it is unclear how to include regularization while ensuring convergence.

The challenges described above apply regardless of the form of the regularizer. Additional difficulties arise when one uses nonsmooth regularizers such as total variation (TV) [11] and sparsity-promoting ones based on the l_1 -norm [12]. These regularizers are not differentiable everywhere precluding optimization by conventional gradient-descent methods (e.g., NCG). Differentiable approximations (e.g., using “corner-rounding” [12, Sec. VI.A], [13, App. A]) can be employed, but even with such modifications the Hessian of the regularizer can have very high curvature leading to slow convergence of conventional gradient-descent methods [14]. While some state-of-the-art algorithms such as (M)FISTA [15], [16] and split-Bregman-type schemes (that split only the regularization term) [17], [18] are able to handle nonsmooth regularizers exactly (i.e., without corner rounding), when applied to X-ray CT, they must minimize a cost function that involves the original statistical data-fidelity term and are in turn hindered by the shift-variance of its Hessian (see Sections II-B, IV).

In this work, we propose to use a variable-splitting technique that not only decouples the regularization term in the spirit of [17], but also dissociates the statistical and geometrical components in the data-fidelity term. This forms the key feature of our approach that enables us to “isolate” the shift-variant element in the statistical data-fidelity term thereby alleviating the problem. Our splitting procedure uses auxiliary constraint variables to transform the original PWLS problem into an equivalent constrained optimization task that we solve using the classical method-of-multipliers [19], [20] and alternating direction optimization [21]–[23] frameworks. This leads to an alternating direction method of multipliers (ADMM) algorithm for solving the original PWLS problem that, apart from requiring simple operations (such as inverting a diagonal matrix, solving 1D denoising problems), involves the solving of a “nearly” shift-invariant linear system, which is amenable to FFT-based preconditioning using cone-type filters [7]. Experimental results with synthetic and real *in vivo* human data indicate that the proposed ADMM converges faster than conventional

(NCG and ordered subsets) and state-of-the-art (MFISTA and split-Bregman) methods, illustrating the efficacy of our splitting scheme and the potential of cone-filter preconditioners for accelerating the proposed ADMM. The proposed ADMM can also be used with a variety of convex regularization criteria (see Section VI-A) including smooth edge-preserving regularizers and nonsmooth ones such as TV and ℓ_1 -regularization.

The paper is organized as follows. In Section II, we mathematically formulate X-ray CT reconstruction as a PWLS problem and briefly discuss drawbacks of some existing algorithms for X-ray CT. Section III discusses the proposed splitting strategy and the development of the ADMM algorithm in detail. In Section IV, we compare our ADMM algorithm with the split-Bregman technique applied for CT, schematically. Section V is dedicated to numerical experiments and results, while Section VI discusses possible extensions of this work to 3D CT and other statistical models. Finally, we draw our conclusions in Section VII.

II. Statistical X-ray CT Reconstruction

A. Problem Formulation

For CT, an accurate statistical model for the data is quite complicated [24], [25] and is often replaced by a Gaussian approximation [1], [2] with a suitable diagonal weighting term \mathbf{W} whose components $\{w_j\}$ are inversely proportional to the measurement variances [1], [2]. We consider a penalized weighted least-squares (PWLS) formulation of statistical CT reconstruction [1]:

$$\mathbf{P0} : \arg \min_{\mathbf{x}} \{J(\mathbf{x}) \triangleq J_{\text{data}}(\mathbf{y}, \mathbf{A}\mathbf{x}) + \Psi(\mathbf{R}\mathbf{x})\}, \quad (1)$$

$$J_{\text{data}}(\mathbf{y}, \mathbf{A}\mathbf{x}) \triangleq \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_{\mathbf{W}}^2, \quad (2)$$

where \mathbf{y} is the $M \times 1$ data vector (log of transmission data), \mathbf{A} is the $M \times N$ system matrix, $\mathbf{A}\mathbf{x}$ represents the forward projection operation (e.g., line integrals), $\mathbf{W} = \text{diag}\{w_j\}$ is a $M \times M$ diagonal matrix consisting of statistical weights,¹ and $\|\mathbf{u}\|_{\mathbf{W}}^2 \triangleq \mathbf{u}^T \mathbf{W} \mathbf{u}$. We use a general family of regularizers of the form [12]

$$\Psi(\mathbf{R}\mathbf{x}) = \lambda \sum_r \kappa_r \Phi_r \left(\sum_{p=1}^P \|[\mathbf{R}_p \mathbf{x}]_r\|^m \right), \quad (3)$$

where $\lambda > 0$ is the regularization parameter, $\kappa_r > 0 \forall r$ are user-provided weights that govern the spatial resolution in the reconstructed output [26], Φ_r are potential functions, the $R \times N$ matrix $\mathbf{R} \triangleq [\mathbf{R}_1^T \dots \mathbf{R}_P^T]^T$ constitutes regularization operators \mathbf{R}_p (e.g., finite differences, frames, etc) of size $L \times N$, where $R = PL$. We concentrate on values of m and instances of Φ that result in a convex regularizer Ψ in (3).

The above general regularizer is in the ‘‘analysis’’ form [27], i.e., Ψ is specified as a function of the reconstructed image \mathbf{x} . The method proposed in this paper can also be easily extended to handle ‘‘synthesis’’ forms [27], e.g., by writing $\mathbf{x} = \mathbf{S}\boldsymbol{\theta}$ and considering $\mathcal{J}(\boldsymbol{\theta}) = J_{\text{data}}(\mathbf{y}, \mathbf{A}\mathbf{S}\boldsymbol{\theta}) + \Phi(\boldsymbol{\theta})$ in $\mathbf{P0}$, for some potential function Φ and synthesis operator \mathbf{S} . We focus on the

¹For simplicity we used $w_j = e^{-y_j}$ in our experiments.

analysis form (3) as it includes popular nonsmooth criteria such as TV (for $\Phi_r(x) = \sqrt{x}$ and $m = 2$), analysis ℓ_1 -wavelets (for $\Phi_r(x) = x$, $m = 1$) and a variety of smooth convex edge-preserving regularizers (e.g., Huber [28], [29], Fair [6], [30] etc).

B. Previous Approaches

Conventional gradient-descent methods, e.g., NCG, for $\mathbf{P0}$ depend on the Hessian of J_{data} : $\mathbf{H}_{\text{data}} = \mathbf{A}^T \mathbf{W} \mathbf{A}$, which is highly shift-variant in CT particularly due to the large dynamic range of \mathbf{W} . As a result, it becomes difficult to precondition and accelerate such methods [6]. Fessler *et al.* [6] directly attacked $\mathbf{P0}$ using NCG and proposed a shift-variant preconditioner to tackle \mathbf{H}_{data} . But their preconditioner is data-dependent and requires at least one pair of FFT-iFFT operations per NCG-iteration.

Iterative shrinkage-thresholding (IST) [31] and its variants ((M)FISTA [15], [16], and (M)TWIST [32]) that are applicable to $\mathbf{P0}$ depend on the Lipschitz constant L_{data} of $J_{\text{data}}(\mathbf{y}, \mathbf{Ax})$:

$$L_{\text{data}} = \sigma_{\max}\{\mathbf{A}^T \mathbf{W} \mathbf{A}\}, \quad (4)$$

where σ_{\max} represents the maximum eigenvalue. The convergence speed of these algorithms is primarily determined by (4): A large value of L_{data} results in small gradient steps [15, Sec. 1.1] leading to slow convergence. Since \mathbf{W} has a large dynamic range and due to the (approximately) $1/r$ -type decay of the elements of $\mathbf{A}^T \mathbf{A}$, L_{data} can be large for CT decreasing convergence speed of IST-type algorithms. Optimization transfer-based methods (e.g., [33, Sec. IV-B.1]) face a similar issue in that the surrogate functions end up having high curvature [5] due to \mathbf{W} , which again leads to small update-steps and slow convergence.

In summary, the weighting term \mathbf{W} , although crucial for improving reconstruction quality, poses a challenge for optimization. Compared to $\mathbf{A}^T \mathbf{W} \mathbf{A}$, the term $\mathbf{A}^T \mathbf{A}$ is “more” shift-invariant and is appropriate for preconditioning using cone filters. This property has been used to accelerate un-weighted least-squares reconstruction for tomographic image reconstruction [7]. Therefore, our idea to mitigate the shift-variance of \mathbf{H}_{data} is to untangle \mathbf{W} from \mathbf{H}_{data} thereby making the resulting problem “more” shift-invariant and suitable to circulant preconditioning. To do so, we adopt a variable-splitting strategy.

Variable splitting (VS) refers to the process of introducing auxiliary constraint variables to separate coupled components in the cost function [12], [17], [18], [34]–[42]. This procedure transforms the original minimization problem into an equivalent constrained optimization problem that can be effectively solved using classical constrained optimization schemes [19], [20]. The VS approach is appealing as it renders the resulting constrained problem tractable to alternating minimization schemes that decouple it in terms of the auxiliary variables and simplify optimization [12], [17], [18], [34], [36], [37], [39]–[42].

The VS approach has become popular recently for solving reconstruction problems in image processing [17], [34]–[37], MRI [12], [39], [40] and CT [18], [41], [42]. Many authors have focussed on splitting the regularization term [17], [18], [34], [37], [39]–[41] as it is hard to tackle in inverse problems (especially nonsmooth ones such as TV and ℓ_1 -regularization). Splitting the regularization term enables one to handle it exactly (i.e., without the need for “corner-rounding” [12, Sec. VI.A], [13, App. A] for nonsmooth criteria) via simple denoising problems [12], [17], [37], [42]. However, in PWLS problems for CT, the data-term adds to the complexity (as it leads to a shift-variant hessian \mathbf{H}_{data}) and therefore demands attention. So in this work, besides splitting the regularization term, we also split the data-term.

III. Proposed Method

A. Equivalent Constrained Optimization Problem

We introduce auxiliary constraint variables $\mathbf{u} \in \mathbb{R}^M$ and $\mathbf{v} \in \mathbb{R}^R$ and write $\mathbf{P0}$ as the following equivalent constrained problem:

$$\begin{aligned} \arg \min_{\mathbf{x}, \mathbf{u}, \mathbf{v}} \{f(\mathbf{u}, \mathbf{v}) \triangleq J_{\text{data}}(\mathbf{y}, \mathbf{u}) + \Psi(\mathbf{v})\} \\ \text{s.t. } \mathbf{u} = \mathbf{Ax}, \quad \mathbf{v} = \mathbf{Rx}, \end{aligned} \quad (5)$$

where \mathbf{u} separates the effect of \mathbf{W} on \mathbf{Ax} and \mathbf{v} splits the regularization term as in [17]. Afonso *et al.* [36] and Figueiredo *et al.* [38] have utilized data-term-splitting in the context of image restoration [36], [38] and reconstruction from partial Fourier observations [36]. However, our emphasis here is on CT reconstruction where \mathbf{u} plays an important role: It leads to a sub-problem that is “nearly” shift-invariant and suitable to preconditioning using cone filters [42] as explained in Section III-C.

In general, the proposed splitting strategy (5) can be applied to any PWLS problem of the form $\mathbf{P0}$ so as to exploit shift-invariant features in the data-model, e.g., deconvolution of blurred images corrupted with non-stationary noise.

Before proceeding, we rewrite (5) concisely as

$$\mathbf{P1} : \arg \min_{\mathbf{x}, \mathbf{z}} f(\mathbf{z}) \text{ s.t. } \mathbf{z} = \mathbf{Cx}, \quad (6)$$

where

$$\mathbf{z} \triangleq \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix}, \quad \mathbf{C} \triangleq \begin{bmatrix} \mathbf{A} \\ \mathbf{R} \end{bmatrix}. \quad (7)$$

Since $\mathbf{P1}$ is equivalent to $\mathbf{P0}$, solving $\mathbf{P1}$ for \mathbf{x} yields the desired reconstruction in (1).

B. Method of Multipliers

To solve $\mathbf{P1}$, we use the classical framework of the method of multipliers [19], [20] and construct an augmented Lagrangian (AL) function [12], [19], [20], [42]

$$\mathcal{L}(\mathbf{x}, \mathbf{z}, \gamma) \triangleq f(\mathbf{z}) + \gamma^\top (\mathbf{z} - \mathbf{Cx}) + \frac{\mu}{2} \|\mathbf{z} - \mathbf{Cx}\|_{\Lambda}^2 \quad (8)$$

that combines a multiplier term $\gamma^\top (\mathbf{z} - \mathbf{Cx})$ with Lagrange multiplier

$\gamma \triangleq [\gamma_{\mathbf{u}}^\top \quad \gamma_{\mathbf{v}}^\top]^\top \in \mathbb{R}^{M+R}$ and a quadratic penalty term $\frac{\mu}{2} \|\mathbf{z} - \mathbf{Cx}\|_{\Lambda}^2$, where $\mu > 0$ is the AL penalty parameter and $\Lambda > \mathbf{0}$ is a symmetric weighting matrix. The multiplier term can be absorbed into the penalty term in (8) (by completing the square) for ease of manipulation leading to

$$\mathcal{L}(\mathbf{x}, \mathbf{z}, \eta) \triangleq f(\mathbf{z}) + \frac{\mu}{2} \|\mathbf{z} - \mathbf{Cx} - \eta\|_{\Lambda}^2 + c_{\eta}, \quad (9)$$

where $\eta \triangleq [\eta_{\mathbf{u}}^\top \quad \eta_{\mathbf{v}}^\top]^\top = -\frac{1}{\mu} \Lambda^{-1} \gamma$ and $c_{\eta} \triangleq -\frac{\mu}{2} \|\eta\|_{\Lambda}^2$ is a constant independent of \mathbf{x} and \mathbf{z} . Unlike standard approaches [36], [38] that set $\Lambda = \mathbf{I}_{M+R}$, we propose to use

$$\Lambda = \begin{bmatrix} \mathbf{I}_M & \mathbf{0} \\ \mathbf{0} & \nu \mathbf{I}_R \end{bmatrix}, \quad (10)$$

where $\nu > 0$. This is crucial in CT because the elements of \mathbf{A} and \mathbf{R} can differ by several orders of magnitude and it is imperative to balance them to avoid numerical instabilities in the resulting algorithm and to achieve faster convergence [42].

The classical AL scheme for solving $\mathbf{P1}$ alternates between a joint-minimization step and an update step [12, Sec. III]:

$$(\mathbf{x}^{(j+1)}, \mathbf{z}^{(j+1)}) = \arg \min_{\mathbf{x}, \mathbf{z}} \mathcal{L}(\mathbf{x}, \mathbf{z}, \eta^{(j)}), \quad (11)$$

$$\eta^{(j+1)} = \eta^{(j)} - (\mathbf{z}^{(j+1)} - \mathbf{C}\mathbf{x}^{(j+1)}), \quad (12)$$

respectively. Unlike pure penalty methods, remarkably, the AL formalism does not require increasing $\mu \rightarrow \infty$ to ensure convergence of (11)–(12) to a solution of $\mathbf{P1}$ [19].

C. Alternating Direction Minimization

It is numerically appealing to replace the more difficult joint-minimization step (11) by alternating direction optimization that decouples (11) as [21]–[23]

$$\mathbf{x}^{(j+1)} = \arg \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{z}^{(j)}, \eta^{(j)}), \quad (13)$$

$$\mathbf{z}^{(j+1)} = \arg \min_{\mathbf{z}} \mathcal{L}(\mathbf{x}^{(j+1)}, \mathbf{z}, \eta^{(j)}). \quad (14)$$

Thus, at the j th iteration, instead of (11)–(12), we perform (ignoring constant terms)

$$\mathbf{x}^{(j+1)} \stackrel{\varepsilon_x^{(j)}}{\approx} \arg \min_{\mathbf{x}} \|\mathbf{z}^{(j)} - \mathbf{C}\mathbf{x} - \eta^{(j)}\|_{\Lambda}^2, \quad (15)$$

$$\mathbf{z}^{(j+1)} \stackrel{\varepsilon_z^{(j)}}{\approx} \arg \min_{\mathbf{z}} f(\mathbf{z}) + \|\mathbf{z} - \mathbf{C}\mathbf{x}^{(j+1)} - \eta^{(j)}\|_{\Lambda}^2, \quad (16)$$

$$\eta^{(j+1)} = \eta^{(j)} - (\mathbf{z}^{(j+1)} - \mathbf{C}\mathbf{x}^{(j+1)}), \quad (17)$$

where we write $\mathbf{x} \stackrel{\varepsilon}{\approx} \mathbf{x}^*$ to mean that $\|\mathbf{x} - \mathbf{x}^*\|_2 \leq \varepsilon$, i.e., we allow for inexact updates in (15)–(16) in the spirit of [43]. Although (15)–(16) is an approximation to (11), the following theorem adapted from [43, Theorem 8] to $\mathbf{P1}$ guarantees convergence of (15)–(17) to a solution of (5) (and $\mathbf{P0}$).

Theorem 1—Consider $\mathbf{P1}$ in (6) where f is closed, proper, convex² and \mathbf{C} has full column-rank. Let $\boldsymbol{\eta}^{(0)} \in \mathbb{R}^{M+R}$, $\mu > 0$,

²A convex function h is closed if and only if it is lower semi-continuous (LSC) [44, pp. 51–52] and is proper if $h(\mathbf{x}) < +\infty$ for at least one \mathbf{x} and $h(\mathbf{x}) > -\infty \forall \mathbf{x}$ [44, p. 24]. It can be shown that the convex functions J_{data} , Ψ (for a variety of regularizers such as TV and ℓ_1 -regularization), and their sum, $f(5)$, are LSC and proper [38].

$$\sum_j \varepsilon_{\mathbf{x}}^{(j)} < \infty, \quad \text{and} \quad \sum_j \varepsilon_{\mathbf{z}}^{(j)} < \infty. \quad (18)$$

If **P1** has a solution $(\mathbf{x}^*, \mathbf{z}^*)$, then the sequence of updates $\{(\mathbf{x}^{(j)}, \mathbf{z}^{(j)})\}_j$ generated by (15)–(17) converges to $(\mathbf{x}^*, \mathbf{z}^*)$. If **P1** has no solution, then at least one of the sequences $\{(\mathbf{x}^{(j)}, \mathbf{z}^{(j)})\}_j$ or $\{\eta^{(j)}\}_j$ diverges.

The result of Eckstein *et al.* [43, Theorem 8] uses an AL function with $\mathbf{\Lambda} = \mathbf{I}$, so we apply [43, Theorem 8] to (9) with $\mathbf{\Lambda}$ in (10) through a simple change of variables.³ For CT, it can be readily ensured that **C** has full column-rank for a variety of regularization operators **R**. In the sequel, we explain how to perform the minimizations in (15)–(16).

Firstly, we see that due to the structure of $f(\mathbf{z})$ and **C**, (16) further dissociates into the following:

$$\mathbf{u}^{(j+1)} \stackrel{\varepsilon_{\mathbf{u}}^{(j)}}{=} \arg \min_{\mathbf{u}} \left\{ J_{\text{data}}(\mathbf{y}, \mathbf{u}) + \frac{\mu}{2} \|\mathbf{u} - \mathbf{A}\mathbf{x}^{(j+1)} - \eta_{\mathbf{u}}^{(j)}\|_2^2 \right\}, \quad (19)$$

$$\mathbf{v}^{(j+1)} \stackrel{\varepsilon_{\mathbf{v}}^{(j)}}{=} \arg \min_{\mathbf{v}} \left\{ \Psi(\mathbf{v}) + \frac{\mu\nu}{2} \|\mathbf{v} - \mathbf{R}\mathbf{x}^{(j+1)} - \eta_{\mathbf{v}}^{(j)}\|_2^2 \right\}. \quad (20)$$

These sub-problems are independent of each other and can therefore be solved simultaneously, where $[\varepsilon_{\mathbf{z}}^{(j)}]^2 = [\varepsilon_{\mathbf{u}}^{(j)}]^2 + [\varepsilon_{\mathbf{v}}^{(j)}]^2 \forall j$. Sub-problem (19) is quadratic and has a closed form solution:

$$\mathbf{u}^{(j+1)} = \mathbf{D}_{\mu}^{-1} (\mathbf{W}\mathbf{y} + \mu(\mathbf{A}\mathbf{x}^{(j+1)} + \eta_{\mathbf{u}}^{(j)})), \quad (21)$$

where $\mathbf{D}_{\mu} \triangleq (\mathbf{W} + \mu\mathbf{I}_M)$. Since **W** is diagonal, \mathbf{D}_{μ} can be inverted exactly, so that $\varepsilon_{\mathbf{u}}^{(j)} = 0$ in (19) $\forall j$.

Minimization w.r.t. **v** (20) corresponds to a denoising problem that can be solved efficiently and/or exactly for a variety of instances of (3) including TV: This has been elucidated by many authors [12], [17], [36]–[40], [45], e.g., the techniques developed in [12, Sec. IV.A-2 – IV.A-6] can be directly applied to (20). For brevity, we concentrate on two particular instances of (3) and solve (20) exactly so that $\varepsilon_{\mathbf{v}}^{(j)} = 0$ in (20) $\forall j$.

- Analysis ℓ_1 -regularization [$m = 1, \Phi_r(x) = x \forall r$ in (3)]:

$$\Psi(\mathbf{R}\mathbf{x}) = \lambda \sum_{r=1}^R \kappa_r \|\mathbf{R}\mathbf{x}\|_r \quad (22)$$

with the shift-invariant Haar wavelet transform (excluding the approximation level) for **R**, which is a sparsity-promoting criterion [12], [27], [36], [38].

- Smooth edge-preserving regularization [$P = 1, m = 1, \Phi_r = \Phi_{\text{FP}} \forall r$ in (3)]:

³Writing $\mathbf{z}_0 = \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{z}$, $\eta_0 = \mathbf{\Lambda}^{\frac{1}{2}} \eta$, and $\mathbf{M} = \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{C}$, it is easy to see that (15)–(17) solve the constrained problem $\arg \min_{\mathbf{x}, \mathbf{z}_0} f(\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{z}_0)$ s.t. $\mathbf{z}_0 = \mathbf{M}\mathbf{x}$ that is equivalent to **P0** using the AL function $\mathcal{L}_0(\mathbf{x}, \mathbf{z}_0, \eta_0) = f(\mathbf{\Lambda}^{-\frac{1}{2}} \mathbf{z}_0) + \frac{\mu}{2} \|\mathbf{z}_0 - \mathbf{M}\mathbf{x} - \eta_0\|_2^2$ with an unweighted penalty term.

$$\Psi(\mathbf{R}\mathbf{x}) = \lambda \sum_{r=1}^R \kappa_r \Phi_{\text{FP}}(|[\mathbf{R}\mathbf{x}]_r|) \quad (23)$$

using the Fair potential $\Phi_{\text{FP}}(x) = x/\delta - \log(1+x/\delta)$ with $\delta > 0$ [30] (also the smoothed Laplace function in [45, Eq. 4.11]) and finite-differences for \mathbf{R} . This regularizer ensures a unique solution to $\mathbf{P}\mathbf{0}$ as Φ_{FP} is strictly convex. It has also been successfully applied to PWLS problems in tomography [6].

For these regularizers, (20) separates into R 1D minimization problems in terms of the components $\{v_r\}_{r=1}^R$ of \mathbf{v} :

$$v_r^{(j+1)} = \arg \min_{v_r} \Psi(v_r) + \frac{\mu\nu}{2} (v_r - \varrho_r^{(j)})^2, \quad (24)$$

where $\varrho_r^{(j)}$ is the r th component of $\varrho^{(j)} \triangleq \mathbf{R}\mathbf{x}^{(j+1)} + \eta_{\mathbf{v}}^{(j)}$. For (22), the solution of (24) is given by the shrinkage rule⁴ [46]

$$v_r^{(j+1)} = \text{sign}\{\varrho_r^{(j)}\} \max\left(|\varrho_r^{(j)}| - \frac{\lambda\kappa_r}{\mu\nu}, 0\right). \quad (25)$$

For (23), (24) leads to a quadratic equation in v_r [45, Eq. 4.13] that yields

$$v_r^{(j+1)} = \text{sign}\{\varrho_r^{(j)}\} \frac{\zeta_r^{(j)} + \sqrt{(\zeta_r^{(j)})^2 + 4\delta|\varrho_r^{(j)}|}}{2}, \quad (26)$$

where $\zeta_r^{(j)} \triangleq |\varrho_r^{(j)}| - \delta - \lambda\kappa_r/(\delta\mu\nu)$.

Having addressed (16), we now consider (15) which can be easily solved analytically:

$$\mathbf{x}^{(j+1)\star} = \mathbf{G}_{\mathbf{v}}^{-1} (\mathbf{A}^{\top} (\mathbf{u}^{(j)} - \eta_{\mathbf{u}}^{(j)}) + \nu \mathbf{R}^{\top} (\mathbf{v}^{(j)} - \eta_{\mathbf{v}}^{(j)})), \quad (27)$$

where $\mathbf{x}^{(j+1)\star}$ represents the exact solution to (15) and

$$\mathbf{G}_{\mathbf{v}} \triangleq \mathbf{C}^{\top} \mathbf{\Lambda} \mathbf{C} = (\mathbf{A}^{\top} \mathbf{A} + \nu \mathbf{R}^{\top} \mathbf{R}) \quad (28)$$

is non-singular because $\mathbf{\Lambda} > \mathbf{0}$ and \mathbf{R} is chosen so that \mathbf{C} has full column-rank. Although (27) is an exact analytical solution, the enormous size of $\mathbf{G}_{\mathbf{v}}$ for CT makes it impossible to store and “invert” $\mathbf{G}_{\mathbf{v}}$ exactly. So we propose to use the conjugate gradient (CG) method for (27) and obtain an approximate update $\mathbf{x}^{(j+1)} \stackrel{\varepsilon_{\mathbf{x}}^{(j)}}{\approx} \mathbf{x}^{(j+1)\star}$. Since $\mathbf{G}_{\mathbf{v}}$ is non-singular, we have that

$$\varepsilon_{\mathbf{x}}^{(j)} \leq \frac{\rho_{\mathbf{x}}^{(j)}}{\sigma_{\min}\{\mathbf{G}_{\mathbf{v}}\}}, \quad (29)$$

⁴An analytical update formula similar to (25) is available for the TV regularizer that is based on a vector shrinkage-rule, see e.g., [12, Sec. IV.A-6].

where $\rho_x^{(j)}$ is the corresponding residue and $\sigma_{\min}\{\mathbf{G}_\nu\} > 0$ is the minimum eigenvalue of \mathbf{G}_ν that depends only on \mathbf{A} and \mathbf{R} and can be precomputed e.g., using stochastic techniques [47] or the Power method.⁵ Therefore, using (29), one can monitor $\varepsilon_x^{(j)}$ in the CG-loop and design a suitable stopping rule to satisfy (18).

D. Preconditioning Using Cone Filter

We see that \mathbf{G}_ν contains $\mathbf{A}^T\mathbf{A}$, which is “nearly” shift-invariant, so for shift-invariant⁶ $\mathbf{R}^T\mathbf{R}$, \mathbf{G}_ν is amenable to preconditioning using suitable cone filters [6], [7]. We constructed a circulant matrix $\tilde{\mathbf{G}}_\nu$ from the central column of \mathbf{G}_ν :

$$\tilde{\mathbf{G}}_\nu = \text{circ}\{\mathbf{G}_\nu \mathbf{e}_c\}, \quad (30)$$

and used its inverse, $\tilde{\mathbf{G}}_\nu^{-1}$, as the preconditioner, where \mathbf{e}_c is a standard basis vector of \mathbb{R}^N corresponding to the center pixel of the image and $\text{circ}\{\mathbf{a}\}$ represents the construction of a circulant matrix from a vector \mathbf{a} . The proposed preconditioner $\tilde{\mathbf{G}}_\nu^{-1}$ corresponds to a cone-type filter that amplifies high spatial frequencies and accelerates convergence of both the CG-loop for (27) and the overall ADMM scheme as demonstrated in Section V.

Implementing $\tilde{\mathbf{G}}_\nu^{-1}$ requires only one FFT-iFFT per CG iteration and its construction⁷ uses a product with $\mathbf{R}^T\mathbf{R}$ and only one forward-backward projection that can be performed offline as $\tilde{\mathbf{G}}_\nu$ is independent of \mathbf{W} . In our experiments, we applied at most two preconditioned CG (PCG) iterations with warm starting [12] and found that $\rho_x^{(j)}$ decreased sufficiently rapidly. Based on (15)–(29), we present our algorithm in Fig. 1 for solving $\mathbf{P1}$ (and thus $\mathbf{P0}$). In principle, Steps 4 and 5 of ADMM may be executed in parallel as they are independent of each other, but in our implementation, we chose to execute all the steps sequentially for simplicity.

E. Selection of μ and ν

The parameters μ and ν do not affect the solution of $\mathbf{P1}$, but only regulate the convergence speed of the proposed ADMM [12], [35, Sec. 4.4]. In general, choosing appropriate values for AL penalty parameters (such as μ and ν) is a nontrivial and application-dependent task. Several empirical rules have been put forth by many authors for setting AL penalty parameters (to obtain good convergence speeds for AL-based iterative reconstruction schemes) in many applications, see e.g., [37], [38] for image restoration, [17] for denoising and compressed-sensing MRI, and [12], [39] for parallel MRI reconstruction.

In this paper, Step 3 is the only inexact step of the proposed ADMM. So the computational speed of ADMM is primarily determined by how efficiently (27) is solved, which in turn is governed by ν . We use an empirical rule for selecting ν that is based on [17]: Since ν balances $\mathbf{A}^T\mathbf{A}$ and $\mathbf{R}^T\mathbf{R}$ in \mathbf{G}_ν that have disjoint non-trivial null-spaces, the condition number $\kappa(\mathbf{G}_\nu)$ of \mathbf{G}_ν exhibits a minimum for some $\nu_{\min} > 0$: $\nu_{\min} = \arg \min_\nu \kappa(\mathbf{G}_\nu)$. It was suggested in [17] to use this property to choose AL penalty parameters to ensure quick convergence of the CG-algorithm for solving a linear system such as (27). For \mathbf{A} implemented using the distance-driven (DD) projector [49] and \mathbf{R} in (22)–(23), $\nu_{\min} \approx 10^5$,

⁵Since the Power method (PM) iteratively estimates the maximum eigenvalue (in absolute magnitude) of a matrix [48, p. 488], an estimate $\hat{\sigma}_{\max}\{\mathbf{G}_\nu\}$ of $\sigma_{\max}\{\mathbf{G}_\nu\}$ is first computed by applying PM on \mathbf{G}_ν . Next, applying PM on $\mathbf{K} \triangleq \mathbf{G}_\nu - \hat{\sigma}_{\max}\{\mathbf{G}_\nu\}\mathbf{I}_N$ yields $|\hat{\sigma}_{\min}\{\mathbf{G}_\nu\} - \hat{\sigma}_{\max}\{\mathbf{G}_\nu\}|$ (as $\sigma_{\min}\{\mathbf{G}_\nu\} - \sigma_{\max}\{\mathbf{G}_\nu\}$ is the largest eigenvalue of \mathbf{K} in absolute magnitude) from which $\sigma_{\min}\{\mathbf{G}_\nu\}$ can be easily obtained.

⁶The matrix $\mathbf{R}^T\mathbf{R}$ is circulant when periodic boundary conditions are used for \mathbf{R} in (22)–(23).

⁷We only store the frequency response corresponding to $\tilde{\mathbf{G}}_\nu^{-1}$ to save memory.

which yielded a very small $(\lambda/\mu\nu)$ in (20) and subsequently resulted in slow convergence of ADMM in our experiments. On the other extreme, setting $\nu = 1$ (corresponding to the standard case of $\mathbf{A} = \mathbf{I}_{M+R}$) yields a poorly conditioned \mathbf{G}_ν that was not favorable either.

Based on our experience with 2D CT experiments, we found the empirical rule⁸

$\nu_{emp} = \frac{1}{100} \arg \min_{\nu_0} \kappa(\tilde{\mathbf{G}}_{\nu_0})$ to yield good overall convergence speeds for ADMM, where $\tilde{\mathbf{G}}_\nu$ is the circulant matrix in (30). We also observed that ADMM was slightly more robust to the choice of μ than ν . We selected $\mu = \text{median}\{w_i\}$ to avoid outliers in \mathbf{W} ; this yielded a well-conditioned \mathbf{D}_μ (with $\kappa\{\mathbf{D}_\mu\} \in [10, 40]$) that improved the numerical stability of ADMM.

IV. Comparison with the Split-Bregman Approach

The split-Bregman (SB) method [17] uses constraint variables to split the regularization term alone. For (1), this corresponds to using only $\mathbf{v} = \mathbf{R}\mathbf{x}$ which leads to following equivalent constrained problem

$$\arg \min_{\mathbf{x}, \mathbf{v}} \{f_{\text{SB}}(\mathbf{x}, \mathbf{v}) \triangleq J_{\text{data}}(\mathbf{y}, \mathbf{A}\mathbf{x}) + \Psi(\mathbf{v})\} \quad (31)$$

$$\text{s.t. } \mathbf{v} = \mathbf{R}\mathbf{x}. \quad (32)$$

This type of splitting has been investigated for CT reconstruction in [18], [41]. Applying the Bregman iterations [17, Eq. 3.7–3.8] with alternating minimization [17, Sec. 3.1] to (31) yields the following SB scheme:⁹

$$\mathbf{x}^{(j+1)} = \arg \min_{\mathbf{x}} \left\{ \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_w^2 + \frac{\mu}{2} \|\mathbf{v}^{(j)} - \mathbf{R}\mathbf{x} - \eta_{\mathbf{v}}^{(j)}\|_2^2 \right\}, \quad (33)$$

$$\mathbf{v}^{(j+1)} = \arg \min_{\mathbf{v}} \left\{ \Psi(\mathbf{v}) + \frac{\mu}{2} \|\mathbf{v} - \mathbf{R}\mathbf{x}^{(j+1)} - \eta_{\mathbf{v}}^{(j)}\|_2^2 \right\}, \quad (34)$$

$$\eta_{\mathbf{v}}^{(j+1)} = \eta_{\mathbf{v}}^{(j)} - (\mathbf{v}^{(j+1)} - \mathbf{R}\mathbf{x}^{(j+1)}). \quad (35)$$

The minimization in (34) is same as that in (20), so the techniques described for (20) apply to (34) as well. The main difference between the proposed method (15)–(17) and the SB scheme (33)–(35) is in the way \mathbf{x} is updated. The minimization in (33) leads to

$$\mathbf{x}^{(j+1)\star} = \mathbf{B}_\mu^{-1} [\mathbf{A}^\top \mathbf{W} \mathbf{y} + \mu \mathbf{R}^\top (\mathbf{v}^{(j)} - \eta_{\mathbf{v}}^{(j)})], \quad (36)$$

where $\mathbf{x}^{(j+1)\star}$ represents the exact solution to (36) and

$$\mathbf{B}_\mu \triangleq (\mathbf{A}^\top \mathbf{W} \mathbf{A} + \mu \mathbf{R}^\top \mathbf{R}). \quad (37)$$

⁸It would be ideal to consider $\kappa(\mathbf{G}_\nu)$ instead of $\kappa(\tilde{\mathbf{G}}_\nu)$ for selecting ν_{emp} , but estimating $\kappa(\mathbf{G}_\nu)$ (e.g., using the Power method) for a given ν for CT is computationally expensive (ignoring the fact that it is independent of \mathbf{W} and could be computed offline). But as \mathbf{G}_ν is approximately shift-invariant, $\kappa(\tilde{\mathbf{G}}_\nu) \approx \kappa(\mathbf{G}_\nu)$, which leads to ν_{emp} .

⁹Theorem 1 may not be applicable to the SB scheme (33)–(35) as the constraint matrix, which is simply \mathbf{R} in this case, usually does not have full column-rank. Convergence of SB-type schemes are studied in [17].

The matrix \mathbf{B}_μ contains the shift-variant component \mathbf{W} that makes standard preconditioners (including cone filters) less effective for CG-based solving of (36). Nevertheless, we used PCG for (36) with a circulant preconditioner $\tilde{\mathbf{G}}_\mu^{-1}$ (obtained by setting $\mu \equiv \nu$ in (30)) in our implementation of the SB scheme and found that it improved upon the standard CG method for (36). We selected¹⁰ $\mu = L_{\text{data}}/(100 \sigma_{\max}\{\mathbf{R}^T \mathbf{R}\})$ for SB¹¹ (see (4) for definitions of L_{data} and σ_{\max}). This choice is motivated by the discussion pertaining to ν in Section III-E.

In principle, it is possible to construct a shift-variant preconditioner for \mathbf{B}_μ in the spirit of [6], but such a preconditioner would invariably be data-dependent and may be computationally involved. Our approach (15)–(17) provides a simple and effective alternative using an extra constraint variable \mathbf{u} in (5): Compared to the SB scheme (33)–(35), our method requires only an extra trivial operation of inverting a diagonal matrix \mathbf{D}_μ in (21).

V. Experimental Results

We present numerical results for 2D CT reconstruction from simulated NCAT phantom data and *in vivo* human head data. The proposed ADMM is also applicable, in principle, to 3D CT reconstruction (see Section VI-A). We implemented the following algorithms in Matlab and conducted the experiments on a quad-core PC with 3.07 GHz Intel Xeon processors and 12 GB RAM.

- **NCG- n** : unpreconditioned nonlinear conjugate gradient algorithm with n line-search iterations that monotonically decrease the cost function J [6],
- **MFISTA- n** : Monotone Fast Iterative Shrinkage-Thresholding Algorithm [16] with n iterations for solving auxiliary denoising sub-problems similar to [16, Eq. 3.13],
- **OS- n** : Ordered subsets algorithm [5] with n blocks,
- **SB-(P)CG- n** : Split-Bregman scheme from Section IV with n (P)CG iterations for solving (36),
- **ADMM-(P)CG- n** : Proposed ADMM with n (P)CG iterations for solving (27).

MFISTA is a state-of-the-art method developed by Beck *et al.* [16] for image restoration that is readily applicable to $\mathbf{P0}$ with the Lipschitz constant L_{data} in (4). Beck *et al.* [15] also proposed a back-tracking strategy that does not require explicit computation of L_{data} , but we chose to estimate and use¹¹ L_{data} both for ease of implementation and because it is the smallest possible value [15, Ex. 2.2] that yields the fastest convergence for MFISTA. We applied the Chambolle-type method [50] for the inner-step (i.e., computing the proximal map [16, Eq. 3.13]) of MFISTA as that does not require smoothing of (“corners” of) ℓ_1 -regularizers such as (22).

Since our task is to solve $\mathbf{P0}$, we fixed the cost function J (that led to a visually appealing reconstruction) and focussed on the convergence speed of the algorithms. We quantified the convergence rate using the normalized ℓ_2 -distance between $\mathbf{x}^{(j)}$ and \mathbf{x}^\star :

$$\xi(j) = 20 \log_{10} \left(\frac{\|\mathbf{x}^{(j)} - \mathbf{x}^\star\|_2}{\|\mathbf{x}^\star\|_2} \right), \quad (38)$$

¹⁰Similar to ν_{\min} , one could consider $\mu_{\min} = \frac{1}{100} \arg \min_{\mu} \kappa(\mathbf{B}_\mu)$ for SB, but estimating $\kappa(\mathbf{B}_\mu)$ is impractical mainly due to its dependence on \mathbf{W} . We chose to use the above rational-form for μ , which yields a rough estimate of μ_{\min} .

¹¹We estimated L_{data} using the Power method applied to $\mathbf{A}^T \mathbf{W} \mathbf{A}$. Since L_{data} is \mathbf{W} -dependent, its use is less appealing for practical applications in CT.

where \mathbf{x}^\star is a solution to $\mathbf{P}\mathbf{0}$ obtained numerically by running one of the above algorithms as described next. Since the algorithms have different computation load per (outer) iteration, we evaluated $\xi(j)$ as a function of algorithm run-time¹² t_j , i.e., the time elapsed from start until iteration j . We also plot $\xi(j)$ as a function of the iteration index j for completeness. We used the DD-projector [49] (with 8 threads) for implementing matrix-vector products such as $\mathbf{A}\mathbf{x}$, $\mathbf{A}^\top \mathbf{u}$ and initialized all the algorithms with the image reconstructed using FBP (with the ramp filter) in all experiments.

Products with \mathbf{A} and \mathbf{A}^\top (corresponding to forward- and back-projections, respectively) are computation intensive in CT reconstruction problems and dominate the overall computation load of a reconstruction algorithm.¹³ NCG and MFISTA both require only one product with \mathbf{A} and \mathbf{A}^\top , respectively, per iteration. The OS method breaks products with \mathbf{A} and \mathbf{A}^\top in terms of block-rows of \mathbf{A} and block-columns of \mathbf{A}^\top , respectively, and cycles through each block once per every iteration, so effectively, OS also requires only one product with \mathbf{A} and \mathbf{A}^\top , respectively, per iteration. However, for each block, the OS method demands the evaluation of the gradient of the regularization term that increases computation time per iteration as indicated in Table I. For the SB scheme, we employ (P)CG for “inverting” \mathbf{B}_μ (that depends on $\mathbf{A}^\top \mathbf{W} \mathbf{A}$) in (36), so **SB-(P)CG- n** requires n products with \mathbf{A} and \mathbf{A}^\top , respectively, per iteration of (33)-(35). In the case of ADMM, we apply (P)CG at Step 3 (see Fig.1) for “inverting” \mathbf{G}_v in (27), but that step also requires a product with \mathbf{A}^\top in the RHS of (27), so overall **ADMM-(P)CG- n** uses n products with \mathbf{A} and $n+1$ products with \mathbf{A}^\top per iteration of Steps 3–7 in Fig.1. Table I summarizes this discussion and also shows the mean computation time per iteration (averaged over 10 iterations) of the above algorithms. Although the proposed ADMM-(PCG) requires more forward- and back-projections per iteration (and accordingly exhibits higher computation time per iteration) compared to other algorithms (with the exception of the OS method) in Table I, we demonstrate in the sequel that it converges faster in terms of algorithm runtime.

A. Simulation with the NCAT Phantom

We used a 1024×1024 2D slice of the NCAT phantom [51] and numerically generated a 888×984 -view noisy sinogram with GE LightSpeed fan-beam geometry [52] corresponding to a monoenergetic source with 2.5×10^4 incident photons per ray and no background events. We used the ℓ_1 -regularization in (22) with $\kappa_r = \omega_{(r \bmod N)}$, where

$\omega_n \triangleq \sqrt{[\mathbf{A}^\top \mathbf{W} \mathbf{1}]_n / [\mathbf{A}^\top \mathbf{1}]_n}$ is based on [26]. We reconstructed 512×512 images over a FOV of 65 cm; we obtained \mathbf{x}^\star by running 5000 iterations of **MFISTA-25** as it does not require “corner-rounding” and is therefore guaranteed to converge to a solution of $\mathbf{P}\mathbf{0}$. NCG cannot directly handle nonsmooth criteria such as (22) without smoothing it [13, App. A], so we used a smoothing value of 10^{-6} cm^{-1} . The FBP reconstructions in Figs. 2b, 2c corresponding to the ramp and Hanning filters, respectively, are either noisy or blurred and streaked with artifacts. The ℓ_1 -regularized reconstruction \mathbf{x}^\star in Fig. 2d preserves image features and has lower RMSE than both FBP outputs.

We plot $\xi(j)$ for various algorithms as a function of time in Fig. 3a. The SB-CG scheme appears to converge the slowest, while SB-PCG is faster indicating that the circulant preconditioner $\tilde{\mathbf{G}}_\mu^{-1}$ provides a moderate acceleration of CG for (36). MFISTA is slower than most of the algorithms for the reason explained in Section II. The CG-version of the proposed method, ADMM-CG, is slightly faster than MFISTA and SB-(P)CG but slower

¹²We excluded the computation time spent on estimating L_{data} for MFISTA in the plots. Even with this “unfair advantage” the ADMM method was much faster than MFISTA.

¹³NCG, MFISTA, OS and SB require the evaluation of $\mathbf{A}^\top \mathbf{W} \mathbf{y}$ (e.g., see RHS of (36) for the SB scheme), but this quantity needs to be computed only once, so we ignore this computation need for these schemes.

than NCG. The preconditioned version ADMM-PCG is the fastest among all algorithms illustrating that the cone-filter preconditioner $\tilde{\mathbf{G}}_v^{-1}$ is very effective in accelerating convergence of CG applied to (27) and ADMM-PCG. This is also corroborated by Fig. 3c where for a given number of iterations, ADMM-PCG produces a reconstruction that is closest to \mathbf{x}^\star in terms of $\xi(j)$. Figs. 3b, 3d further substantiate the reconstruction speed-up of ADMM-PCG over other methods, where (both in terms of algorithm run-time and number of iterations) it rapidly leads to a RMSE-value close to $\text{RMSE}(\mathbf{x}^\star)$.

B. Experiments with a *in vivo* Human Head Data-set

In this experiment, we used a *in vivo* human head data-set acquired with a GE scanner using 120 kVp source potential and 585 mA tube current with 0.6 s rotation. We reconstructed a 1024×1024 2D slice with 50 cm FOV and 0.625 mm thickness from a 888×984 -view sinogram. For Ψ in (1), we used the strictly convex regularizer (23) (with $\delta = 10$ HU) that guarantees a unique solution \mathbf{x}^\star to $\mathbf{P}\mathbf{0}$. As NCG generally had faster convergence than MFISTA in our experiments, we obtained \mathbf{x}^\star by running 5000 iterations of **NCG-10**. Fig. 4 shows the reconstruction results for this experiment. The regularized solution \mathbf{x}^\star in Fig. 4c has reduced noise and better preserves the anatomical features compared to the FBP reconstructions in Figs. 4a, 4b obtained using the ramp and Hanning filters, respectively.

Figs. 5a–d plot $\xi(j)$ as a function of t_j and iteration index j for all algorithms considered in this work. Here, we additionally compare the standard OS algorithm (that is not guaranteed to converge) in Figs. 5b, 5d, where we used the implementation from [53] available currently for regularization criteria such as (23). The OS algorithm is faster than all algorithms (including ADMM-PCG) for the first few iterations but it does not converge to the minimizer as expected. In practice, it may be advantageous to run a few iterations of OS and use its output to initialize a more sophisticated iterative algorithm. Figs. 5a, 5b indicate that the convergence trends for MFISTA, NCG, SB-(P)CG and ADMM-CG are generally similar to those in Figs. 3a, 3b. ADMM-PCG again provides notable reconstruction speed-up compared to all algorithms. This substantiates the potential of the cone-filter preconditioner (30) for the proposed ADMM and also demonstrates the benefit of our splitting scheme (5).

VI. Discussion

A. Memory Requirements

Splitting-based algorithms simplify optimization at the expense of manipulating and storing auxiliary constraint variables (and corresponding Lagrange multipliers in the AL formalism) and therefore have additional memory requirements compared to conventional algorithms such as NCG. Although this does not pose much concern for 2D reconstruction problems, it can represent a significant memory overhead for 3D problems. Specifically, the SB (Section IV) and the proposed ADMM¹⁴ (in Fig. 1) schemes use the constraint $\mathbf{v} = \mathbf{R}\mathbf{x}$ that requires the storage of $2P$ vectors (\mathbf{v} and $\boldsymbol{\eta}_v$) of size $L \times 1$. For instance, typically, the size of an image-volume in 3D CT is $N = 512 \times 512 \times 512$ (≈ 1 GB of memory when stored in double-precision format in Matlab). Then, for finite-differences with $P = 13$ (there are 13 nearest-neighbors on one side of any voxel), this corresponds to storing at least 26 image-volumes (≈ 26 GB of memory) that might set a practical limitation on these methods from an implementation perspective.

¹⁴For ADMM, we have to additionally store two $M \times 1$ vectors, \mathbf{u} and the associated multiplier $\boldsymbol{\eta}_u$. This additional memory requirement is moderate for 2D CT and can be high for 3D CT depending on the size of the data.

A quick remedy is to consider the TV regularizer with finite-differences only along the three orthogonal directions ($P=3$ corresponds to 6 image-volumes) which considerably reduces the memory load. Alternatively, one could also consider using an orthonormal transform (such as orthonormal wavelets¹⁵) for \mathbf{R} , so $P=1$ and $L=N$. The SB and ADMM¹⁴ schemes would then require storing only 2 image-volumes (corresponding to \mathbf{v} and η_N). Moreover, an orthonormal \mathbf{R} satisfies $\mathbf{R}^T \mathbf{R} = \mathbf{I}_N$ that facilitates ADMM: \mathbf{G}_v in (28) becomes $\mathbf{G}_v = (\mathbf{A}^T \mathbf{A} + \nu \mathbf{I}_N)$ that is still “nearly” shift-invariant and can be effectively preconditioned using circulant preconditioners. With orthonormal wavelets, one also has the option of excluding the approximation coefficients from the regularization (as they are not sparse) by using scale dependent regularization parameters [50] and setting those parameters corresponding to the approximation level to zero.

B. Inclusion of Nonnegativity Constraint

In CT, a nonnegativity constraint is often imposed [1, Eq. 18], [11, Sec. 2.2] to model the positivity of the attenuation coefficient that is being reconstructed. Although we have not considered such a constraint in $\mathbf{P0}$, it can be easily accommodated [38] as follows. We start with

$$\arg \min_{\mathbf{x}} \{J(\mathbf{x}) \triangleq J_{\text{data}}(\mathbf{y}, \mathbf{A}\mathbf{x}) + \Psi(\mathbf{R}\mathbf{x}) + g(\mathbf{x})\}, \quad (39)$$

where g is an indicator function

$$g(\mathbf{x}) \triangleq \begin{cases} 0, & \text{if } \mathbf{x} \geq 0, \\ +\infty, & \text{otherwise,} \end{cases} \quad (40)$$

that imposes the nonnegativity constraint, $\mathbf{x} \geq 0$, taken component-wise in (40). We then consider the following equivalent constrained version [38] that has an additional constraint compared to (5):

$$\begin{aligned} \arg \min_{\mathbf{x}, \mathbf{u}, \mathbf{v}, \mathbf{w}} \{f(\mathbf{u}, \mathbf{v}, \mathbf{w}) \triangleq J_{\text{data}}(\mathbf{y}, \mathbf{u}) + \Psi(\mathbf{v}) + g(\mathbf{w})\} \\ \text{s.t. } \mathbf{u} = \mathbf{A}\mathbf{x}, \quad \mathbf{v} = \mathbf{R}\mathbf{x}, \quad \mathbf{w} = \mathbf{x}. \end{aligned} \quad (41)$$

Writing $\mathbf{z} \triangleq [\mathbf{u}^T \quad \mathbf{v}^T \quad \mathbf{w}^T]^T$, $\mathbf{C} \triangleq [\mathbf{A}^T \quad \mathbf{R}^T \quad \mathbf{I}_N]^T$ and using $\Lambda = \text{diag}\{\mathbf{I}_M, \nu_1 \mathbf{I}_R, \nu_2 \mathbf{I}_N\}$ in (9), we can design an ADMM-type algorithm similar to (15)–(17) for solving (41). It can be shown that the updates for \mathbf{u} and \mathbf{v} in this algorithm will be similar to (19) and (20), respectively, while the update of \mathbf{x} would involve the “inversion” of $\mathbf{G}_{\nu_1 \nu_2} = (\mathbf{A}^T \mathbf{A} + \nu_1 \mathbf{R}^T \mathbf{R} + \nu_2 \mathbf{I}_N)$ and that of \mathbf{w} would require a simple projection onto the positive orthant [38, Eq. 32]. Since $\mathbf{G}_{\nu_1 \nu_2}$ is also “nearly” shift-invariant, a cone-filter-type preconditioner similar to $\tilde{\mathbf{G}}_v^{-1}$ [see (30)] can be used for effective preconditioning of $\mathbf{G}_{\nu_1 \nu_2}^{-1}$. Moreover, the above \mathbf{C} has full column-rank, so this algorithm also satisfies Theorem 1 and is guaranteed to converge to a solution of (41) and (39).

C. Poisson-Likelihood Model for X-ray CT Reconstruction

The proposed strategy of splitting the data-term [i.e., the use of \mathbf{u} in (5) and (41)] is also applicable for X-ray CT reconstruction using the Poisson-likelihood (PL) statistical model

¹⁵Quality-wise, shift-invariant wavelets are preferable to orthonormal ones [54], but due to their over-complete nature, they require significantly more memory (similar to finite differences) than orthonormal wavelets.

[5, Eq. 1] that may be more suitable for low-dose acquisitions. It can be shown that splitting the PL data-term yields separable 1D problems in $\{u_i\}_{i=1}^M$ that can be solved simultaneously similar to [38, Eq. 30]. However, the PL model for X-ray CT may preclude exact updates like (21) for $\{u_i\}$. Moreover, the general PL model [5, Eq. 1] includes background events and can be (“mildly”) nonconvex, so Theorem 1 cannot be directly applied to an ADMM-type algorithm developed for this problem. We plan to explore cost functions involving the PL model [5, Eq. 1] for transmission tomography reconstruction as part of future extensions to this work.

VII. Summary and Conclusions

Statistical X-ray CT reconstruction using penalized weighted least-squares (PWLS) criteria involve a diagonal weighting matrix \mathbf{W} that poses a hindrance to several optimization methods due to its huge dynamic range and highly shift-variant nature. In this work, we employed a variable-splitting technique that, in addition to separating the regularization term like [17], also dissociates the statistical (\mathbf{W}) and the system (\mathbf{A}) components in the data term to decouple and mitigate the effect of \mathbf{W} . We applied the method of multipliers [19] with alternating minimization [21]–[23] for the resulting equivalent constrained problem and developed an alternating direction method of multipliers (ADMM) algorithm that chiefly involves three simple operations at each iteration: (i) inverting a diagonal matrix that depends on \mathbf{W} , (ii) minimizing a set of 1D auxiliary denoising-cost-functions that can be performed efficiently and/or exactly for a variety of regularizers, and (iii) solving a “nearly” shift-invariant linear system (involving $\mathbf{A}^T\mathbf{A}$) using FFT-based preconditioning with cone-type filters [7].

The proposed ADMM algorithm is guaranteed to converge to a solution of the original PWLS problem under a mild condition on the accuracy of operation (iii) above. We demonstrated using simulations and experiments with real *in vivo* human data that cone-filter-type preconditioners are very effective for solving the linear system in (iii) and that the preconditioned version of the proposed ADMM converges faster than conventional (NCG and ordered subsets) and state-of-the-art (MFISTA and split-Bregman) algorithms for CT. The proposed ADMM can handle a variety of regularization criteria for 2D CT reconstruction and is also applicable to 3D CT reconstruction, perhaps by using certain memory-conserving regularizers.

Acknowledgments

This work was supported by the National Institutes of Health under Grant R01-HL-098686.

The authors would like to thank GE Healthcare, particularly J. Seamans, J. B. Thibault, and B. DeMan for providing CT sinograms and code for the DD-projector, and M. McGaffin, University of Michigan, for proofreading the manuscript.

References

1. Thibault JB, 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(11):4526–44. [PubMed: 18072519]
2. Sauer K, Bouman C. A local update strategy for iterative reconstruction from projections. *IEEE Trans Sig Proc.* Feb; 1993 41(2):534–48.
3. Benson, TM.; Man, BKBD.; Fu, L.; Thibault, J-B. Block-based iterative coordinate descent; *Proc IEEE Nuc Sci Symp Med Im Conf*; 2010. p. 2856-9.
4. Kamphuis C, Beekman FJ. Accelerated iterative transmission CT reconstruction using an ordered subsets convex algorithm. *IEEE Trans Med Imag.* Dec; 1998 17(6):1001–5.

5. Erdogan H, Fessler JA. Ordered subsets algorithms for transmission tomography. *Phys Med Biol.* Nov; 1999 44(11):2835–51. [PubMed: 10588288]
6. Fessler JA, Booth SD. Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction. *IEEE Trans Im Proc.* May; 1999 8(5):688–99.
7. Clinthorne NH, Pan TS, Chiao PC, Rogers WL, Stamos JA. Preconditioning methods for improved convergence rates in iterative reconstructions. *IEEE Trans Med Imag.* Mar; 1993 12(1):78–83.
8. Delaney AH, Bresler Y. A fast and accurate Fourier algorithm for iterative parallel-beam tomography. *IEEE Trans Im Proc.* May; 1996 5(5):740–53.
9. Wallis JW, Miller TR. Rapidly converging iterative reconstruction algorithms in single-photon emission computed tomography. *J Nuc Med.* Oct; 1993 34(10):1793–800.
10. Nuyts J, De Man B, Dupont P, Defrise M, Suetens P, Mortel-mans L. Iterative reconstruction for helical CT: A simulation study. *Phys Med Biol.* Apr; 1998 43(4):729–37. [PubMed: 9572499]
11. Sidky EY, Pan X. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Phys Med Biol.* Sep; 2008 53(17):4777–808. [PubMed: 18701771]
12. Ramani S, Fessler JA. Parallel MR image reconstruction using augmented Lagrangian methods. *IEEE Trans Med Imag.* Mar; 2011 30(3):694–706.
13. Lustig M, Donoho D, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Mag Res Med.* Dec; 2007 58(6):1182–95.
14. Ramani, S.; Fessler, JA. An accelerated iterative reweighted least squares algorithm for compressed sensing MRI. *Proc. IEEE Intl. Symp. Biomed. Imag.*; 2010. p. 257-60.
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. Beck A, Teboulle M. Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *IEEE Trans Im Proc.* Nov; 2009 18(11):2419–34.
17. Goldstein T, Osher S. The split Bregman method for L1-regularized problems. *SIAM J Imaging Sci.* 2009; 2(2):323–43.
18. Vandeghinste, B.; Goossens, B.; Beenhouwer, JD.; Pizurica, A.; Philips, W.; Vandenberghe, S.; Staelens, S. Split-Bregman-based sparse-view CT reconstruction. *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med*; 2011. p. 431-4.
19. Bertsekas DP. Multiplier methods: A survey. *Automatica.* Mar; 1976 12(2):133–45.
20. Nocedal, J.; Wright, SJ. *Numerical optimization.* Springer; New York: 1999.
21. Gabay D, Mercier B. A dual algorithm for the solution of nonlinear variational problems via finite element approximations. *Computers and Mathematics with Applications.* 1976; 2:17–40.
22. Glowinski R, Marroco A. Sur l'approximation par element finis d'ordre un, et la resolution, par penalisation-dualité, d'une classe de problemes de Dirichlet nonlineares. *Revue Française d'Automatique, Informatique et Recherche Opérationelle* 9. 1975; R-2:41–76.
23. Fortin, M.; Glowinski, R. On decomposition-coordination methods using an augmented Lagrangian. In: Fortin, M.; Glowinski, R., editors. *Augmented Lagrangian Methods: Applications to the Solution of Boundary-Value Problems.* Studies in Mathematics and Its Applications. Vol. 15. Elsevier; North-Holland, Amsterdam: 1983. p. 97-146.
24. Elbakri, IA.; Fessler, JA. Efficient and accurate likelihood for iterative image reconstruction in X-ray computed tomography. *Proc. SPIE 5032, Medical Imaging 2003: Image Proc.*; 2003. p. 1839-50.
25. Whiting, BR. Signal statistics in x-ray computed tomography. *Proc. SPIE 4682, Medical Imaging 2002: Med. Phys.*; 2002. p. 53-60.
26. Fessler JA, Rogers WL. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Trans Im Proc.* Sep; 1996 5(9):1346–58.
27. Elad, Michael; Milanfar, Peyman; Rubinstein, Ron. Analysis versus synthesis in signal priors. *Inverse Prob.* 2007; 23:947–68.
28. Huber, PJ. *Robust statistics.* Wiley; New York: 1981.
29. Nikolova M, Ng MK. Analysis of half-quadratic minimization methods for signal and image recovery. *SIAM J Sci Comp.* 2005; 27(3):937–66.

30. Fair RC. On the robust estimation of econometric models. *Ann Econ Social Measurement*. Oct. 1974 2:667–77.
31. 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(11):1413–57.
32. Bioucas-Dias JM, Figueiredo MAT. A new twIST: two-step iterative shrinkage/Thresholding algorithms for image restoration. *IEEE Trans Im Proc*. Dec; 2007 16(12):2992–3004.
33. Ramani S, Thévenaz P, Unser M. Regularized interpolation for noisy images. *IEEE Trans Im Proc*. Feb; 2010 29(2):543–58.
34. Wang Y, Yang J, Yin W, Zhang Y. A new alternating minimization algorithm for total variation image reconstruction. *SIAM J Imaging Sci*. 2008; 1(3):248–72.
35. Ng MK, Weiss P, Yuan X. Solving constrained total-variation image restoration and reconstruction problems via alternating direction methods. *SIAM J Sci Comp*. 2010; 32(5):2710–36.
36. Afonso MV, Bioucas-Dias JM, Figueiredo MAT. An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems. *IEEE Trans Im Proc*. 2011; 20(3):681–95.
37. Afonso MV, Bioucas-Dias José M, Figueiredo Mário AT. Fast image recovery using variable splitting and constrained optimization. *IEEE Trans Im Proc*. Sep; 2010 19(9):2345–56.
38. Figueiredo MAT, Bioucas-Dias JM. Restoration of Poissonian images using alternating direction optimization. *IEEE Trans Im Proc*. 2010; 19(12):3133–45.
39. Ye X, Chen Y, Huang F. Computational acceleration for MR image reconstruction in partially parallel imaging. *IEEE Trans Med Imag*. 2011; 30(5):1055– 63.
40. Ye X, Chen Y, Lin W, Huang F. Fast MR image reconstruction for partially parallel imaging with arbitrary k -space trajectories. *IEEE Trans Med Imag*. 2011; 30(3):575– 85.
41. Xu Q, Mou X, Wang G, Sieren J, Hoffman EA, Yu H. Statistical interior tomography. *IEEE Trans Med Imag*. May; 2011 30(5):1116–28.
42. Ramani, S.; Fessler, JA. Convergent iterative CT reconstruction with sparsity-based regularization. *Image Recon in Rad and Nuc Med; Proc. Intl. Mtg. on Fully 3D*; 2011. p. 302-5.
43. Eckstein J, Bertsekas DP. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Math Prog. Apr*; 1992 55(1–3):293–318.
44. Rockafellar, RT. *Convex analysis*. Princeton University Press; Princeton: 1970.
45. Chaux C, Combettes PL, Pesquet J-C, Wajs Valérie R. A variational formulation for frame-based inverse problems. *Inverse Prob*. Aug; 2007 23(4):1495–518.
46. Chambolle A, DeVore RA, Lee NY, Lucier BJ. Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage. *IEEE Trans Im Proc*. Mar; 1998 7(3):319–35.
47. Dixon JD. Estimating extremal eigenvalues and condition numbers of matrices. *SIAM J Num Anal*. 1983; 20(4):812–814.
48. Anton, H.; Rorres, C. *Elementary linear algebra - applications version*. John Wiley & Sons; 2010.
49. De Man, B.; Basu, S. Distance-driven projection and backprojection. *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*; 2002. p. 1477-80.
50. Selesnick, IW.; Mário, A.; Figueiredo, T. spie-7446. *Wavelets XIII: 2009. Signal restoration with overcomplete wavelet transforms: comparison of analysis and synthesis priors*; p. 74460D
51. Segars WP, Tsui BMW. Study of the efficacy of respiratory gating in myocardial SPECT using the new 4-D NCAT phantom. *IEEE Trans Nuc Sci*. Jun; 2002 49(3):675–9.
52. Wang J, Li T, Lu H, Liang Z. Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose X-ray computed tomography. *IEEE Trans Med Imag*. Oct; 2006 25(10):1272–83.
53. Fessler, JA. *Image Reconstruction Toolbox*. University of Michigan; 2011. <http://www.eecs.umich.edu/~fessler/code/index.html>
54. Figueiredo MAT, Nowak RD. Wavelet-based image estimation: an empirical Bayes approach using Jeffrey’s noninformative prior. *IEEE Trans Im Proc*. Sep; 2001 10(9):1322–31.

1. Select $\mathbf{x}^{(0)}$, $\mu, \nu > 0$ and set $j = 0$
 2. Set $\mathbf{u}^{(0)} = \mathbf{Ax}^{(0)}$, $\mathbf{v}^{(0)} = \mathbf{Rx}^{(0)}$, and $\eta_{\mathbf{u}}^{(0)} = \eta_{\mathbf{v}}^{(0)} = \mathbf{0}$
- Repeat:**
3. Obtain $\mathbf{x}^{(j+1)}$ by applying (P)CG iterations to (27)
 4. Compute $\mathbf{u}^{(j+1)}$ using (21)
 5. Compute $\mathbf{v}^{(j+1)}$ using (24)-(26)
 6. $\eta_{\mathbf{u}}^{(j+1)} = \eta_{\mathbf{u}}^{(j)} - (\mathbf{u}^{(j+1)} - \mathbf{Ax}^{(j+1)})$
 7. $\eta_{\mathbf{v}}^{(j+1)} = \eta_{\mathbf{v}}^{(j)} - (\mathbf{v}^{(j+1)} - \mathbf{Rx}^{(j+1)})$
 8. Set $j = j + 1$
- Until stop criterion is met**

Fig. 1.
ADMM for statistical X-ray CT reconstruction.




Fig. 2. Simulation with the NCAT phantom: (a) Noisefree NCAT phantom (in cm^{-1}), (b) FBP reconstruction with ramp filter, also the initial guess $\mathbf{x}^{(0)}$ for all iterative algorithms, (c) FBP reconstruction with Hanning filter, and (d) ℓ_1 -regularized reconstruction, also the solution \mathbf{x}^\star to $\mathbf{P0}$. Images in (a)–(d) have been normalized to the same color scale [as that of (a)] indicated beside (d). The ℓ_1 -regularized reconstruction (d) is less noisy and has almost no streaky-artifacts compared to both FBP results.

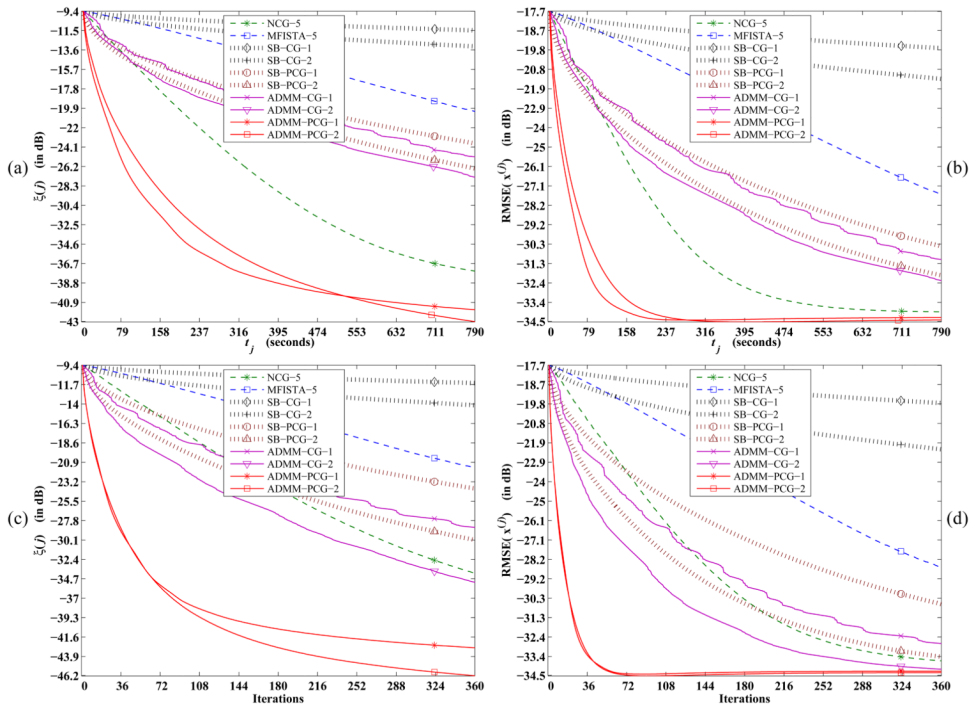


Fig. 3. Simulation with the NCAT phantom: (a), (b) Plot of $\xi(j)$ and $\text{RMSE}(\mathbf{x}^j)$, respectively, as a function of time t_j and (c), (d) Plot of $\xi(j)$ and $\text{RMSE}(\mathbf{x}^j)$, respectively, with respect to iterations, for various algorithms considered in this work. The unpreconditioned version of the proposed method, ADMM-CG, converges slightly faster than MFISTA and the split-Bregman scheme SB-(P)CG but is slower than NCG as seen in (a) and (b). But the preconditioned version, ADMM-PCG, converges rapidly both in terms of $\xi(j)$ and RMSE indicating that the cone-filter-preconditioner ($\tilde{\mathbf{G}}_v^{-1}$ in Section III-E) greatly accelerates convergence of the proposed ADMM.




Fig. 4. Experiment with the *in vivo* human head data-set: (a) FBP reconstruction with the ramp filter, also the initial guess $\mathbf{x}^{(0)}$ for all iterative algorithms, (b) RBP reconstruction with Hanning filter, and (c) PWLS reconstruction with the strictly convex regularizer (23), also the unique solution \mathbf{x}^{\star} to $\mathbf{P}\mathbf{0}$. Images in (a)–(c) are displayed in Hounsfield units indicated beside (c). The regularized reconstruction (c) is less noisy and preserves anatomical features compared to both FBP results.




Fig. 5. Experiment with the *in vivo* human head data-set: (a), (b) Plot of $\xi(j)$ as a function of time t_j and (c), (d) Plot of $\xi(j)$ with respect to iterations, for various algorithms considered in this work. MFISTA and SB-CG appear to be the slowest. The proposed ADMM-(P)CG is generally faster than the split-Bregman scheme SB-(P)CG as seen in (a). Although ADMM-CG converges slower than NCG as seen in (b), the preconditioned version ADMM-PCG is the fastest among the considered algorithms, illustrating the benefit of the cone-filter-based preconditioner ($\tilde{\mathbf{G}}_\nu$ in Section III-E) for the proposed ADMM.

TABLE I

Computation Time and Number of Projections required per Iteration of Algorithms Compared in Section V

Algorithm	Time/Iteration (in seconds)		Projection operations/Iteration	
	Section V-A	Section V-B	Forward (A)	Backward (A^T)
NCG-5	1.56	4.85	1	1
NCG-10	-	8.83	1	1
MFISTA-5	2.49	8.87	1	1
MFISTA-25	5.23	-	1	1
OS-4	-	10.19	1 (effective)	1 (effective)
OS-41	-	61.84	1 (effective)	1 (effective)
SB-CG-1	2.29	6.22	1	1
SB-CG-2	3.29	8.93	2	2
SB-PCG-1	2.29	6.25	1	1
SB-PCG-2	3.30	9.07	2	2
ADMM-CG-1	3.29	8.91	1	2
ADMM-CG-2	4.31	11.61	2	3
ADMM-PCG-1	3.32	8.94	1	2
ADMM-PCG-2	4.34	11.70	2	3