

Accelerating ordered-subsets image reconstruction for X-ray CT using double surrogates

Jang Hwan Cho and Jeffrey A. Fessler

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

ABSTRACT

Conventional ordered-subsets (*OS*) methods for regularized image reconstruction involve computing the gradient of the regularizer for every subset update. When dealing with large problems with many subsets, such as in 3D X-ray CT, computing the gradient for each subset update can be very computationally expensive. To mitigate this issue, some investigators use unregularized iterations followed by a denoising operation after updating all subsets.¹ Although such methods save computation, their convergence properties are uncertain, and since they may not be minimizing any particular cost function it becomes more difficult to design regularization parameters. Furthermore, it is known that inserting filtering steps into unregularized algorithms can lead to undesirable spatial resolution properties.² Our goal here is to reduce the computational cost without inducing such problems. We propose a new OS-type algorithm that is derived using optimization transfer principles. The proposed method allows the gradient of the regularizer to be updated less frequently, and thus reduces the computational expense when many subsets are used. Our derivation leads to a correction term that accounts for the fact that the regularizer gradient is updated less frequent than every sub-iteration. Simulations and a phantom experiment show that the proposed method reconstructed images with compatible image quality within reduced computation time.

Keywords: ordered-subsets, double surrogates, cone-beam tomography, iterative reconstruction, regularization

1. INTRODUCTION

Model-based image reconstruction (*MBIR*) methods, also known as statistical image reconstruction methods, can incorporate accurate system models and take the stochastic characteristics of noise into account, and thus have the potential to improve CT image quality and reduce patient dose compared to conventional methods such as filtered back-projection (*FBP*). However, iterative algorithms for penalized-likelihood (*PL*) image reconstruction require considerable computations, and the computational expense is one of the great challenges for practical utilization of MBIR methods.

Ordered-subsets (*OS*) algorithms, also known as block-iterative or incremental gradient methods, are popular in the field of statistical image reconstruction due to their significant acceleration in initial iterations. The basic idea of OS methods is to group the measurement data into an ordered sequence of subsets or blocks and utilize only one subset of the data for each update instead of using the entire measurements. Unregularized OS reconstruction methods are used routinely for PET and SPECT scans.³ For regularized image reconstruction problems, conventional OS algorithms calculate the gradient of the regularizer for every subset update.⁴ For large problems with large number of subsets, such as cone-beam or helical CT image reconstruction, calculating the gradient for each subset update can be very computationally expensive.

In this paper, we propose a new OS algorithm that is derived from optimization transfer principles and that allows us to compute the gradient of the regularizer less frequently. The proposed method reduces computational cost with little impact on the convergence rate leading to an overall acceleration. We apply the proposed algorithm to penalized weighted least squares (*PWLS*) image reconstruction for cone-beam X-ray computed tomography (*CT*). The method easily could be adapted for other statistical models.

Supported in part by NIH grant R01-HL-098686 and by GE Healthcare.

Jang Hwan Cho: janghcho@umich.edu, 734-615-5735

Jeffrey A. Fessler: fessler@umich.edu, 734-763-1434

2. METHODS

2.1 Ordered-Subsets (OS) algorithms

For most penalized-likelihood image reconstruction problems, the objective function and its gradient can be rewritten as the following partially separable forms:⁵

$$\Psi(\mathbf{x}) = \sum_{m=1}^M \Psi_m(\mathbf{x}), \quad \nabla_{\mathbf{x}} \Psi(\mathbf{x}) = \sum_{m=1}^M \nabla_{\mathbf{x}} \Psi_m(\mathbf{x}),$$

where typically $\Psi_m(\mathbf{x})$ corresponds to a subset of the projection data. Most iterative algorithms use the gradient of the objective function to obtain a minimizer, and many can be written as the form

$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - \alpha_n \mathbf{D}_0(\mathbf{x}^{(n)}) \nabla_{\mathbf{x}} \Psi(\mathbf{x}^{(n)}), \quad (1)$$

where $\alpha_n > 0$ is a relaxation parameter and \mathbf{D}_0 is a scaling matrix, which is typically diagonal. For many algorithms, the empirical findings have suggested that in early iterations, one can replace the gradient of the entire objective function with that of only a part of the objective function and still have $\mathbf{x}^{(n+1)}$ be better than $\mathbf{x}^{(n)}$ but with less computational cost. Such methods are called incremental gradient methods in the optimization literature.^{6,7} The term ordered-subsets has been used in tomography fields, because only a subset of projection views are used in the update steps and the ordering of the views is important.^{3,4,8} For $\mathbf{x}^{(n)}$ far from the minimizer, $\hat{\mathbf{x}}$, if we select the subsets to be “balanced” in some appropriate sense then the following conditions hold:

$$\nabla_{\mathbf{x}} \Psi_1(\mathbf{x}) \cong \nabla_{\mathbf{x}} \Psi_2(\mathbf{x}) \cong \dots \cong \nabla_{\mathbf{x}} \Psi_M(\mathbf{x}),$$

or equivalently,

$$\nabla_{\mathbf{x}} \Psi(\mathbf{x}) \approx M \nabla_{\mathbf{x}} \Psi_m(\mathbf{x}), \quad \forall m. \quad (2)$$

Thus, instead of (1), a typical OS algorithm has the form shown in Table 1.

Table 1. Ordered-subsets Algorithm

Initialize $\mathbf{x}^{(0)}$	
for $n = 0, 1, \dots$,	$\mathbf{x}^{(n)} := \mathbf{x}^{(n,0)}$
for $m = 1, \dots, M$	$\mathbf{x}^{(n,m)} = \mathbf{x}^{(n,m-1)} - \alpha_n M \mathbf{D}(\mathbf{x}^{(n,m-1)}) \nabla_{\mathbf{x}} \Psi_m(\mathbf{x}^{(n,m-1)}),$
end	$\mathbf{x}^{(n+1)} := \mathbf{x}^{(n,M)}$
end	

We refer to each update in Table 1 as the m th subset update, or m th sub-iteration of the n th iteration. One complete iteration is performed by cycling through all the subsets indexed by m so that all data is utilized. In tomography problems, the subsets are selected so that projections within each subset correspond to angularly downsampled projections. It was suggested that the ordering of the subsets that makes projections corresponding to one subset as “orthogonal” as possible to previously used projections is preferable.^{9,10} Despite their success in speeding up the initial convergence, ordinary OS algorithms are not convergent in general but rather approach a suboptimal limit-cycle without relaxation, i.e., when $\alpha_n = \alpha$. To address this issue, several families of convergent OS type algorithms have been proposed,¹¹⁻¹⁴ although those modifications tend to slow down convergence. In this paper, we focus on the initial convergence characteristics of OS algorithms rather than their final convergence properties.

Consider a PWLS objective function of the form

$$\Psi(\mathbf{x}) = \mathbf{L}(\mathbf{x}) + \mathbf{R}(\mathbf{x}), \quad \mathbf{L}(\mathbf{x}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_{\mathbf{W}}^2, \quad \mathbf{R}(\mathbf{x}) = \beta \sum_{k=1}^K \psi_k([\mathbf{C}\mathbf{x}]_k), \quad (3)$$

where \mathbf{A} is the system matrix, \mathbf{x} is the discretized version of the object being imaged, $\mathbf{W} = \text{diag}\{w_i\}$ is a statistical weighting matrix, β is the regularization coefficient that controls the resolution-noise tradeoff, ψ_k is the potential function, and \mathbf{C} is a matrix performs finite differences between neighboring voxels. The data-fit term can be rewritten as the following separable form:

$$\mathbf{t}(\mathbf{x}) = \sum_{m=1}^M \mathbf{t}_m(\mathbf{x}),$$

where \mathbf{t}_m is the data-fit term associated with the m th subset. OS algorithm for PWLS objective function is shown in Table 2. For *separable quadratic surrogates* method, we can define the scaling matrix \mathbf{D} as follows:

$$\mathbf{D}(\mathbf{x}) = [\mathbf{D}_L(\mathbf{x}) + \mathbf{D}_R(\mathbf{x})]^{-1}.$$

Typical choices⁴ for \mathbf{D}_L and \mathbf{D}_R are $\mathbf{D}_L = \text{diag}(\mathbf{A}'\mathbf{W}\mathbf{A}\mathbf{1})$ and $\mathbf{D}_R(\mathbf{x}) = \text{diag}(|\mathbf{C}'|' \text{diag}(|\omega(\mathbf{C}\mathbf{x})|) |\mathbf{C}'| \mathbf{1})$ where $\omega(t) = \dot{\psi}(t)/t$ is Huber's curvature.¹⁵ Notice that we need to calculate the regularization gradient, $\nabla_{\mathbf{x}} \mathbf{R}(\mathbf{x}^{(n,m-1)})$, for every sub-iteration, which may cause each update to be very computationally expensive when dealing with large problems with many subsets.

Table 2. Ordered-subsets Algorithm for PWLS problem

Initialize $\mathbf{x}^{(0)}$ for $n = 0, 1, \dots$, <div style="text-align: center; margin: 10px 0;"> $\mathbf{x}^{(n)} := \mathbf{x}^{(n,0)}$ </div> for $m = 1, \dots, M$ <div style="text-align: center; margin: 10px 0;"> $\mathbf{x}^{(n,m)} = \mathbf{x}^{(n,m-1)} - \alpha_n \mathbf{D}(\mathbf{x}^{(n,m-1)}) \left(M \nabla_{\mathbf{x}} \mathbf{t}_m(\mathbf{x}^{(n,m-1)}) + \nabla_{\mathbf{x}} \mathbf{R}(\mathbf{x}^{(n,m-1)}) \right),$ </div> end <div style="text-align: center; margin: 10px 0;"> $\mathbf{x}^{(n+1)} := \mathbf{x}^{(n,M)}$ </div> end
--

2.2 Ordered-subsets with Double Surrogates

Consider a general PL objective function of the form

$$\Psi(\mathbf{x}) = \mathbf{t}(\mathbf{x}) + \mathbf{R}(\mathbf{x}), \quad (4)$$

where $\mathbf{t}(\mathbf{x})$ is the data-fit term and $\mathbf{R}(\mathbf{x})$ is the regularizer. We assume that the data-fit term $\mathbf{t}(\mathbf{x})$ has a quadratic surrogate of the form:

$$\mathbf{t}(\mathbf{x}) \leq \phi_L(\mathbf{x}; \tilde{\mathbf{x}}) = \mathbf{t}(\tilde{\mathbf{x}}) + \nabla \mathbf{t}(\tilde{\mathbf{x}})(\mathbf{x} - \tilde{\mathbf{x}}) + \frac{1}{2}(\mathbf{x} - \tilde{\mathbf{x}})' \mathbf{D}_L(\tilde{\mathbf{x}})(\mathbf{x} - \tilde{\mathbf{x}}), \quad \forall \tilde{\mathbf{x}},$$

with an appropriate diagonal matrix \mathbf{D}_L . We also assume the regularizer $\mathbf{R}(\mathbf{x})$ has a quadratic surrogate of the form:

$$\mathbf{R}(\mathbf{x}) \leq \phi_R(\mathbf{x}; \bar{\mathbf{x}}) = \mathbf{R}(\bar{\mathbf{x}}) + \nabla \mathbf{R}(\bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}}) + \frac{1}{2}(\mathbf{x} - \bar{\mathbf{x}})' \mathbf{D}_R(\bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}}), \quad \forall \bar{\mathbf{x}},$$

with an appropriate diagonal matrix \mathbf{D}_R . Then we define the following double-surrogate function:

$$\phi(\mathbf{x}; \tilde{\mathbf{x}}, \bar{\mathbf{x}}) \triangleq \phi_L(\mathbf{x}; \tilde{\mathbf{x}}) + \phi_R(\mathbf{x}; \bar{\mathbf{x}}). \quad (5)$$

By construction, this quadratic surrogate has the following properties:

$$\begin{aligned} \Psi(\mathbf{x}) &= \phi(\mathbf{x}; \mathbf{x}, \mathbf{x}) \\ \Psi(\mathbf{x}) &\leq \phi(\mathbf{x}; \tilde{\mathbf{x}}, \bar{\mathbf{x}}), \quad \forall \tilde{\mathbf{x}}, \bar{\mathbf{x}}. \end{aligned}$$

These properties generalize those of usual optimization transfer methods.^{16,17} For subsequent use, note that from (5):

$$\nabla_{\mathbf{x}} \phi(\mathbf{x}; \tilde{\mathbf{x}}, \bar{\mathbf{x}}) = \nabla \mathbf{t}(\tilde{\mathbf{x}}) + \mathbf{D}_L(\tilde{\mathbf{x}})(\mathbf{x} - \tilde{\mathbf{x}}) + \nabla \mathbf{R}(\bar{\mathbf{x}}) + \mathbf{D}_R(\bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}}),$$

so the minimizer of the double surrogate function is given as follows:

$$\arg \min_{\mathbf{x}} \phi(\mathbf{x}; \tilde{\mathbf{x}}, \bar{\mathbf{x}}) = [\mathbf{D}_L(\tilde{\mathbf{x}}) + \mathbf{D}_R(\bar{\mathbf{x}})]^{-1} (\mathbf{D}_L(\tilde{\mathbf{x}})\tilde{\mathbf{x}} + \mathbf{D}_R(\bar{\mathbf{x}})\bar{\mathbf{x}} - \nabla \mathbf{L}(\tilde{\mathbf{x}}) - \nabla \mathbf{R}(\bar{\mathbf{x}})).$$

A standard optimization transfer algorithm works as follows:

$$\begin{aligned} \mathbf{x}^{(n+1)} &= \arg \min_{\mathbf{x}} \phi(\mathbf{x}; \mathbf{x}^{(n)}, \mathbf{x}^{(n)}) \\ &= \mathbf{x}^{(n)} - [\mathbf{D}_L(\mathbf{x}^{(n)}) + \mathbf{D}_R(\mathbf{x}^{(n)})]^{-1} (\nabla \mathbf{L}(\mathbf{x}^{(n)}) + \nabla \mathbf{R}(\mathbf{x}^{(n)})). \end{aligned}$$

One can show this also decreases Ψ monotonically, i.e., $\Psi(\mathbf{x}^{(n+1)}) \leq \Psi(\mathbf{x}^{(n)})$. Furthermore, this algorithm converges under suitable conditions.¹⁶ For the case of *separable quadratic surrogates (SQS)*, usually this type of algorithm converges undesirably slowly.⁴ The conventional *ordered-subsets (OS)* approach to accelerate convergence is to make the following approximation:

$$\nabla \mathbf{L}(\tilde{\mathbf{x}}) \approx \mathbf{D}_m \nabla \mathbf{L}_m(\tilde{\mathbf{x}}),$$

where \mathbf{L}_m is the data-fit term corresponding to the m th subset of the projection views, and \mathbf{D}_m is a suitable diagonal matrix, which often simply is $M\mathbf{I}$ for M subsets as suggested in (2). Using this approximation, we define the following *approximate* surrogate function:

$$\phi_m(\mathbf{x}; \tilde{\mathbf{x}}, \bar{\mathbf{x}}) \triangleq \mathbf{L}(\tilde{\mathbf{x}}) + \mathbf{D}_m \nabla \mathbf{L}_m(\tilde{\mathbf{x}})(\mathbf{x} - \tilde{\mathbf{x}}) + \frac{1}{2}(\mathbf{x} - \tilde{\mathbf{x}})' \mathbf{D}_L(\tilde{\mathbf{x}})(\mathbf{x} - \tilde{\mathbf{x}}) + \phi_R(\mathbf{x}; \bar{\mathbf{x}}).$$

For a conventional regularized ordered-subsets method, the minimization step for each subset is given as follows:

$$\begin{aligned} \mathbf{x}^{(n,m)} &= \arg \min_{\mathbf{x}} \phi_m(\mathbf{x}; \mathbf{x}^{(n,m-1)}, \mathbf{x}^{(n,m-1)}) \\ &= \mathbf{x}^{(n,m-1)} - [\mathbf{D}_L(\mathbf{x}^{(n,m-1)}) + \mathbf{D}_R(\mathbf{x}^{(n,m-1)})]^{-1} (\mathbf{D}_m \nabla \mathbf{L}(\mathbf{x}^{(n,m-1)}) + \nabla \mathbf{R}(\mathbf{x}^{(n,m-1)})), \\ \mathbf{x}^{(n,0)} &:= \mathbf{x}^{(n)}, \quad \mathbf{x}^{(n+1)} := \mathbf{x}^{(n,M)}, \end{aligned} \tag{6}$$

for $m = 1, \dots, M$.

This iteration is undesirably slow because it computes the regularization gradient $\nabla \mathbf{R}$ for every subset. To reduce this expense, we propose to exploit the generality of the double surrogate (5) by using the following novel update:

$$\begin{aligned} \mathbf{x}^{(n,m)} &= \arg \min_{\mathbf{x}} \phi_m(\mathbf{x}; \mathbf{x}^{(n,m-1)}, \mathbf{x}^{(n)}) \\ &= \mathbf{D}_\phi(\mathbf{x}^{(n,m-1)}, \mathbf{x}^{(n)}) (\mathbf{D}_L(\mathbf{x}^{(n,m-1)})\mathbf{x}^{(n,m-1)} + \mathbf{D}_R(\mathbf{x}^{(n)})\mathbf{x}^{(n)} - \mathbf{D}_m \nabla \mathbf{L}(\mathbf{x}^{(n,m-1)}) - \nabla \mathbf{R}(\mathbf{x}^{(n)})) \\ &= \mathbf{x}^{(n,m-1)} - \mathbf{D}_\phi(\mathbf{x}^{(n,m-1)}, \mathbf{x}^{(n)}) \left(\mathbf{D}_m \nabla \mathbf{L}(\mathbf{x}^{(n,m-1)}) + \nabla \mathbf{R}(\mathbf{x}^{(n)}) + \underbrace{\mathbf{D}_R(\mathbf{x}^{(n)})(\mathbf{x}^{(n,m-1)} - \mathbf{x}^{(n)})}_{\text{new term}} \right), \end{aligned} \tag{7}$$

where $\mathbf{D}_\phi(\mathbf{x}^{(n,m-1)}, \mathbf{x}^{(n)}) = [\mathbf{D}_L(\mathbf{x}^{(n,m-1)}) + \mathbf{D}_R(\mathbf{x}^{(n)})]^{-1}$. This new iteration (7) utilizes the *same* regularizer gradient for all subsets. Compared to (6), the updates in (7) are quite similar except for an extra term that compensates for not updating the regularizer gradient. Table 3 summarizes the proposed algorithm. It requires storing the previous image \mathbf{x}^{last} .

In above description, we updated the regularizer gradient only after *all* subsets were updated. Obviously, the regularizer gradient can be updated as frequently as needed and we denote the update frequency as U_f (see algorithm in Table 3). Updating the regularizer gradient less often will reduce the computational cost at the expense of the convergence rate in early iterations.

The proposed algorithm was evaluated on a PWLS image reconstruction problem for cone-beam X-ray CT. We have investigated the trade-off between the update frequency and the computational expense per iteration.

Table 3. General PL Ordered Subsets Method with Double Surrogate

Initialize $\mathbf{x}^{(0)}$	
for $n = 0, 1, \dots,$	$\mathbf{x}^{(n)} := \mathbf{x}^{(n,0)}, \quad \mathbf{x}^{\text{last}} := \mathbf{x}^{(n)}$
for $m = 1, \dots, M$	
if $\text{mod}(m - 1, U_f) = 0$	$\mathbf{x}^{\text{last}} = \mathbf{x}^{(n,m-1)}$
end	
$\mathbf{x}^{(n,m)} = \arg \min_{\mathbf{x}} \phi_m(\mathbf{x}; \mathbf{x}^{(n,m-1)}, \mathbf{x}^{\text{last}})$	
end	
$\mathbf{x}^{(n+1)} := \mathbf{x}^{(n,M)}$	
end	

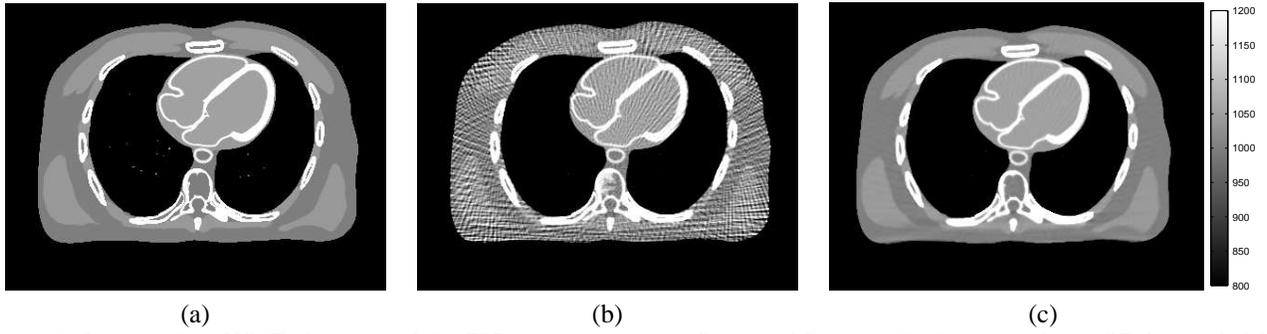


Figure 1. Images of (a) XCAT phantom and (b) FBP reconstruction (c) Converged Image (\mathbf{x}^∞) from cone-beam CT data with 164 projection views.

3. RESULTS

The proposed algorithm was investigated on a 3D cone-beam CT image reconstruction problem with limited view angles. The XCAT phantom¹⁸ was used as a static object, and the image was reconstructed to a $512 \times 512 \times 50$ grid with pixel size $\Delta_x = \Delta_y = 0.9766$ mm and $\Delta_z = 0.625$ mm.

We simulated a 3rd-generation axial cone-beam CT system using the separable footprint projector.¹⁹ The simulated system has $N_s = 888$ channels and $N_t = 32$ detector rows spaced by $\Delta_s = 1.0239$ mm and $\Delta_t = 1.09878$ mm per view, and 164 evenly spaced view angles over 360° , which corresponds to an undersampling factor of 6. The source to detector distance was 949 mm, and the source to rotation center distance was 541 mm. We also included a quarter detector offset in the channel direction to reduce aliasing.

For the edge-preserving regularizer, we used a certainty-based penalty²⁰ to obtain more uniform resolution and a q-Generalized Gaussian MRF (q-GGMRF)²¹ as the penalty function to provide edge-preserving regularizer. The regularization parameter β was selected such that the target PSF has a full-width at half-maximum (FWHM) of approximately 1.4 mm.²⁰ We generated the noisy sinogram with Poisson noise, and used weighting $w_i = \exp(-[A\mathbf{x}]_i)$.

To assess the convergence speed of the proposed method, we computed the normalized mean squared difference between the image estimate at the n th iteration, \mathbf{x}_n , with both the fully converged solution, \mathbf{x}^∞ , and true image, \mathbf{x}^{true} . We calculated \mathbf{x}^∞ using consecutive steps of ordered-subsets with decreasing number of subsets. We used 41, 10, and 1 subsets with 100, 100, and 1000 iterations respectively.

This type of “relaxed” OS is guaranteed to converge because the final stage uses just 1 subset for which (6) is convergent.⁴ Fig. 1 shows the images of true phantom, Filtered Back Projection (FBP) reconstruction,²² and fully converged image (\mathbf{x}^∞). For the FBP reconstruction, the ramp filter was associated with a Hanning apodization window of 2048-point in length to attenuate the high frequency noise. Due to limited view angles, conventional FBP reconstruction, which was used as the initial condition for our OS reconstruction, shows severe streaking artifacts compared to the true phantom. The fully converged image has much less artifacts, thus illustrating the benefits of statistical image reconstruction in limited view

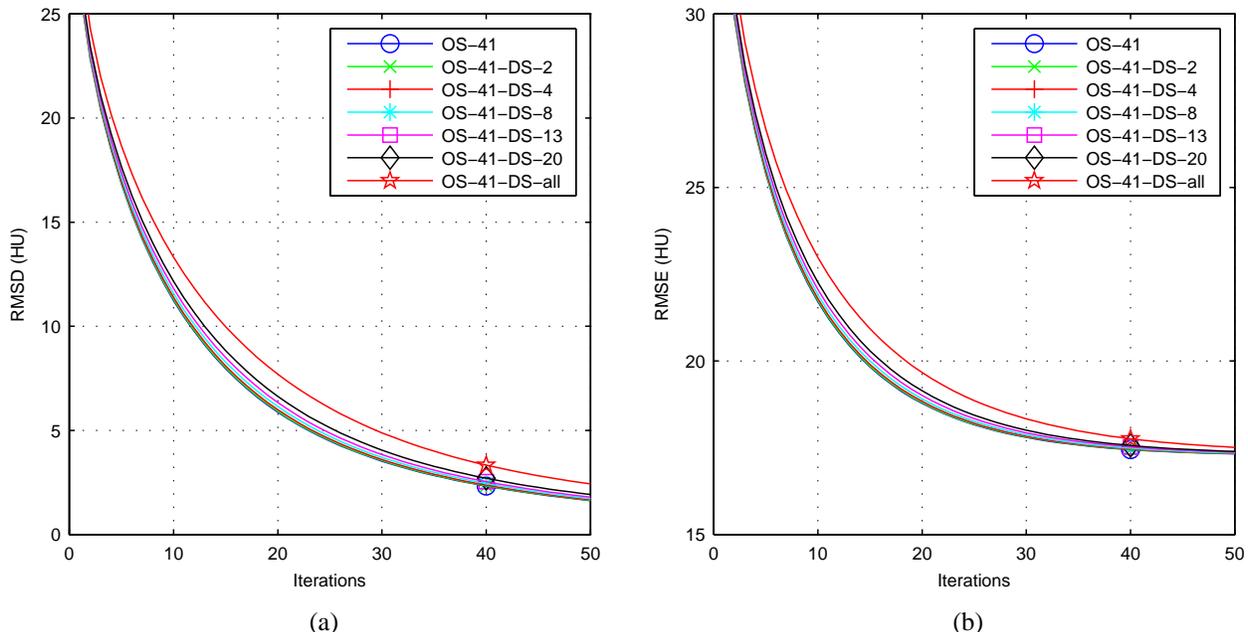


Figure 2. Convergence rates at each iteration for different regularizer update frequency. OS-41-DS-n indicates OS with 41 subsets and $U_f = n$, and “n = all” means only updating once after all subset updates are done. (a) $x^{(n)}$ with respect to x^∞ and (b) $x^{(n)}$ with respect to x^{true}

angle problems.

We divided the projections into 41 subsets, which corresponds to 4 views per subset. This is a rather aggressive selection compared to conventional choices to try to accelerate convergence significantly.

The regularizer gradient was updated at different frequencies to see its effect to convergence and computation time. Fig. 2 illustrates that our proposed method gives similar root-mean-square (RMS) differences, $\sqrt{\frac{1}{N} \sum_{j=1}^N (x_j^{(n)} - x_j^\infty)^2}$, as the conventional OS even when we update the regularizer infrequently per iteration.

On the other hand, Fig. 3 illustrates that the computational expense required to obtain the same level of RMS differences was reduced by the proposed method. With $U_f = 13$, which gives the best result in this case, the proposed method converges about 3 times faster than the conventional OS. By observing the reconstructed image at the same time point, Fig. 4 clearly shows that the proposed method is converging faster. There exists tradeoff between convergence rates and computational expense, and for our case calculating the regularizer gradient for every 13 subset updates gave the most efficient results. For different problems, the optimal update frequency will differ. However, it is noticeable that regardless of the update frequency, the proposed method is converging faster than the conventional OS. As the problem gets larger and the number of subsets increase, the computational expenses required to calculate the gradient of the regularizer becomes much more dominant. Therefore, we can expect substantial benefits from our proposed method for such problems.

4. CONCLUSIONS

We presented a simple modification of the conventional OS method that allows the regularizer gradient to be less frequently updated. The method provides very significant acceleration when applied to large problems with many subsets, such as cone-beam or helical CT image reconstruction, while still providing good reconstructed images.

Simulations demonstrated that a good reconstruction with compatible quality was achieved within much less computation time. However, the reduction in the computational expense may depend on the size of the problem and the complexity of updating the regularizer compared to the complexity of forward and back-projection. When implemented on graphics processing unit (GPU), the trade-offs may differ.

We only focussed on the initial convergence characteristics of the ordered-subsets method. Since the conventional OS algorithm is inherently not globally convergent, neither is the proposed method. We can apply the double surrogate idea to

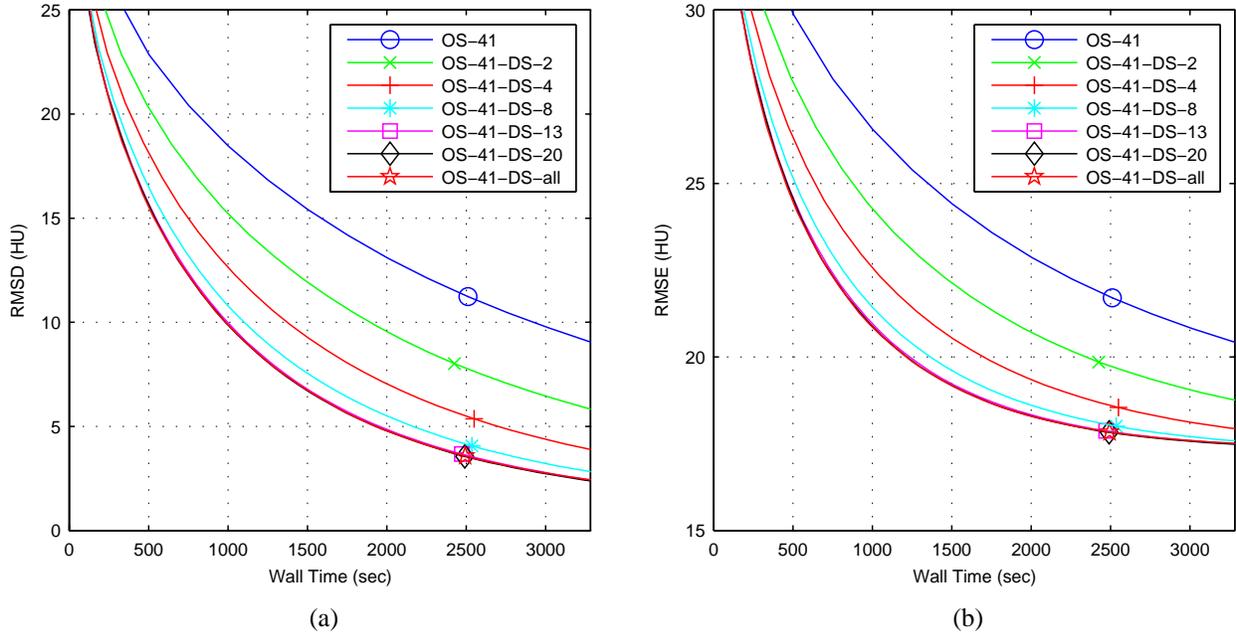


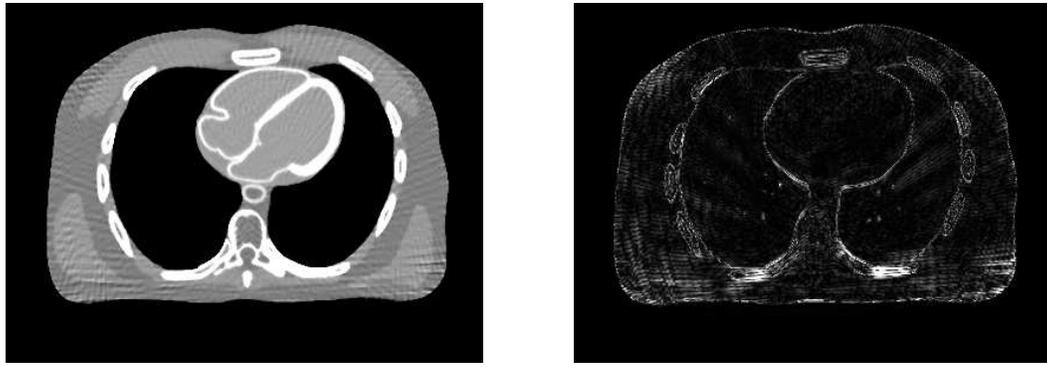
Figure 3. Convergence rates versus time for different regularizer update frequency. (a) $\mathbf{x}^{(n)}$ with respect to \mathbf{x}^{∞} and (b) $\mathbf{x}^{(n)}$ with respect to \mathbf{x}^{true}

other block iterative algorithms like incremental optimization transfer (IOT), which is globally convergent.²³ Further research will address the systematic way to determine the optimal update frequency and extension of the idea to other algorithms.

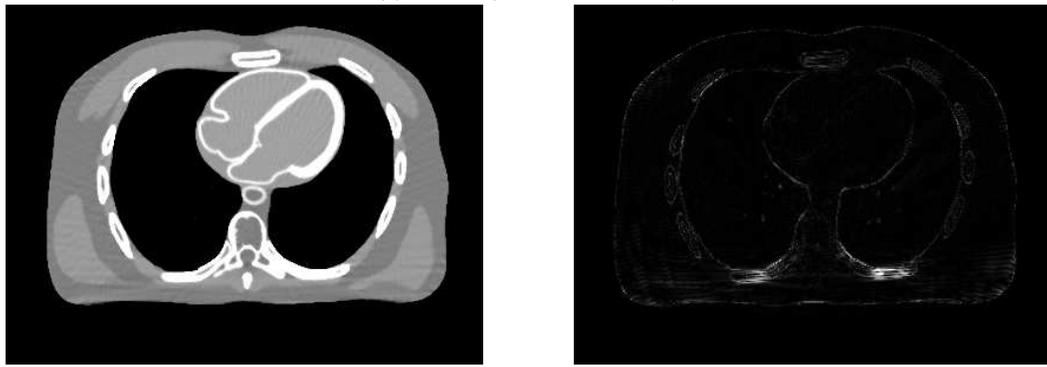
REFERENCES

- [1] Xu, W. and Mueller, K., "Evaluating popular non-linear image processing filters for their use in regularized iterative CT," in [*Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*], 2864–5 (2010).
- [2] Mustafovic, S. and Thielemans, K., "Object dependency of resolution in reconstruction algorithms with inter-iteration filtering applied to PET data," *IEEE Trans. Med. Imag.* **23**, 433–46 (Apr. 2004).
- [3] Hudson, H. M. and Larkin, R. S., "Accelerated image reconstruction using ordered subsets of projection data," *IEEE Trans. Med. Imag.* **13**, 601–9 (Dec. 1994).
- [4] Erdoğan, H. and Fessler, J. A., "Ordered subsets algorithms for transmission tomography," *Phys. Med. Biol.* **44**, 2835–51 (Nov. 1999).
- [5] Nocedal, J., "Large scale unconstrained optimization," in [*The State of the Art in Numerical Analysis*], Watson, A. and Duff, I., eds., Oxford University Press, Oxford (1996).
- [6] Kibardin, V. M., "Decomposition into functions in the minimization problem," *Avtomatika i Telemekhanika* **9**, 66–79 (Sept. 1979). Translation: p. 1311–23 in Plenum Publishing Co. "Adaptive Systems".
- [7] Bertsekas, D. P., "A new class of incremental gradient methods for least squares problems," *SIAM J. Optim.* **7**, 913–26 (Nov. 1997).
- [8] Kamphuis, C. and Beekman, F. J., "Accelerated iterative transmission CT reconstruction using an ordered subsets convex algorithm," *IEEE Trans. Med. Imag.* **17**, 1001–5 (Dec. 1998).
- [9] Herman, G. T. and Meyer, L. B., "Algebraic reconstruction techniques can be made computationally efficient," *IEEE Trans. Med. Imag.* **12**, 600–9 (Sept. 1993).
- [10] Hamaker, C. and Solmon, D. C., "The angles between the null spaces of X rays," *J. Math. Anal. Applic.* **62**, 1–22 (Jan. 1978).
- [11] De Pierro, A. R. and Yamagishi, M. E. B., "Fast EM-like methods for maximum 'a posteriori' estimates in emission tomography," *IEEE Trans. Med. Imag.* **20**, 280–8 (Apr. 2001).
- [12] Kudo, H., Nakazawa, H., and Saito, T., "Convergent block-iterative method for general convex cost functions," in [*Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med.*], 247–250 (1999).
- [13] Neal, R. and Hinton, G. E., "A view of the EM algorithm that justifies incremental, sparse and other variants," in [*Learning in Graphical Models*], Jordan, M. I., ed., 255–68, Kluwer, Dordrecht (1998).
- [14] Blatt, D., Hero, A. O., and Gauchman, H., "A convergent incremental gradient method with a constant step size," *SIAM J. Optim.* **18**(1), 29–51 (2007).

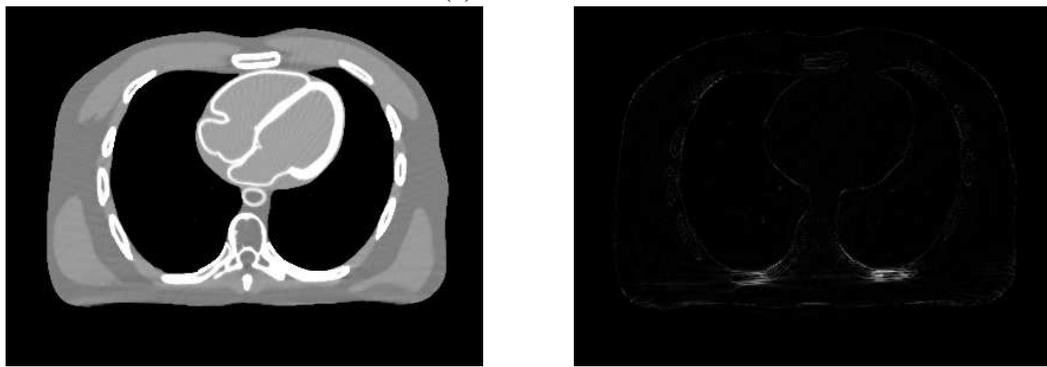
- [15] Huber, P. J., [*Robust statistics*], Wiley, New York (1981).
- [16] Jacobson, M. W. and Fessler, J. A., "An expanded theoretical treatment of iteration-dependent majorize-minimize algorithms," *IEEE Trans. Im. Proc.* **16**, 2411–22 (Oct. 2007).
- [17] Hunter, D. R. and Lange, K., "A tutorial on MM algorithms," *American Statistician* **58**, 30–7 (Feb. 2004).
- [18] Segars, W. P., Mahesh, M., Beck, T. J., Frey, E. C., and Tsui, B. M. W., "Realistic CT simulation using the 4D XCAT phantom," *Med. Phys.* **35**, 3800–8 (Aug. 2008).
- [19] Long, Y., Fessler, J. A., and Balter, J. M., "3D forward and back-projection for X-ray CT using separable footprints," *IEEE Trans. Med. Imag.* **29**, 1839–50 (Nov. 2010).
- [20] Fessler, J. A. and Rogers, W. L., "Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs," *IEEE Trans. Im. Proc.* **5**, 1346–58 (Sept. 1996).
- [21] Thibault, J.-B., Sauer, K., Bouman, C., and Hsieh, J., "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," *Med. Phys.* **34**, 4526–44 (Nov. 2007).
- [22] Feldkamp, L. A., Davis, L. C., and Kress, J. W., "Practical cone beam algorithm," *J. Opt. Soc. Am. A* **1**, 612–9 (June 1984).
- [23] Ahn, S., Fessler, J. A., Blatt, D., and Hero, A. O., "Convergent incremental optimization transfer algorithms: Application to tomography," *IEEE Trans. Med. Imag.* **25**, 283–96 (Mar. 2006).



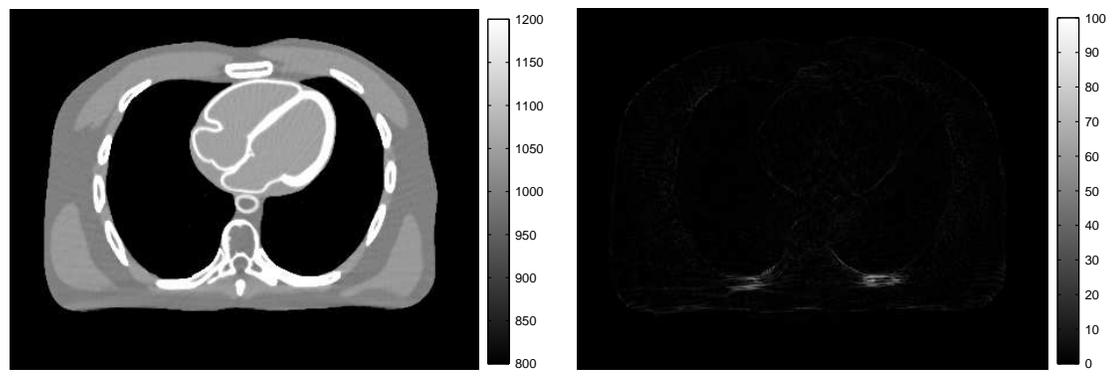
(a) OS-41 (conventional OS)



(b) OS-41-DS-4



(c) OS-41-DS-13



(d) OS-41-DS-all

Figure 4. Comparing the convergence speed of OSDS with different update frequencies. Left column: Images at the same time point (4000 sec after initialization). Right column: Absolute difference images with respect to $\mathbf{x}^{(\infty)}$.