Cine cone beam CT reconstruction using low-rank matrix factorization: algorithm and a proof-of-principle study

Jian-Feng Cai[*]   Xun Jia[†]   Hao Gao[‡]   Steve B. Jiang[§]   Zuowei Shen[¶]

Hongkai Zhao[‖]

Abstract

Respiration-correlated CBCT, commonly called 4DCBCT, provides respiratory phase-resolved CBCT images. A typical 4DCBCT consists of a small number (usually 10 or less) of CBCT images, each corresponding to a particular breathing phase and reconstructed using all projections falling in that phase bin over the 4DCBCT scan. Therefore, 4DCBCT represents averaged patient images over one breathing cycle and the 4th dimension is actually breathing phase instead of time. In many clinical applications, it is desirable to obtain true 4DCBCT with the 4th dimension being time, i.e., each constituent CBCT image corresponds to an instantaneous projection. Theoretically it is impossible to reconstruct a CBCT image from a single projection. However, if all the constituent CBCT images of a 4DCBCT scan share a lot of redundant information and can be represented with or approximated by linear combinations of a small number of basis images, it might be possible to make a good reconstruction of these images by exploring sparsity and coherence/redundancy of these CBCT images. These CBCT images are not completely time resolved since they are represented by a small number of basis images, but they can exploit both local and global temporal coherence of the patient anatomy automatically and contain much more temporal variation information of the patient geometry than the conventional 4DCBCT. We propose in this work a computational model and algorithms for the reconstruction of this type of semi-time-resolved CBCT, called cine-CBCT, based on low rank approximation that can utilize the underlying temporal coherence both locally and globally, such as slow variation, periodicity or repetition, in those cine-CBCT images. In particular, let each column of the matrix U represent a CBCT image to be reconstructed and the total number of columns is the same as the number of projections, the key assumption of our method is that U can be approximated well by a matrix whose rank is much smaller than the number of projections in the matrix factorization form $U \approx LR$. The columns of L forms the basis for all cine-CBCT images. The number of columns for the matrix L constraints the rank of approximation and hence implicitly imposing a temporal coherence condition among all the images in cine-CBCT. We further impose desired image properties by enforcing sparsity of the tight wavelet frame transform of L. In this matrix factorization formulation both the low rank approximation and the sparse representation in wavelet transform domain reduces the true degree of freedom of U significantly. A split Bregman method is used to solve the problem. In this paper we focus on presenting this new algorithm and showing the proof of principle using simulation studies on an NCAT phantom.

1   Introduction

When cone beam computed tomography (CBCT) is applied to thorax or upper abdomen regions, motion-induced artifacts, such as blurring or distortion, become a serious problem, because different x-ray projections

[*]Department of Mathematics, The University of Iowa, Iowa City, IA 52242, USA

[†]Department of Radiation Oncology, University of Texas Southwestern Medical Center, Dallas, TX 75235

[‡]School of Biomedical Engineering and Department of Mathematics, Shanghai Jiao Tong University, Shanghai 200240, CHINA

[§]Department of Radiation Oncology, University of Texas Southwestern Medical Center, Dallas, TX 75235

[¶]Department of Mathematics, National University of Singapore, Singapore 119076

[‖]Department of Mathematics, University of California Irvine, Irvine, CA 92697, USA

correspond to different volumetric CBCT images due to patient respiratory motion. To overcome this problem, four-dimensional CBCT (4DCBCT) [17, 21, 30] has been developed. In such a modality, all x-ray projections are first retrospectively grouped into different respiratory phase bins according to a breathing signal tagged on every projection image. A set of CBCT images are then reconstructed, each at a breathing phase, under the assumption that the projections placed into each bin correspond to the same or similar CBCT image. The number of phase bins is usually empirically chosen by a user based on the consideration of balancing temporal resolution and image quality. On one hand, a high temporal resolution requires a large number of phase bins, which leads to insufficient number of projections available to each phase and hence degraded quality of reconstructed CBCT images. On the other, a small number of phase bins results in low temporal resolution, as well as a relatively large bin width and the associated residual motion artifacts in the reconstructed images. To maintain clinically acceptable temporal resolution and image quality, 4DCBCT acquisition protocols such as slow gantry rotations and multiple gantry rotations have been proposed [20–22], and novel reconstruction or image processing techniques have been employed [1, 2, 25, 26].

Conventional 4DCBCT reconstruction approach reconstructs CBCT images at different phases independent of each other [30]. This straightforward method neglects the temporal correlations of CBCT images at different phases. In contrast, a collaborative reconstruction scheme has been recently proposed [11, 14, 16, 24, 29, 31]. Among them, [14] and [31] proposed a temporal non-local means method to constraint that the reconstructed CBCT images at neighboring phases must contain repetitive anatomical features. The locations of these features are allowed to vary among phases. [29] and [24] first reconstruct an average CBCT image using all projections at all phases and impose a similarity constraint between this average image and the CBCT at each phase. These methods explicitly enforces the similarity among reconstructed CBCT images. [11] utilizes robust PCA techniques for the 4DCBCT reconstruction problem. This method restores a matrix whose columns are the 4DCBCT images at all phases under the assumption that, after a proper transformation, this matrix can be decomposed as the sum of a low-rank matrix corresponding to an almost static background and a sparse matrix corresponding to a moving foreground. The low-rank constraint achieved by minimizing the nuclear norm [5] of the associated matrix implicitly imposes the inter-phase similarity of the background. However, in the context of anatomical motion in thorax or upper abdomen, there is no clear distinction between foreground and background. Besides, as pointed out in [9], the nuclear norm minimization always finds a matrix factorization with mutually orthogonal factors. In this sense, it may have some unfavourable bias on the resulting low-rank matrices.

Although commonly called 4DCBCT, respiration-correlated CBCT [10] is a more accurate name for this imaging technique, since its 4th dimension is actually the respiratory phase rather than time. Clinically, it is more preferable to reconstruct true 4DCBCT with the 4th dimension being time, i.e., one CBCT image is reconstructed based on the corresponding instantaneous projection, if possible. This time-resolved CBCT, called cine-CBCT in this work to avoid confusion, encounters an apparent technical barrier of reconstructing a CBCT image based on only one projection where the information needed for a good reconstruction is severely insufficient. One way to overcome this barrier is to borrow some useful information of a particular patient's anatomy from prior images. Previously we have tried to utilize prior 4DCT images of the same patient [18, 19]. Particularly, we have built a lung motion model for the patient by performing a principle component analysis(PCA) on the motion vector fields obtained from deformable registration on prior 4DCT images. We discovered that only a few principle components are sufficient to represent the lung motion to a satisfactory degree of accuracy. The reconstruction of a volumetric image based on one acquired x-ray projection is then achieved by finding those PCA coefficients, so that the projection of the corresponding CBCT matches the acquired one. However, this method heavily depends on how similar the current CBCT images are to the prior 4DCT images and how accurately the PCA parameters can be determined using one projection.

In this paper, we will propose another way to overcome this technical barrier without using any prior images. We will show the principle of proof of a computational model and numerical algorithms for cine-CBCT reconstruction based on a low rank assumption of the matrix which includes all cine-CBCT images as its columns. The low rank approximation, which uses a form of matrix factorization, maximally and automatically exploits the temporal coherence of the patient anatomy locally and globally without using any

prior knowledges. Moreover, sparsity condition is also enforced under appropriate transforms, *i.e.* wavelet transform in space, to satisfy desired properties for the sequence of CBCT images. Simulation studies have demonstrated promising results for our method for this challenging problem.

2   Methods and Materials

2.1   Assumption

Let us denote a patient CBCT image at time t by $x(t)$, which can be viewed as a point embedded in a high dimensional space R^N with its pixel values being the coordinate, where N is the total number of pixels in each CBCT image. $x(t)$ travels along a smooth trajectory as patient breathes. A cine-CBCT reconstruction problem tries to recover points on this trajectory, each corresponding to a moment at an x-ray projection measurement. The key assumption in our approach is that the trajectory stays piecewisely close to a low dimensional subspace in the very high dimensional embedding space. Equivalently, when we form a matrix $\boldsymbol{U}$ with each column being an CBCT image, the resulting matrix is approximately of low-rank. We believe this assumption is valid due to the fact that body motion is far from a free form motion, which is verified by the fact that only a few principle components are sufficient to represent the lung motion to a satisfactory degree of accuracy as studied in [18,19]. Based on this assumption, we use a matrix factorization formulation $\boldsymbol{U} \approx \boldsymbol{LR}$. The number of columns for the matrix $\boldsymbol{L}$ constraints the rank of approximation. We further impose desired image properties by enforcing sparsity of the tight wavelet frame transform of $\boldsymbol{L}$. In this matrix factorization formulation both the low rank approximation and the sparse representation in wavelet transform domain reduces the true degree of freedom of $\boldsymbol{U}$ significantly. In all previous approaches mentioned in the Section of Introduction, reduction of degrees of freedom are exploited explicitly or implicitly from redundancy and coherence in space and time as well as sparse representation. In particular, 4DCBCT based on phase binning explicitly used the low rank assumption, i.e., the rank is equal to the number of phases. However, as mentioned before, the number of phases, i.e., the rank, is usually empirically chosen and more importantly coherence in time or among different phases are not explored at all.

Another important practical issue is how well the small number of image basis for the low dimensional subspace and the factorization $\boldsymbol{U} \approx \boldsymbol{LR}$ can be recovered from the projection data. Intuitively, this requires each basis image or equivalently all cine-CBCT images collectively have a sufficient number of projections covering a sufficiently large scanning angle, which is also required for the construction of a single CBCT image. As shown in experiments in Section 4.3, reconstruction from data using different projections can produce different results for the same ground truth.

2.2   Model

In this proof-of-principle study, we consider the reconstruction of a 2D slice of the CBCT to illustrate the principles of the new algorithm. The basic ideas, however, can be easily generalized into 3D contexts. Let the unknown 2D image be a function u defined in $\mathbb{R}^2$. When the X-ray source is placed with angle θ, the projection measured by the imager at location z is

$$f_\theta(z) = \mathcal{P}_\theta u(z) := \int_0^{\ell_z} u(\boldsymbol{x}_\theta + s\boldsymbol{r}_z)ds \tag{1}$$

where $\boldsymbol{x}_\theta \in \mathbb{R}^2$ is the coordinate of the X-ray source, and $\boldsymbol{r}_z \in \mathbb{S}^2$ and ℓ_z are respectively the direction and the length of the line connecting the X-ray source and the location z on the imager. The operator $\mathcal{P}$ in (1) is also known as Radon transform. If $f_\theta(z)$ is sampled with respect to z, the resulting projection data can be written as a vector $\boldsymbol{f}_\theta \in \mathbb{R}^M$ that obeys

$$\boldsymbol{P}_\theta \boldsymbol{u} = \boldsymbol{f}_\theta. \tag{2}$$

Here $\boldsymbol{P}_\theta \in \mathbb{R}^{M \times N}$ and $\boldsymbol{u} \in \mathbb{R}^N$ are the discretization of $\mathcal{P}_\theta$ and u in (1) respectively. Suppose that we have T projections where the X-ray source is placed with angles $\theta_1, \theta_2, \ldots, \theta_T$ respectively. Let $\boldsymbol{P}_i$ and $\boldsymbol{f}_i$ stand

for $\boldsymbol{P}_{\theta_i}$ and $\boldsymbol{f}_{\theta_i}$ respectively. Then, by putting (2) with different angles together, the CT projection can be written into a system of linear equations

$$\boldsymbol{P}\boldsymbol{u} = \boldsymbol{f}, \tag{3}$$

where $\boldsymbol{P} = [\boldsymbol{P}_1; \ \boldsymbol{P}_2; \ \ldots; \ \boldsymbol{P}_T] \in \mathbb{R}^{MT \times N}$ and $\boldsymbol{f} = [\boldsymbol{f}_1; \ \boldsymbol{f}_2; \ \ldots; \ \boldsymbol{f}_T] \in \mathbb{R}^{MT}$. In other words, the CT image reconstruction problem is to recover image $\boldsymbol{u}$ from its partial Radon transform.

In the context of cine-CBCT reconstruction, instead of only one unknown image u as in the CBCT problem, there is a set of unknown images, each associated with a projection, denoted by $\{\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_T\}$ and the projection condition in (3) may be modified to

$$\boldsymbol{P}_i \boldsymbol{u}_i = \boldsymbol{f}_i, \quad i = 1, 2, \ldots, T. \tag{4}$$

If we write all the unknown images in a matrix form $\boldsymbol{U} = [\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_T]$, (4) can be rewritten into a compact matrix equation

$$\mathcal{P}\boldsymbol{U} = \boldsymbol{F} \tag{5}$$

where $\mathcal{P}\boldsymbol{U} = [\boldsymbol{P}_1 \boldsymbol{u}_1, \ldots, \boldsymbol{P}_T \boldsymbol{u}_T]$ and $\boldsymbol{F} = [\boldsymbol{f}_1, \ldots, \boldsymbol{f}_T]$.

With low rank approximation assumption for $\boldsymbol{U}$ as discussed in Section 2.1, the number of intrinsic unknowns of (5) are possibly less than the number of measurements, and therefore it is possible to reconstruct T CBCT images in $\boldsymbol{U}$ from T projections. To incorporate the underlying low-rank assumption into the cine-CBCT reconstruction process, we would like to explore matrix factorization based low-rank models in this study. More specifically, we enforce a decomposition form of the unknown matrix $\boldsymbol{U}$ of images as $\boldsymbol{U} = \boldsymbol{L}\boldsymbol{R}$ where $\boldsymbol{L} \in \mathbb{R}^{N \times K}$ and $\boldsymbol{R} \in \mathbb{R}^{K \times T}$ for a small integer K. From basic linear algebra knowledge, any matrix $\boldsymbol{U}$ of rank K can be represented in $\boldsymbol{U} = \boldsymbol{L}\boldsymbol{R}$ and conversely, the rank of $\boldsymbol{U}$ is at most K given the factorization form. In our algorithm, the value of K is specified by the user and our method regarding the selection of K will be presented later. We would like to point out that there is another popular approach for imposing the low rank condition that penalizes the nuclear norm of the matrix $\boldsymbol{U}$, namely the sum of the singular values. Compare these two methods, the latter always finds orthogonal basis because of the involved singular value decomposition process. In contrast, our factorization method does not require this implicit orthogonality and hence attains the advantage of avoiding unfavourable bias on the resulting low-rank matrices [9]. In some sense, the approach of penalizing the nuclear norm of $\boldsymbol{U}$ is like Principle Component Analysis and our approach is like Independent Component Analysis. As a consequence, each of our basis, i.e., each column of $\boldsymbol{L}$, represents more meaningful images and hence we can enforce some desired image properties on $\boldsymbol{L}$, such as sparsity under wavelet transform.

This matrix factorization approach also allows for the additional regularizations on $\boldsymbol{L}$ and $\boldsymbol{R}$, so that $\boldsymbol{U} = \boldsymbol{L}\boldsymbol{R}$ carries desirable physical properties. First, the columns of the matrix $\boldsymbol{L}$, corresponding to images, can be interpreted as a basis that efficiently represents all the columns of $\boldsymbol{U}$. Notice that images usually have sparse coefficients under suitable transforms such as wavelet tight frames [6, 27]. Let $\mathcal{D}$ be such a transform. Since the columns of $\boldsymbol{L}$ are images, we want $\mathcal{D}\boldsymbol{L}$ to be sparse. It is well-known that ℓ_1-norm minimization leads to sparse solutions. Therefore, the ℓ_1-norm $\|\mathcal{D}\boldsymbol{L}\|_1$ is penalized. Second, the rows of $\boldsymbol{R}$ are the coefficients of $\boldsymbol{U}$ under the basis $\boldsymbol{L}$, which reflect the respiratory motion. As the respiratory motion varies and could be irregular sometimes, there is no universal regularization that is suitable for all kinds of respiratory motion. We simply penalize its Frobenius norm $\|\boldsymbol{R}\|_F^2 := \sum_{i,j} \boldsymbol{R}_{ij}^2$ to balance the energy of $\boldsymbol{L}$ and $\boldsymbol{R}$. By enforcing sparse representation of $\boldsymbol{L}$ in wavelet domain, the intrinsic degrees of freedom are further reduced. Both low rank and sparse conditions can be regarded as some kind of regularization that makes the whole problem less under-determined or ill-posed.

Altogether, we propose to reconstruct the cine-CBCT by solving an optimization problem

$$\min_{\boldsymbol{L}, \boldsymbol{R}} \lambda \|\mathcal{D}\boldsymbol{L}\|_1 + \|\boldsymbol{R}\|_F^2, \quad \text{s.t. } \mathcal{P}(\boldsymbol{L}\boldsymbol{R}) = \boldsymbol{F}, \tag{6}$$

In practice, (5) is barely satisfied, because there is always unavoidable noise in the measurements $\boldsymbol{F}$. Moreover, in reality the matrix $\boldsymbol{U}$ is only approximately low-rank and explicitly enforcing a low-rank representation leads to error. For these considerations, we solve

$$\min_{\boldsymbol{L}, \boldsymbol{R}} \lambda \|\mathcal{D}\boldsymbol{L}\|_1 + \|\boldsymbol{R}\|_F^2, \quad \text{s.t. } \|\mathcal{P}(\boldsymbol{L}\boldsymbol{R}) - \boldsymbol{F}\|_F^2 \leq \sigma^2, \tag{7}$$

4

where $\|\cdot\|_F$ is the Frobenius norm, namely $\|\boldsymbol{A}\|_F = \sqrt{\sum_{i,j} \boldsymbol{A}_{i,j}^2}$ for a matrix $\boldsymbol{A}$. In (7), σ is a parameter to control to what extent the violation of (5) is allowed.

2.3 Algorithm

Let us first consider the algorithm for solving Equation (6). We use a split Bregman method [8, 12] (also known as augmented Lagrangian method) to solve this problem. The augmented Lagrangian of (6) is

$$E(\boldsymbol{L}, \boldsymbol{R}, \boldsymbol{Z}) = \lambda\|\mathcal{D}\boldsymbol{L}\|_1 + \|\boldsymbol{R}\|_F^2 + \langle \boldsymbol{Z}, \mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\rangle + \frac{\mu}{2}\|\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\|_F^2 \tag{8}$$

where $\langle \cdot, \cdot \rangle$ is the inner product, and $\boldsymbol{Z}$ is Lagrange multipliers. With appropriate fixed $\boldsymbol{Z}$, an optimal $\boldsymbol{L}, \boldsymbol{R}$ can be found by simply minimizing $E(\boldsymbol{L}, \boldsymbol{R}, \boldsymbol{Z})$ with respect to $(\boldsymbol{L}, \boldsymbol{R})$. Therefore, the trick is to determine $\boldsymbol{Z}$. In the augmented Lagrangian algorithm, we use

$$\begin{cases} \begin{cases} \boldsymbol{L} \leftarrow \arg\min_{\boldsymbol{L}} E(\boldsymbol{L}, \boldsymbol{R}, \boldsymbol{Z}), \\ \boldsymbol{R} \leftarrow \arg\min_{\boldsymbol{R}} E(\boldsymbol{L}, \boldsymbol{R}, \boldsymbol{Z}), \end{cases} \\ \boldsymbol{Z} \leftarrow \boldsymbol{Z} + (\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}). \end{cases} \tag{9}$$

The two subproblems in (9) are solved by split Bregman algorithm again in a similar fashion to (8) and (9). Let us take the first one as an example. By introducing an auxiliary variable $\boldsymbol{C}$, the first subproblem in (9) is equivalent to

$$\min_{\boldsymbol{L}} \|\boldsymbol{C}\|_1 + \langle \boldsymbol{Z}, \mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\rangle + \frac{\mu}{2}\|\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\|_F^2, \quad \text{s.t. } \boldsymbol{C} = \mathcal{D}\boldsymbol{L}, \tag{10}$$

and the associated augmented Lagrangian is

$$E_1(\boldsymbol{L}, \boldsymbol{C}, \boldsymbol{Z}_1) = \lambda\|\boldsymbol{C}\|_1 + \langle \boldsymbol{Z}, \mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\rangle + \frac{\mu}{2}\|\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\|_F^2 + \langle \boldsymbol{Z}_1, \boldsymbol{C} - \mathcal{D}\boldsymbol{L}\rangle + \frac{\mu_1}{2}\|\boldsymbol{C} - \mathcal{D}\boldsymbol{L}\|_F^2.$$

Then, a split Bregman method for solving the first subproblem in (9) is as follows

$$\begin{cases} \boldsymbol{L} \leftarrow \arg\min_{\boldsymbol{L}} E_1(\boldsymbol{L}, \boldsymbol{C}, \boldsymbol{Z}_1), \\ \boldsymbol{C} \leftarrow \arg\min_{\boldsymbol{C}} E_1(\boldsymbol{L}, \boldsymbol{C}, \boldsymbol{Z}_1), \\ \boldsymbol{Z}_1 \leftarrow \boldsymbol{Z}_1 + (\boldsymbol{C} - \mathcal{D}\boldsymbol{L}). \end{cases} \tag{11}$$

All steps involved in (11) can be solved easily by either entry-wise soft-thresholding (for $\boldsymbol{C}$) and linear equation solvers (for $\boldsymbol{L}$), e.g. conjugate gradient method. We use the same method for solving the second subproblem in (9). The full algorithm is summarized in Algorithm 1, where $\mathcal{T}$ is the soft-thresholding operator defined by $[\mathcal{T}_\alpha(\boldsymbol{A})]_{ij} = \text{sign}([\boldsymbol{A}]_{ij}) \cdot \max\{|[\boldsymbol{A}]_{ij}| - \alpha, 0\}$.

For any true or semi time resolved CBCT reconstruction, both storage and computational cost is a challenge. An important feature of our formulation and computational algorithm is to explore the low rank property. In our numerical computation, the huge matrix is stored and manipulated in a low-rank decomposition form, which can significantly reduce the memory size and computational cost. In particular, the computational cost of Algorithm 1 is mainly on solving linear equations in Step 1(a)(i) and 1(b)(i), where the conjugate gradient (CG) algorithm is employed. For moderately large μ_1/μ and μ_2/μ, the condition number of these linear equations is small, and therefore only a few iterations are necessary for the convergence of the CG algorithm. Consequently, each step in Algorithm 1 can be done in several matrix-vector products. The computational time for the reconstruction of images with 128×128 pixels and 360 projections is typically a few hours, on a computer with Core i7-2600 CPU at 3.40GHz, 8GB memory, and Matlab R2013a.

One complexity of this problem comes from the non-convexity of the objective function in (6). Since there may exist local minima, we have to make sure our algorithm get a desired solution. Yet, in our numerical experiments, this non-convexity issue is not found to be a problem for the following two reasons. First of all, our algorithm may find the global minimum of (6). As pointed out in [4], the augmented Lagrangian

Algorithm 1

1: Repeat for s_{out} times.

   (a) Solve the first subproblem in (9) by performing the following steps for s_{in} times.

      (i) Update $\boldsymbol{L} \leftarrow \arg\min_{\boldsymbol{L}} \frac{\mu}{2}\|\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F} + \boldsymbol{Z}/\mu\|_F^2 + \frac{\mu_1}{2}\|\mathcal{D}\boldsymbol{L} - \boldsymbol{C} - \boldsymbol{Z}_1/\mu_1\|_F^2$.

      (ii) Update $\boldsymbol{C} \leftarrow \mathcal{T}_{\lambda/\mu_1}(\mathcal{D}\boldsymbol{L} - \boldsymbol{Z}_1/\mu_1)$.

      (iii) Update $\boldsymbol{Z}_1 = \boldsymbol{Z}_1 + (\boldsymbol{C} - \mathcal{D}\boldsymbol{L})$.

   (b) Solve the second subproblem in (9) by performing the following steps for s_{in} times.

      (i) Update $\boldsymbol{R} \leftarrow \arg\min_{\boldsymbol{R}} \frac{\mu}{2}\|\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F} + \boldsymbol{Z}/\mu\|_F^2 + \frac{\mu_2}{2}\|\boldsymbol{R} - \boldsymbol{B} - \boldsymbol{Z}_2/\mu_2\|_F^2$.

      (ii) Update $\boldsymbol{B} \leftarrow (\boldsymbol{R} - \boldsymbol{Z}_2/\mu_2)/(1 + 1/\mu_2)$.

      (iii) Update $\boldsymbol{Z}_2 = \boldsymbol{Z}_2 + (\boldsymbol{B} - \boldsymbol{R})$.

2: If $\|\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\|_F^2$ is small enough, then return.
3: Update $\boldsymbol{Z} \leftarrow \boldsymbol{Z} + (\mathcal{P}(\boldsymbol{LR} - \boldsymbol{F}))$.
4: Goto Step 1.

---

algorithm attains the global minimum for non-convex objectives in low-rank factorization form under suitable assumptions. Similar algorithms to ours has been used in, e.g., SDPLR (Semi-definite programming via low rank factorization) [3]. Secondly, we have also chosen the initial guess to the iterative algorithm in the Algorithm 1 carefully as following. We first solve a convex minimization problem

$$\min_{\boldsymbol{U}} \frac{1}{2}\|\mathcal{P}\boldsymbol{U} - \boldsymbol{F}\|_F^2 + \lambda\|\boldsymbol{U}\|_*, \tag{12}$$

where $\|\boldsymbol{U}\|_*$ is the nuclear norm, i.e., the summation of the singular values of $\boldsymbol{U}$. The nuclear norm minimization is able to find a lowest-rank solution $\boldsymbol{U}_*$ of (5) within a precision [5]. Then, the initial guess in the Algorithm 1 is chosen to be the best rank-K approximation to $\boldsymbol{U}_*$. Let $\boldsymbol{U}_* = \boldsymbol{W}\boldsymbol{\Sigma}\boldsymbol{V}^T$ be a singular value decomposition. If the rank of $\boldsymbol{U}_*$ is less than K, we choose $\boldsymbol{L}^{(0)} = \boldsymbol{W}\boldsymbol{\Sigma}^{1/2}$ and $\boldsymbol{R}^{(0)} = \boldsymbol{\Sigma}^{1/2}\boldsymbol{V}^T$ as the initial guesses. If the rank of $\boldsymbol{U}_*$ exceeds K, we choose $\boldsymbol{L}^{(0)} = \boldsymbol{W}_K\boldsymbol{\Sigma}_K^{1/2}$ and $\boldsymbol{R}^{(0)} = \boldsymbol{\Sigma}_K^{1/2}\boldsymbol{V}_K^T$ as the initial guesses, where $\boldsymbol{W}_K, \boldsymbol{V}_K$ are the first K columns of $\boldsymbol{W}$ and $\boldsymbol{V}$ respectively, and $\boldsymbol{\Sigma}_K$ is the $K \times K$ principle submatrix of $\boldsymbol{\Sigma}$. By this way, $\boldsymbol{L}^{(0)}\boldsymbol{R}^{(0)}$ is the best rank-K approximation of $\boldsymbol{U}_*$.

As for the problem in (7), we still use the Algorithm 1, but we stop the iteration as soon as $\|\mathcal{P}(\boldsymbol{LR}) - \boldsymbol{F}\|_F^2 \leq \sigma^2$. By this way, we can get a quite good approximate solution to (9); see [23] for a detailed discussion. This approach has been previously used and discussed in similar mathematical problems but in other contexts, e.g., [7, 8, 23].

2.4 Choice of K

One practical issue in our algorithm is the selection of the parameter K to control the matrix rank. In practice, we achieve this goal by the following procedure.

We first set K to be a large enough number, e.g. $K = 20$, and run our algorithm. Recall that the columns of $\boldsymbol{L}$ forms a basis to represent the images in $\boldsymbol{U}$. Because of the sufficiently large K value in this trial run, the algorithm is forced to generate a set of K basis vectors containing those good ones for $\boldsymbol{U}$, as well as the unnecessary bad ones. This is reflected by the fact that some columns of $\boldsymbol{L}$ in the solution make significant contributions to $\boldsymbol{U}$, while the other columns contribute little. For those unnecessary columns in $\boldsymbol{L}$, their presence introduces a small but observable amount of signals into the finally reconstructed images

Figure 1: The ground truth cine-CBCT images. From left to right: frame $20, 81, 160, 256, 334$ The display window is $(-1000, 420)$ HU.

in $\boldsymbol{U}$ and compromises its quality. It is therefore desirable to eliminate them from the first place by setting the value of K to be the number of the significant columns. A second reconstruction is then performed with this properly chosen K value.

It remains to determine which columns are significant and which ones are negligible. Since the contribution of the i-th column of $\boldsymbol{L}$ to $\boldsymbol{U}$ is $\boldsymbol{L}(:,i)\boldsymbol{R}(i,:)$, any norm of this project reflects the importance of this basis vector $\boldsymbol{L}(:,i)$. The larger norm is, the more significant $\boldsymbol{L}(:,i)$ is. In the experiments, we use the ∞-norm to identify the significance, which is the maximum of row absolute sums of a matrix, or equivalently, the pixel that has the largest absolute sum over time. The column $\boldsymbol{L}(:,i)$ is considered to be negligible, if $\|\boldsymbol{L}(:,i)\boldsymbol{R}(i,:)\|_\infty$ is close to zero, and significant otherwise.

2.5 Experiments

We tested our cine-CBCT reconstruction algorithm on a digital NURBS-based cardiac-torso (NCAT) phantom [28], which generates a patient body in thorax region with detailed anatomical features and a realistic motion pattern. The patient respiratory period is 4 seconds. The CBCT gantry rotates about the patient at a constant speed for a full rotation in 59 seconds, in which 360 X-ray projections are acquired. At each X-ray projection acquisition, we compute the NCAT phantom image at the specific time point with a resolution of 128×128, and the projection is then computed using a ray-tracing algorithm at the associated projection angle with a detector resolution of 256 bins. The patient breathing period and the gantry rotation period are chosen to be incommensurate deliberately, so that all the 360 patient images are distinct, although some of them visually look close. Some of the underlying true images are shown in Figure 1. Under this setup, the size of $\boldsymbol{U}$ is 16384×360, and the size of $\boldsymbol{F}$ is 256×360.

Two experiments are performed in this feasibility study, namely the measured projections $\boldsymbol{F}$ generated from a *full-fan* and *half-fan* scan. In each case, we will first present the results for the selection of the parameter K and then the reconstruction results. Apart from visual inspection of the reconstructed cine-CBCT images, we have also quantitatively assessed the restored image quality using relative error as a metric defined by $\|\boldsymbol{L}\boldsymbol{R} - \boldsymbol{U}_{true}\|_F / \|\boldsymbol{U}_{true}\|_F$, where $\boldsymbol{L}, \boldsymbol{R}$ are the outputs of the proposed method, and $\boldsymbol{U}_{true}$ is the matrix consisting of ground truth NCAT images. Unless specified, the parameters used throughout our experiments are $\lambda = 4000$, $\mu = 1$, $\mu_1 = \mu_2 = 10^5$, and $s_{in} = s_{out} = 10$. We stop the algorithm when $\|\mathcal{P}(\boldsymbol{L}\boldsymbol{R}) - \boldsymbol{F}\|_F / \|\boldsymbol{F}\|_F \le \epsilon$ and ϵ is chosen case by case.

3 Experimental Results

3.1 Full-fan scan

For this case, the plot of $\|\boldsymbol{L}(:,i)\boldsymbol{R}(i,:)\|_\infty$ in a descending order as a function of the column index i is first depicted in Figure 2, with a trial run of $K = 20$. Clearly, the 8th column and beyond contribute little to the reconstructed images and hence should be removed. As such, $K = 7$ is selected and the corresponding results are shown in Figure 3 and Figure 4. Comparing the restored images in Figure 3 with the corresponding ground truth images in Figure 1, it is found that our algorithm is able to capture the motions of the anatomy

Figure 2: $\|\boldsymbol{L}(:,i)\boldsymbol{R}(i,:)\|_\infty$ in descending order. The result is produced with $K = 20$.

and restore the structures, even those small ones inside the lung. Meanwhile, observable artifacts inside the heart and at its boundary also exist. Quantitatively, the relative error of the restored cine-CBCT images is 3.89%.

To further look into the reconstruction results, we plot the columns of $\boldsymbol{L}$ in Figure 4(a), where each column is reshaped into a 128×128 image. These images form the basis to represent all the reconstructed cine-CBCT images. It is observed that the first one is similar to an image averaged over all the cine-CBCT images. Its presence provides an overall structure that is common to all the images in the cine-CBCT. Meanwhile, other basis images represent differences between images of $\boldsymbol{U}$. We also plot the corresponding coefficients in $\boldsymbol{R}$ in Figure 4(b), which attain a periodically variation pattern, indicating the patient respiratory motion.

3.2 Noisy data from a full-fan scan

For the case with noise in the projection data from a full-fan scan, the parameter K is selected as $K = 4$ according to Figure 5. Figure 6(a) depicts the reconstructed images. The relative error of the restored cine-CBCT images is 6.81%. The columns of $\boldsymbol{L}$ and the rows of $\boldsymbol{R}$ are plotted in Figure 6(b) and (c), respectively. Again, we see that the columns of $\boldsymbol{L}$ are meaningful basis corresponding to the average image as well as the variations between them, and the rows of $\boldsymbol{R}$ oscillate periodically.

3.3 Half-fan scan

For the case with noiseless data from a half-fan scan, the parameter K is selected as $K = 7$ according to Figure 7. In Figure 8, we plot the reconstructed images, the columns of $\boldsymbol{L}$, and the rows of $\boldsymbol{R}$. The relative error of the restored cine-CBCT images is 6.62%, and we see again that the columns of $\boldsymbol{L}$ are meaningful basis corresponding to the average image as well as the variations between them, and the rows of $\boldsymbol{R}$ oscillate periodically.

(a) The restored cine-CBCT images. The display window is $(-1000, 420)$ HU.

(b) The absolute value of the differences to the ground truth. The display window is $(0, 280)$ HU.

Figure 3: The restored cine-CBCT images from a full-fan scan and the absolute value of the differences to the ground truth. The relative error is 3.89%. From left to right: frame $20, 81, 160, 256, 334$.

(a) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.

(b) The rows of $\boldsymbol{R}$.

Figure 4: $\boldsymbol{L}$ and $\boldsymbol{R}$ in the cine-CBCT image reconstruction from a full-fan scan.

Figure 5: $\|\boldsymbol{L}(:,i)\boldsymbol{R}(i,:)\|_\infty$ in descending order. The result is produced with $K = 20$ and noise at 0.5 mAs/projection.

(a) The restored cine-CBCT image with the relative error 6.81%. From left to right: frame $20, 81, 160, 256, 334$. The display window is $(-1000, 420)$ HU.


(b) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.

(c) The rows of $\boldsymbol{R}$.

Figure 6: Result of cine-CBCT image reconstruction with noise at 0.5 mAs/projection from a full-fan scan.

Figure 7: $\|\boldsymbol{L}(:,i)\boldsymbol{R}(i,:)\|_\infty$ in descending order. The result is produced with $K = 20$ from a half-fan scan.

4 Validations on our model

4.1 Validation on the low-rank assumption

To demonstrate validity of the low rank assumption, we have conducted a study on a digital NCAT phantom, where typical images of such a phantom are in Fig. 1. A sequence of T images are generated at a transverse slice during a breathing cycle. Each of these images is represented by a vector of length N, where N is the number of pixels in one image. We then form a $N \times T$ matrix using these vectors as columns. Singular value decomposition of this matrix is conducted. For a case with $T = 100$ per breathing cycle, there are in total 100 non-zero singular values (SVs), indicating that the matrix is strictly speaking of full rank. Yet, only the first few SVs dominate, as illustrated in the left panel of Fig. 9. This implies that the matrix is approximately of low rank. In particular, the 8 largest SVs accounts for over 97% of the sum of all the SVs. Moreover, this ratio, namely the required number of leading SVs to account over 97% of the sum of all SVs to the total number of SVs (or images), decreases dramatically as the number of images T sampled in a breathing cycle increases, as shown in the right panel of Fig. 9. This simple numerical experiment demonstrates the validity of our assumption to a certain extent and the rest of our reconstruction work is based on this assumption. However, we remark that whether this assumption holds in real clinical cases is subject to further validations.

4.2 Exactly Low-rank Ground Truth

We have also conducted one experiment where the ground truth 4DCT images are indeed forms a low-rank matrix. This experiments show that our model works better for exactly low-rank ground truths than with approximately low-rank ones.

We generate the true images $\hat{\boldsymbol{U}}_1$ by $\hat{\boldsymbol{U}}_1 = \hat{\boldsymbol{L}}_1\hat{\boldsymbol{R}}_1$, where $\hat{\boldsymbol{L}}_1$ and $\hat{\boldsymbol{R}}_1$ are from Figure 4 (a) and (b) respectively. By this way, the true cine-CBCT image $\hat{\boldsymbol{U}}_1$ is an approximation to the cine-CBCT image in Section 2.5 whose respiratory motion is rather regular. Clearly, the rank of $\hat{\boldsymbol{U}}_1$ is 7 and the dimension is

(a) The restored cine-CBCT image with the relative error 6.62%. From left to right: frame $20, 81, 160, 256, 334$. TThe display window is $(-1000, 420)$ HU.

(b) The absolute value of the differences to the ground truth. The display window is $(0, 710)$ HU.

(c) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.

(d) The rows of $\boldsymbol{R}$.

Figure 8: Result of cine-CBCT image reconstruction from a half-fan scan.

Figure 9: Left: Singular values of a set of 100 NCAT images taken with in one breathing cycle. Right: the ratio of required number of singular values to restore 97% of the NCAT images to the total number of images n_f in a breathing cycle.

16384×360, as $\hat{L}_1$ consists of 7 basis images of size 128×128 and $\hat{R}_1$ of 7 coefficients of length 360. Then, we restore L and R from $F = \mathcal{P}\hat{U}_1$ by Algorithm 1. The restored basis images and the ground truth images are compared in Figure 10. The relative error of the restored cine-CBCT image is 1.35%, which is smaller than that in Section 3.1, where the ground truth is not exactly but approximately low-rank. This shows the low-rankness of the ground truth 4DCT images can help reduce the reconstruction error.

4.3   One cycle scan

Our model does not depend on the periodicity of the respiratory motion. To validate this fact, we tested our model on cine-CBCT images generated in one breathe cycle, where no breathe phase is repeated. We use the true cine-CBCT image from the one described in Section 2.5 with a scanning period 4.1 seconds and and still $T = 360$, and frames of this true cine-CBCT image are shown in Figure 11. Therefore, there is only one full cycle of respiratory motion in the true cine-CBCT images. The results are plotted in Figure 12 and Figure 13. Our method can still give a satisfactory result, and the relative error of the restored cine-CBCT images is 9.78%.

For the one cycle scan, we performed two more experiments to see how the projection operator $\mathcal{P}$ affects the quality of the reconstructed images. Firstly, we find that a random permutation on the projection angles could improve the quality of the reconstructed image. This can be seen from the results presented in Figure 14, and the relative error of the restored cine-CBCT images is 5.39%. The reason is that, when we randomly permute the projection angles, the projection angles for similar images are random, making the reconstruction more accurate and robust especially when using one cycle scan. Secondly, increasing the temporal resolution could improve the quality of reconstructed images. To see this, we use the true cine-CBCT image from the one described in the previous paragraph with a scanning period 4.1 seconds and $T = 720$. The results are depicted in Figure 15. The relative error of the restored cine-CBCT images is 7.42%. This is related to the validation test presented in Section 4.1, i.e., the ration of required number of leading SVs to account over a certain percentage of the sum of all SVs to the total number of SVs (or images), decreases dramatically as the number of images T sampled in a breathing cycle increases, as shown in the right panel of Fig. 9.

13

(a) The restored cine-CBCT image with the relative error 1.35%. From left to right: frame $20, 81, 160, 256, 334$. The display window is $(-1000, 420)$ HU.

(b) The absolute value of the differences to the ground truth. The display window is $(0, 280)$ HU.

(c) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.

(d) The rows of $\boldsymbol{R}$.

Figure 10: Result of cine-CBCT image (with exactly rank 7) reconstruction from a full-fan scan.

Figure 11: The ground truth cine-CBCT images generating from one breathe cycle. From left to right: frame $50, 120, 190, 260, 330$. The display window is $(-1000, 420)$ HU.

Figure 12: $\|\boldsymbol{L}(:,i)\boldsymbol{R}(i,:)\|_\infty$ in descending order in Experiment 5. The result is produced with $K = 20$.

4.4   High resolution images

Since the image video is stored and manipulated in a low-rank decomposition form, our method can reconstruct images with high resolution images. To demonstrate this, we test our algorithm on a high resolution correspondence of the datasets in Section 3.1. In particular, we test images with 512×512 pixels and 360 projections, and we still use a detector resolution of 256 bins. The results are depicted in Figure 16. The relative error of the restored cine-CBCT images is 10.11%. The error is larger than that for low-resolution since the linear equation (5) is more under-determined than the low resolution case.

5   Conclusion and Discussions

In this work, we have proposed an computational algorithm to reconstruct cine-CBCT to provide a set of time-dependent images in the thorax region under the assumption that the matrix formed by all of these images to be reconstructed has a low rank approximation. In contrast to the currently used 4DCBCT, where only respiratory phase-resolved imaging is achieved, the cine-CBCT modality retrieves the instantaneous patient anatomy corresponding to each CBCT projection. Cine-CBCT reconstruction is apparently a challenging problem due to the very limited projection information. Yet, by effectively incorporating the underlying local and global temporal coherence of cine-CBCT images into the reconstruction process, we demonstrate the feasibility of this modality. In particular, the low rank approximation is based on matrix factorization form $\boldsymbol{U} = \boldsymbol{LR}$ for the matrix $\boldsymbol{U}$ whose columns are the images to be reconstructed. The dimension of the matrix $\boldsymbol{L}$ constraints the rank of $\boldsymbol{U}$ and hence maximally and automatically exploits a coherence condition among all the images in cine-CBCT. Simulation studies on an NCAT phantom serve as a very preliminary test to demonstrate the feasibility of our approach, where successive reconstructions have been observed in various experiments.

In addition to the apparent advantage of improved temporal resolution in the cine-CBCT over 4DCBCT, the reconstruction of cine-CBCT is not as demanding regarding the amount of projection data as 4DCBCT, at least based on our simulation studies. Because of the phase binning in 4DCBCT, a certain number of

(a) The restored cine-CBCT image with the relative error 9.78%. From left to right: frame $50, 120, 190, 260, 330$. The display window is $(-1000, 420)$ HU.

(b) The absolute value of the differences to the ground truth. The display window is $(0, 1420)$ HU.

(c) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.

(d) The rows of $\boldsymbol{R}$.

Figure 13: Result of cine-CBCT image (one breathe cycle) reconstruction from a full-fan scan.

(a) The restored cine-CBCT image with the relative error 5.39%. From left to right: frame 50, 120, 190, 260, 330. The display window is $(-1000, 420)$ HU.

(b) The absolute value of the differences to the ground truth. The display window is $(0, 710)$ HU.

(c) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.


(d) The rows of $\boldsymbol{R}$.

Figure 14: Result of cine-CBCT image (one breathe cycle) reconstruction from a full-fan scan with randomly permuted projection angles.

17

(a) The restored cine-CBCT image with the relative error 7.42%. From left to right: frame $100, 240, 380, 520, 660$. The display window is $(-1000, 420)$ HU.

(b) The absolute value of the differences to the ground truth. The display window is $(0, 1420)$ HU.

(c) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.


(d) The rows of $\boldsymbol{R}$.

Figure 15: Result of cine-CBCT image (one breathe cycle) reconstruction from a full-fan scan with $T = 720$.

18

(a) The restored cine-CBCT image with the relative error 10.11%. From left to right: frame $20, 81, 160, 256, 334$. The display window is $(-1000, 420)$ HU.

(b) The columns of $\boldsymbol{L}$. Each column is reshaped into an image.

(c) The rows of $\boldsymbol{R}$.

Figure 16: Result of 512×512 cine-CBCT image reconstruction from a full-fan scan.

projections are required in each bin to yield a CBCT image for the bin with a satisfactory quality. Hence, it is usually challenging to get a 4DCBCT with only one-minute scan and in practice, protocols of multiple rotations or a slow rotation is utilized. In contrast, the cine-CBCT effectively incorporates the inter-image correlations into the reconstruction and a set of images with visible major anatomical features are obtained even with a super-fast scan protocol corresponding to only a single breathing period, i.e., a 4 second scan if the patient's breathing period is 4 seconds, as demonstrated in our studies. Another advantage is that our method can deal with respiratory motion quite well in the sense that irregularity does not affect the construction of the regular images too much as shown in the experiments.

Although the 4th dimension in cine-CBCT is more than breathing phase as in conventional 4DCBCT, we would like to point out that cine-CBCT is not completely time resolved in general. A good cine-CBCT reconstruction is based on two key assumptions. First, all cine-CBCT images can be represented (approximately) by linear combinations of a small number of basis images, meaning certain irregular patient anatomy variation may not be captured by cine-CBCT if they are not in the linear space spanned by the extracted basis images. Second, all cine-CBCT images collectively correspond to a sufficient number of projections covering a sufficiently large scanning angle, meaning our algorithm will not work if the scanning angle is smaller than that required by the regular CBCT reconstruction. That being said, cine-CBCT images reconstructed using our algorithm still contain clinically valuable temporal information and represent a quantum leap from the conventional 4DCBCT images.

Despite the success in this feasibility study, there are a few practical issues as for the clinical applications of cine-CBCT. First of all, the low rank assumption has to be verified using clinical data. Second, only one slice of the patient anatomy is reconstructed in this proof-of-principle study. When cone-beam geometry is considered to reconstruct volumetric CBCT images, although the fundamental principles in our algorithm can be easily applied, there potentially come some problems due to the extremely large data size involved. The computational efficiency needs to be improved.

In general, true or semi time resolved CBCT reconstruction poses a challenge to both storage and

computational time. Currently, it takes a few hours to reconstruct those images shown in our study with 128×128 pixels and 360 projections on a typical computer using Matlab. For high-resolution reconstruction, cone beam geometry, and clinical applications, more powerful computational platforms, such as graphics processing unit (GPU) [15], or multi-GPU [32], can be used. GPU is essentially a parallel processing platform with a large number of processing cores. For example, in Algorithm 1, the most computation intensive part is computing forward projection $\mathcal{P}$ and its transpose $\mathcal{P}^T$ when solving the least-square problems. For problems of large sizes, these matrices are too large to be stored. But these operations can be greatly accelerated on the GPU platform by splitting the operations according to projection ray-lines, or image voxels [13]. Other computation intensive tasks, such as matrix-matrix or matrix-vector operations, can be easily parallelized to achieve significant speed up. Another challenge is large memory requirement which can not fit into single GPU. In this case, splitting calculations across multiple GPUs will be the solution, at the cost of inter-GPU communications [32]. It will be our future research project to implement the reconstruction algorithm on GPU.

Acknowledgement

References

[1] F. Bergner, T. Berkus, M. Oelhafen, P. Kunz, T. Pan, R. Grimmer, L. Ritschl, and M. Kachelriess, *An investigation of 4D cone-beam CT algorithms for slowly rotating scanners*, MEDICAL PHYSICS, 37 (2010), pp. 5044–5053.

[2] F. Bergner, T. Berkus, M. Oelhafen, P. Kunz, T. Pan, and M. Kachelriess, *Autoadaptive phase-correlated (AAPC) reconstruction for 4D CBCT*, MEDICAL PHYSICS, 36 (2009), pp. 5695–5706.

[3] S. Burer and R. D. C. Monteiro, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, Math. Program., 95 (2003), pp. 329–357. Computational semidefinite and second order cone programming: the state of the art.

[4] S. Burer and R. D. C. Monteiro, *Local minima and convergence in low-rank semidefinite programming*, Math. Program., 103 (2005), pp. 427–444.

[5] J.-F. Cai, E. J. Candès, and Z. Shen, *A singular value thresholding algorithm for matrix completion*, SIAM J. Optimiz., 20 (2010), pp. 1956–1982.

[6] J.-F. Cai, B. Dong, S. Osher, and Z. Shen, *Image restoration: Total variation, wavelet frames, and beyond*, J. Amer. Math. Soc., 25 (2012), pp. 1033–1089.

[7] J.-F. Cai, S. Osher, and Z. Shen, *Linearized Bregman iterations for frame-based image deblurring*, SIAM J. Imaging Sci., 2 (2009), pp. 226–252.

[8] J.-F. Cai, S. Osher, and Z. Shen, *Split Bregman methods and frame based image restoration*, Multiscale Modeling & Simulation, 8 (2009), pp. 337–369.

[9] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky, *The convex geometry of linear inverse problems*, Arxiv preprint arXiv:1012.0621, (2010).

[10] E. Ford, G. Mageras, E. Yorke, and C. Ling, *Respiration-correlated spiral CT: A method of measuring respiratory-induced anatomic motion for radiation treatment planning*, Medical Physics, 30 (2003), pp. 88–97.

[11] H. Gao, J.-F. Cai, Z. Shen, and H. Zhao, *Robust principle component analysis based four-dimensional computed tomography*, Physics in Medicine and Biology, 56 (2011), pp. 3181–3198.

[12] T. Goldstein and S. Osher, *The split Bregman method for L1 regularized problems*, SIAM Journal on Imaging Sciences, 2 (2009), pp. 323–343.

[13] X. Jia, B. Dong, Y. Lou, and S. Jiang, *GPU-based iterative cone-beam CT reconstruction using tight frame regularization*, Physics in Medicine and Biology, 56 (2011), pp. 3787–3807.

[14] X. Jia, Y. Lou, B. Dong, Z. Tian, and S. Jiang, *4D computed tomography reconstruction from few-projection data via temporal non-local regularization*, Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010, (2010), pp. 143–150.

[15] X. Jia, Y. Lou, R. Li, W. Y. Song, and S. B. Jiang, *Gpu-based fast cone beam ct reconstruction from undersampled and noisy projection data via total variation*, Medical Physics, 37 (2010), pp. 1757–1760.

[16] X. Jia, Z. Tian, Y. Lou, J. Sonke, and S. Jiang, *Four-dimensional cone beam ct reconstruction and enhancement using a temporal non-local means method*, submitted to Medical Physics, (2012).

[17] S. Kriminski, M. Mitschke, S. Sorensen, and et al., *Respiratory correlated cone-beam computed tomography on an isocentric c-arm*, Physics in Medicine and Biology, 50 (2005), pp. 5263–5280.

[18] R. Li, X. Jia, J. H. Lewis, X. Gu, M. Folkerts, C. Men, and S. B. Jiang, *Real-time volumetric image reconstruction and 3D tumor localization based on a single x-ray projection image for lung cancer radiotherapy*, Medical Physics, 37 (2010), pp. 2822–2826.

[19] R. Li, J. H. Lewis, X. Jia, X. Gu, M. Folkerts, C. Men, W. Y. Song, and S. B. Jiang, *3D tumor localization through real-time volumetric x-ray imaging for lung cancer radiotherapy*, Medical Physics, 38 (2011), pp. 2783–2794.

[20] T. Li and L. Xing, *Optimizing 4D cone-beam CT acquisition protocol for external beam radiotherapy*, International Journal of Radiation Oncology Biology Physics, 67 (2007), pp. 1211–1219.

[21] T. F. Li, L. Xing, P. Munro, and et al., *Four-dimensional cone-beam computed tomography using an on-board imager*, Medical Physics, 33 (2006), pp. 3825–3833.

[22] J. Lu, T. M. Guerrero, P. Munro, A. Jeung, P.-C. M. Chi, P. Balter, X. R. Zhu, R. Mohan, and T. Pan, *Four-dimensional cone beam CT with adaptive gantry rotation and adaptive data sampling*, Medical Physics, 34 (2007), pp. 3520–3529.

[23] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, *An iterative regularization method for total variation-based image restoration*, Multiscale Model. Simul., 4 (2005), pp. 460–489 (electronic).

[24] Z. Qi and G.-H. Chen, *Performance studies of four-dimensional cone beam computed tomography*, Physics in Medicine and Biology, 56 (2011), pp. 6709–6721.

[25] Qinghui Zhang, Yu-Chi Hu, Fenghong Liu, K. Goodman, K. Rosenzweig, and G. Mageras, *Correction of motion artifacts in cone-beam CT using a patient-specific respiratory motion model*, Medical Physics, 37 (2010), pp. 2901–9.

[26] S. Rit, J. W. H. Wolthaus, M. van Herk, and J.-J. Sonke, *On-the-fly motion-compensated cone-beam CT using an a priori model of the respiratory motion*, MEDICAL PHYSICS, 36 (2009), pp. 2283–2296.

[27] A. RON AND Z. SHEN, *Affine systems in $L_2(\mathbf{R}^d)$: the analysis of the analysis operator*, J. Funct. Anal., 148 (1997), pp. 408–447.

[28] W. SEGARS, M. MAHESH, T. BECK, E. FREY, AND B. TSUI, *Realistic ct simulation using the 4d xcat phantom*, Medical Physics, 35 (2008), p. 3800.

[29] SHUAI LENG, JIE TANG, J. ZAMBELLI, B. NETT, R. TOLAKANAHALLI, AND GUANG-HONG CHEN, *High temporal resolution and streak-free four-dimensional cone-beam computed tomography*, Physics in Medicine and Biology, 53 (2008), pp. 5653–73.

[30] J. J. SONKE, L. ZIJP, P. REMEIJER, AND ET AL., *Respiratory correlated cone beam ct*, Medical Physics, 32 (2005), pp. 1176–1186.

[31] Z. TIAN, X. JIA, B. DONG, Y. LOU, AND S. JIANG, *Low-dose 4dct reconstruction via temporal nonlocal means*, Medical Physics, 38 (2011), pp. 1359–1365.

[32] H. YAN, X. WANG, L. CERVINO, S. B. JIANG, AND X. JIA, *Towards the clinical implementation of iterative cone beam ct reconstruction for radiation therapy using a multi-gpu system*, Submitted to Med. Phys., (2014).