



HHS Public Access

Author manuscript

IEEE Trans Med Imaging. Author manuscript; available in PMC 2015 July 15.

Published in final edited form as:

IEEE Trans Med Imaging. 2012 October ; 31(10): 1837–1848. doi:10.1109/TMI.2012.2199763.

Model-Based Tomographic Reconstruction of Objects Containing Known Components

J. Webster Stayman,

Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21205 USA

Yoshito Otake,

Departments of Biomedical Engineering and Computer Computer Science, Johns Hopkins University, Baltimore, MD 21205 USA

Jerry L. Prince,

Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218

A. Jay Khanna, and

Department of Orthopaedic Surgery, Johns Hopkins University, Baltimore, MD 21205 USA

Jeffrey H. Siewerdsen

Departments of Biomedical Engineering and Computer Science, Johns Hopkins University, Baltimore, MD 21205 USA

Abstract

The likelihood of finding manufactured components (surgical tools, implants, etc.) within a tomographic field-of-view has been steadily increasing. One reason is the aging population and proliferation of prosthetic devices, such that more people undergoing diagnostic imaging have existing implants, particularly hip and knee implants. Another reason is that use of intraoperative imaging (e.g., cone-beam CT) for surgical guidance is increasing, wherein surgical tools and devices such as screws and plates are placed within or near to the target anatomy. When these components contain metal, the reconstructed volumes are likely to contain severe artifacts that adversely affect the image quality in tissues both near and far from the component. Because physical models of such components exist, there is a unique opportunity to integrate this knowledge into the reconstruction algorithm to reduce these artifacts. We present a model-based penalized-likelihood estimation approach that explicitly incorporates known information about component geometry and composition. The approach uses an alternating maximization method that jointly estimates the anatomy and the position and pose of each of the known components. We demonstrate that the proposed method can produce nearly artifact-free images even near the boundary of a metal implant in simulated vertebral pedicle screw reconstructions and even under conditions of substantial photon starvation. The simultaneous estimation of device pose also provides quantitative information on device placement that could be valuable to quality assurance and verification of treatment delivery.

Index Terms

CT Reconstruction; Implant Imaging; Metal Artifact Reduction; Penalized-Likelihood Estimation

I. Introduction

In tomographic imaging, there are many situations in which portions of the image volume are known *a priori*. Examples in orthopedics include pedicle screws and rods for spine surgery, knee and hip implants for joint replacements, and plates and screws for fixation in trauma cases. In image-guided procedures, surgical tools are often placed within the imaging field. When these components are metallic, measurements whose projections contain those elements can suffer from reduced signal-to-noise ratio due to photon starvation. Similarly, since reconstruction of highly attenuating components involves mathematical inversion of near zero projection values, algorithms tend to be very sensitive to any biases (e.g., due to polyenergetic effects). Both of these effects tend to produce streak artifacts in the reconstructed images [1, 2]. Such artifacts tend to be particularly troublesome since it is often the region immediately surrounding a component that is of diagnostic interest, which is exactly where artifacts tend to be most pronounced. Particular situations where image quality in the neighborhood of a metallic component is critical include the visualization around implants for indications of subsidence or osteolysis [3], assessment of pedicle screw placement to avoid critical structures in the spine [4, 5], and biopsy needle guidance [6].

Various approaches have been developed to mitigate metal streak artifacts [7–15]. Many methods consider measurements through metal to be missing data. The missing data can simply be eliminated from the reconstruction algorithm [14], or may be filled in using values based on the neighborhood of the missing data [8, 9]. However, rarely is the exact knowledge of the metal component used.

Tomographic imaging generally benefits from the incorporation of prior knowledge into the reconstruction algorithm. This is particularly true for situations that involve undersampling and low signal-to-noise. Methods that seek to correct for metal streak artifacts tend to require identification of spatial locations in the volume, or the locations in the projection image where the metal implant lies. This localization typically relies on knowledge that the metal components have a high attenuation coefficient. In effect, this is a relatively weak incorporation of prior knowledge. In penalized-likelihood reconstruction schemes, general knowledge about the image can be included via Gibbs priors or penalty functions [16–19]. In more recent work, very specific image priors that incorporate prior scans of the anatomy have been used in algorithms like PICCS [20] and modified penalized-likelihood approaches [21].

The availability of physical models for surgical tools, fixation hardware, and implants allows for very specific prior knowledge to be incorporated into the reconstruction routine with the potential for additional benefit. Since such components represent manufactured objects, CAD models that completely specify their material composition and structure may be available. In this paper we propose an algorithm that integrates such known physical models into the reconstruction process. Specifically, the model of the object volume itself is

a combination of: 1.) the volume of (unknown) background anatomy; and 2.) the component (or components) known to be in the imaging field-of-view. While the form and attenuation distributions of the components are known (e.g., derived from a CAD model that specifies the shape and material content of the device), the positions and poses are unknown. Thus, the parameters associated with each component registration also become part of the object model. The resulting reconstruction scheme has an objective that is a function of both image parameters and registration parameters, and these two sets of parameters are estimated jointly. This approach bears some similarity with other objective functions that seek joint image reconstruction and registration [22–24]. A preliminary introduction of this reconstruction approach was presented in [25]. The work below provides a comprehensive discussion of the methodology including a derivation of the reconstruction algorithm, a study of the convergence properties, and an investigation of the influence of regularization.

The proposed approach has similarities with work by Snyder, Murphy, *et al.* who also developed a model-based approach for incorporating exact knowledge of a component through a constraint on the objective function [26–28]. The approach outlined in this paper is distinct, utilizes an unconstrained objective function, adopts a regularization term, and is generalized for an arbitrary number of known components whose poses are unknown. A more detailed discussion of the similarities and differences between these two strategies for incorporating prior component knowledge is provided in Section IV.

This paper is outlined as follows. In Section II, we develop a likelihood-based objective function and reconstruction algorithm that models the image volume as the combination of an arbitrary number of known components (with unknown poses) within the unknown background anatomy. In Section III, we illustrate the performance of this method and compare the performance with traditional analytic and iterative approaches, including an investigation of convergence properties and the influence of regularization.

II. Methods

A. Forward Model

Consider the following measurement model for a transmission tomography system

$$\bar{y}_i = b_i \exp(-l_i). \quad (1)$$

Mean projection measurements, \bar{y}_i , are related to line integrals through the object, l_i , via Beer's law. Each measurement has a scalar value, b_i , associated with the unattenuated x-ray fluence for the particular ray path and gain for the specific detector element. For a discretized object (e.g., using a voxel basis), the line integrals can be represented as

$$l_i = \sum_{j=1}^p a_{ij} \mu_j, \quad (2)$$

where a_{ij} represents the contribution of the j^{th} voxel (or other basis), μ_j , to the i^{th} line integral. This model may be written compactly in a matrix-vector form such that

$$\bar{y} - \mathbf{D}\{b\} \exp(-\mathbf{A}\mu) \quad (3)$$

where $\mathbf{D}\{\cdot\}$ represents an operator that forms a diagonal matrix from the vector argument, and the system matrix, \mathbf{A} , represents the collection of all a_{ij} . Ordinarily, (3) represents the complete relationship between the object and the measurements, and it is from (3) that an objective function is derived. We choose to further parameterize the object as the combination of an unknown anatomical background volume, μ_* , and an arbitrary number of components, $\mu_I^{(n)}$, whose attenuation distributions are known. Mathematically, we write

$$\mu = \left(\prod_{n=1}^N \mathbf{D}\{\mathbf{W}(\lambda^{(n)}) s^{(n)}\} \right) \mu_* + \sum_{n=1}^N \mathbf{W}(\lambda^{(n)}) \mu_I^{(n)}, \quad (4)$$

where $\mathbf{W}(\lambda)$ represents a transformation operator parameterized by the vector. While this transformation is general, in this work we focus on three-dimensional (3D) rigid transformations. In this case, the vector λ is comprised of the six elements (translation and rotation values) that define a component's pose. The second term in (4) represents the summation of attenuation values for N known attenuation volumes, $\mu_I^{(n)}$, each having been (rigidly) transformed according to its own parameter vector, $\lambda^{(n)}$. The first term of (4) represents the contribution of the background anatomy, μ_* , which has been modified by multiplication with a set of transformed masks, $s^{(n)}$, corresponding to each known component. Specifically, each mask represents a support region for each component. A typical mask is largely binary with zero values in the interior and ones outside the component but allows for intermediate values to model partial volume effects at the edges. The product operator, \prod , used in (4) and throughout this paper represents an element-by-element product for matrix or vector operands.

B. Likelihood-Based Objective Function

Equations (3) and (4) represent the relationship between the mean measurements and the parameterized volume. Selection of a noise model allows for derivation of a likelihood function. We choose to invoke a Poisson model which yields the following log-likelihood

$$\begin{aligned} L(\mu_*, \Lambda; y) &= \sum_i [y]_i \log ([\bar{y}(\mu_*, \Lambda)]_i - [y]_i) \\ &= \sum_i y_i \log [b_i \exp(-l_i)] - b_i \exp(-l_i) \end{aligned} \quad (5)$$

where we have denoted the collection of the set of registration parameters for all N components as

$$\Lambda = \left\{ \lambda^{(n)} \right\}_{n=1}^N. \quad (6)$$

We choose to employ a penalized-likelihood estimator with the following implicit form

$$\begin{aligned} \{\hat{\mu}_*, \hat{\Lambda}\} &= \arg \max_{\mu_*, \Lambda} \Phi(\mu_*, \Lambda; y) \\ &= \arg \max_{\mu_*, \Lambda} L(\mu_*, \Lambda; y) - \beta R(\mu_*) \end{aligned} \quad (7)$$

where $R(\cdot)$ denotes a penalty function that discourages overly noisy images. We refer to the implicit estimator in (7) as the “known-component reconstruction” (KCR) approach. This estimator jointly approximates both the anatomical background and the set of registration parameters associated with each known component. While the regularization term is general, we will use a pairwise quadratic roughness penalty denoted as

$$R(\mu_*) = \sum_j \sum_{k \in \mathcal{N}_j} \left([\mu_*]_j - [\mu_*]_k \right)^2 \quad (8)$$

where \mathcal{N}_j denotes a first-order neighborhood of voxels around voxel j . Note that this penalty is applied only to the background anatomy image (μ_*), and there is no smoothing between the known components and the anatomy, but there is potential smoothing across the voxels in μ_* , that are masked by the transformed $s^{(n)}$ in (4). These values are ultimately replaced by $\mu_I^{(n)}$ to form the reconstructed image μ .

C. Transformation Operator

The joint estimator in (7) requires the transformation operator, \mathbf{W} , found in (4) to be specified. The formulation in (4) is general and the transformation could potentially represent warping operations or other forms of parameterized deformations. Similarly, \mathbf{W} could represent component transformations that have constrained motions like pivots or hinge points. In this paper we focus on the case of rigid transformations.

Transformation operators associated with registration of a “moving” image to a “fixed” image can often be decomposed into two parts: 1) a mapping of gridded points in the moving image to the positions of those points in the transformed image; and 2) an interpolation operation that uses a neighborhood of voxels about each point in the moving image to generate a voxel value in the transformed image. For a 3D rigid transformation, we may denote the mapping of points from the moving image to the transformed image as

$$\begin{bmatrix} x'_j \\ y'_j \\ z'_j \end{bmatrix} = R_x(\theta) R_y(\psi) R_z(\phi) \left(\begin{bmatrix} x_j \\ y_j \\ z_j \end{bmatrix} - \begin{bmatrix} x_c \\ y_c \\ z_c \end{bmatrix} \right) + \begin{bmatrix} x_c - x_T \\ y_c - y_T \\ z_c - z_T \end{bmatrix}, \quad (9)$$

where $[x_j, y_j, z_j]^T$ represents a voxel location in the moving image, and $[x'_j, y'_j, z'_j]^T$ is the position in the transformed image. The vector $[x_c, y_c, z_c]^T$ is the location of the center of the image volume and $[x_T, y_T, z_T]^T$ are translation parameters. The parameterized matrices $R_x(\theta)$, $R_y(\psi)$, and $R_z(\phi)$ apply rotations about each axis with (θ) , ψ , and ϕ denoting angular values for each rotation. Thus, we may now explicitly define the six element parameter vector for transformation as $\lambda = [x_T, y_T, z_T, \theta, \psi, \phi]^T$. While (9) denotes the mapping of the j^{th} voxel position from $[x_j, y_j, z_j]^T$ to $[x'_j, y'_j, z'_j]^T$, we recognize that the *entire* mapping is defined by the three “moving image” vectors x , y , and z , and the three “transformed image” vectors x' , y' ,

and z' . That is, these vectors represent the collections of all grid points in the moving and transformed images.

With the point mapping defined, it remains to define the interpolation step. We consider a kernel-based interpolation approach that is separable in each dimension. An example 2D kernel-based interpolation is illustrated in Figure 1. For each location, j , in the transformed image, a kernel is applied to a neighborhood of points in the moving image to approximate the value in the transformed image. One-dimensional (1D) (position-dependent) kernels, $k_j(t)$, are applied successively along each dimension in the neighborhood and are summed, so that each pass reduces the dimension by one – for example, in 2D, two 1D kernels applied to get a single voxel value. In 3D, three 1D kernels are used. The application of the kernels illustrated in Figure 1 to obtain all interpolated points in the transformed image may be represented mathematically in a compact form as

$$u = \mathbf{W}(\lambda)v = \mathbf{K}_z(z')\mathbf{K}_y(y')\mathbf{K}_x(x')v, \quad (10)$$

where u and v are vectors denoting the transformed and moving images, respectively. The \mathbf{K} operators represent the 1D kernel applications in each direction for *all* locations and are a function of the particular point mapping. Specifically, if a fourpixel neighborhood 1D kernel is used for each dimension and both u and v have p voxels, the \mathbf{K}_x operator applies a local kernel in the x -direction for each interpolation position to obtain vector of length $16p$. The \mathbf{K}_y operator applies a y -direction kernel to these $16p$ elements to obtain a vector of length $4p$, and the final \mathbf{K}_z operator uses these $4p$ elements and a z -direction kernel to obtain the length p vector that is the transformed image, u .

It remains to define the particular kernels, $k(t)$, used at each interpolation location. While there are many kernel choices, we employed a B-spline kernel [29] defined as follows:

$$k(t) = \begin{cases} \frac{2}{3} - \frac{1}{2}|t|^2(2 - |t|) & |t| < 1 \\ \frac{1}{6}(2 - |t|)^3 & 1 \leq |t| < 2 \\ 0 & \text{Otherwise} \end{cases} \quad (11)$$

This kernel has two particularly advantageous properties: 1) it preserves nonnegativity in the transformed image (for non-negative “moving” images) and thereby avoids negative attenuation coefficients from entering into the estimation; and 2) $k(t)$ is differentiable and therefore allows for gradient-based methods in the optimization of the objective proposed in (7).

In the following section we utilize analytic gradient information to approximate solutions to (7). Since that gradient information necessarily includes derivatives of the transformation operator with respect to the parameters being estimated, we derive those equations here. Specifically, we require the derivative of the \mathbf{W} operator with respect to elements of λ . Using the chain rule one may write

$$\begin{aligned}
& \dot{\mathbf{W}}_m(\lambda)v \triangleq \frac{\partial}{\partial \lambda_m} \mathbf{W}(\lambda)v \\
& = \mathbf{D} \left\{ \frac{\partial}{\partial \lambda_m} z' \right\} \dot{\mathbf{K}}_z(z') \mathbf{K}_y(y') \mathbf{K}_x(x') v \\
& + \mathbf{D} \left\{ \frac{\partial}{\partial \lambda_m} y' \right\} \mathbf{K}_z(z') \dot{\mathbf{K}}_y(y') \mathbf{K}_x(x') v \\
& + \mathbf{D} \left\{ \frac{\partial}{\partial \lambda_m} x' \right\} \mathbf{K}_z(z') \mathbf{K}_y(y') \dot{\mathbf{K}}_x(x') v
\end{aligned} \quad (12)$$

where we denote the derivative transformation operator with respect to the m^{th} element of λ as $\dot{\mathbf{W}}_m(\lambda)v$. The $\dot{\mathbf{K}}$ operators are similar to the original kernel operators except that they apply the derivative kernel. For the B-spline kernel in (11), we may compute the derivative kernel as

$$\dot{k}(t) = \begin{cases} \frac{1}{2} \text{sgn}(t) |t|^2 - t(2 - |t|) & |t| < 1 \\ \frac{1}{2} \text{sgn}(t) (2 - |t|)^2 & 1 \leq |t| < 2 \\ 0 & \text{Otherwise} \end{cases} \quad (13)$$

The remaining partial derivatives in (12) may be computed by differentiating (9) with respect to λ_m . These derivatives are straightforward to compute and are not presented here.

D. Optimization Approach

The previous subsections describe the forward model and all of the elements included in the estimator specified in (7). However, since this estimator is the implicit maximizer of an objective function, an optimization algorithm must be adopted to iteratively approximate a solution. From the development in the previous section, we note that the objective function is differentiable and one may find the analytic expressions for the derivatives with respect to all of the estimation parameters:

$$\nabla_{\mu} L(\mu_*, \Lambda; y) = \left(\prod_{n=1}^N \mathbf{D} \left\{ \mathbf{W}(\lambda^{(n)}) s^{(n)} \right\} \right) \cdot \mathbf{A}^T [\mathbf{D} \{b\} \exp(-\mathbf{A}\mu) - y] \quad (14)$$

$$\frac{\partial}{\partial \lambda_m^{(n)}} L(\mu_*, \Lambda; y) = \left[\mathbf{D} \{ \mu_* \} \mathbf{D} \left\{ \prod_{k \neq n} \mathbf{W}(\lambda^{(k)}) s^{(k)} \right\} \dot{\mathbf{W}}_m(\lambda^{(n)}) s^{(n)} + \dot{\mathbf{W}}_m(\lambda^{(n)}) \mu_I^{(n)} \right]^T \cdot \mathbf{A}^T [\mathbf{D} \{b\} \exp(-\mathbf{A}\mu) - y] \quad (15)$$

Given the analytic gradient of the objective function it would be straightforward to use almost any gradient-based algorithm to jointly solve for the parameters, $\hat{\mu}$ and $\hat{\Lambda}$; however, this would ignore years of development of specialized algorithms customized for tomographic reconstruction. Thus, we propose using an alternating minimization approach where the image volume parameters, μ_* , are updated using a tomography-specific image update, and the registration parameters, Λ , are updated using a traditional gradient-based approach.

Since Λ is relatively small with only $6N$ elements (for rigidly transformed components), it is practical to use a quasi-Newton update approach using the Broyden–Fletcher–Goldfarb–Shanno (BFGS) updates [30]. (If N is large, one might resort to a limited-memory variant [31].) We adopt such an approach that includes a bracketing line search.

For image updates, we choose to use updates based on a paraboloidal surrogates approach. Such updates closely follow the surrogates-based approach introduced by Erdogan and Fessler in [32]; however, since there are significant changes to account for the modified forward model in (3) and (4), a modified derivation that includes the masking operation using $s^{(n)}$, and the additional components, $\mu_I^{(n)}$, is discussed in detail in the Appendix for an unregularized objective function. It is straightforward to modify the updates found in the Appendix using the techniques in [33] to include various penalty functions, to adopt an ordered-subsets subiterations for improved convergence rates, or to use precomputed curvatures for decreasing the computational load.

E. Implementation

To investigate the approach derived above, we implemented all routines in Matlab (The Mathworks, Inc., Natick MA) with calls to custom external libraries for the more computationally intensive aspects of the KCR methodology. Specifically, we implemented both matched Siddon [34] and separablefootprint [35] projectors and backprojectors in C/C++ using CUDA libraries for execution on GPU. The same approach was taken for 3D kernel-based interpolation via GPU.

Pseudocode for the entire alternating maximization approach for KCR is shown in Table 1. The code consists of two separate blocks: the first containing the registration parameter updates and the second containing the ordered-subsets-based volume updates. In this pseudocode, we have used a generalized iteration superscript, $[k, p, m]$, to indicate the main iteration number, k , the registration block subiteration number, p , and the image update block subiteration number, m . The symbol $[\cdot]$ is used for values not dependent on a particular index. Table 1 illustrates a few free parameter choices for the algorithm – namely, the number of registration subiterations, P , the number of ordered-subsets, M , and the overall number of iterations. The scaling factor, σ , is used in the guess for the initial inverse Hessian estimate, $\mathbf{H}^{[0,0,\cdot]}$, and needs to be chosen empirically based on the scale of the objective function, or via an initial line search to find the rough scale for an appropriate step size. Other salient features of this optimization include the reuse of \mathbf{H} between main loop iterations and the inclusion of regularization in the image update using the first and second derivatives of the quadratic roughness penalty (indicated by \hat{R} and \ddot{R} , respectively).

In the case where $M = 1$, image updates should monotonically converge to a local maximum, and when a line search that satisfies the Wolfe conditions is used, the registration updates should monotonically converge to a local maximum. Therefore, under these conditions, we would also expect convergence of the entire algorithm to a local maximum, though to our knowledge no proof of convergence is known. When $M > 1$ we would not expect true convergence, just as ordinary ordered-subsets approaches do not strictly converge. Other algorithmic permutations are possible, including precomputation of curvatures outside the

main loop and mixing registration updates between individual ordered-subsets image updates.

F. Computational Complexity

It is possible to specify approximate bounds on the computational complexity of the iterative approach outlined in Table I. To provide a general assessment of computational complexity not dependent on the speed of a particular hardware platform or software implementation, we have chosen to characterize the computational load in terms of the number of projections and/or backprojections, which are the operations that tend to dominate computation time. Assuming the computation time for projection and backprojection to be approximately equal, we denote the complexity for each of these operations as $O(Z)$. Thus, ignoring regularization, the complexity of a single iteration of a penalized-likelihood reconstruction using separable paraboloidal surrogates and precomputed curvatures is $O(2Z)$.

The KCR approach has added complexity associated with the registration updates. The gradient computation requires one projection-backprojection pair, and each subsequent function evaluation requires one additional projection. For a bracketing line search with an appropriate initial step length, one can expect three function evaluations from BFGS contributing $O(4Z)$ per registration update. Additionally, one must consider the registration operator and its derivatives. Labeling the cost associated with \mathbf{W} in (11) as $O(B)$, and recognizing that the derivative computations in (13) are dominated by the weighted sums of three blurring operations yields an additional $O(3B)$, we find that a typical BFGS iteration would take $O(4Z+12NB)$ for N component registrations. Thus, recalling P is the number of registration subiterations, the total computation time for a single KCR iteration is $O((4P+2)Z+12PNB)$.

In practice, the above complexity estimate is a worst-case scenario. If components are small relative to the imaging volume, projections need only be computed over a small field-of-view. Similarly, angular subsets may be used (much like ordered subsets for the image updates) eliminating P from the first term. Thus, for small components the complexity can approach $O(2Z+12PNB)$ per iteration and if computations are dominated by the projection and backprojection operations, the complexity approaches that of ordinary penalized-likelihood.

III. Results

A. Simulation Studies

To investigate the performance of the KCR approach, we conducted a number of studies using a simulated cone-beam computed tomography (CBCT) system. In this work we focused on the task of imaging for guidance of pedicle screw placement in spine surgery [36]. A CAD model for a polyaxial pedicle screw was obtained for the simulation studies and is illustrated in Figure 2. The model consists of two main components: a pivoting head (blue) and a threaded screw (red). For the anatomical background volume we used a digital phantom derived from a high-quality conventional CT scan of the lumbar region of a custom RANDO phantom (The Phantom Laboratory, Greenwich NY). Multiplanar slices of this phantom are shown in Figure 2.

For all of the investigations we considered an image volume of $300 \times 300 \times 128$ voxels (1 mm cubes). The pedicle screw was similarly voxelized and placed into the image volume using the object model in (4). Note that (4) additionally requires voxelized component masks, $s^{(n)}$, that were also derived from the CAD model. We assumed the pedicle screw to be homogeneous titanium with a linear attenuation coefficient of 0.3 mm^{-1} . For comparison, the soft tissue background in the digital phantom was 0.02 mm^{-1} . The system geometry was chosen to emulate a flat-panel detector C-arm system with a source-to-detector distance of 1200 mm and source-to-axis distance of 600 mm [37]. Projections were 360×150 pixels at $(1.552 \times 1.552) \text{ mm}^2$ pitch simulated over 360° using 360 angles. Simulated data were generated using a monoenergetic forward model and Poisson noise. Exposure levels were fixed at a uniform level of $b = 10^4$ photons per detector element in the unattenuated beam. As demonstrated in the following sections, this exposure level leads to significant photon starvation artifacts for traditional reconstruction approaches.

Three separate screw placement scenarios were investigated: 1) A unilateral monolithic (single component) pedicle screw; 2) bilateral monolithic screws; and 3) a unilateral two-component polyaxial screw. The registration parameters, Λ , were chosen to simulate proper placement of the screw (or screws) in one of the lumbar vertebra.

B. Convergence Properties

To investigate the convergence properties of the KCR approach, we reconstructed simulated data from the unilateral single-component screw placement scenario mentioned above. The initial guess for the background anatomy was a filtered-backprojection (FBP) image truncated at an upper attenuation of 0.03 mm^{-1} (in a coarse attempt to remove the pedicle screw from the anatomy, since the screw should not be present in the background volume, μ^*). A number of initial guesses were attempted for the registration parameters. For initial guesses sufficiently far from the solution, we observed “capture issues” in which the screw position estimate was lost in a local maximum (most often leaving the reconstructed field-of-view entirely so that the problem reduces to a traditional penalized-likelihood estimate). However, we found that initialization for which there was at least modest overlap of the components with the true location *in a portion of the projection data*, the approach converged toward the true pose (except for errors due to noise). For example, initiating KCR with a misregistered guess of pedicle screw location (perhaps totally non-overlapping in the 3D volume), but with significant overlap in at least a few projection angles, the approach was typically able to converge toward an accurate pose estimate. For initial guesses inside the capture range, residual errors in the pose estimate were consistently within one voxel (for translation parameters) and less than 0.5 degrees (in the rotation parameters) for the $b = 10^4$ exposure level.

Figure 3 illustrates sample iterations starting with the initial guess and showing the first four iterations using $P = 10$ and $M = 60$, for registration and image subiterations, respectively. Note that the joint nature of the algorithm is evident, with refinements in both the pose of the component and reduction of streak artifacts in the background image between outer loop iterations.

To better understand the relative effects of different choices of numbers of subiterations, we computed the convergence plots in Figure 4. The two sets of plots represent the convergence rates for a varying number of registration updates for a fixed number of ordered-subset updates, and vice versa. Specifically, we plotted the difference between the objective function at the k^{th} iteration and that at the solution, Φ^* (where the latter value was approximated using 100 iterations of KCR with $P = 10$, and a dynamic number of subsets, as fully described below).

Figure 4(a) shows the first 25 iterations for three approaches using 4, 6, and 10 registration updates with 60 subsets. Objective function differences are shown after both the registration update (halfway between integer k values) and the image update (indicated with the \odot symbol). We note that the convergence rate appears to increase with additional BFGS updates; however, increasing the number above $P = 10$ produced virtually the same convergence plots and suggested little to be gained from further emphasis on the registration updates. We also found that the image updates lead to a greater objective function increase than the registration updates.

In Figure 4(b), the number of registration updates was fixed at $P = 10$, and M was varied across 1, 5, 10, 20, 40, and 60 subsets. In these cases, the convergence rate had a strong dependence on the number of subsets (with performance akin to traditional ordered-subset approaches). While all methods appeared to increase the objective monotonically, we note a plateau beyond $M = 60$, which is expected since ordered-subsets are not strictly convergent. To avoid this plateau and maintain the excellent convergence rate, we implemented an approach in which the number of subsets is changed dynamically. Specifically, we started with 120 subsets and reduced the number of subsets every four outer iterations through the sequence 90, 60, 40, 30, 20, 10, 3, 2, and 1. This method is represented in the lowest plot in Figure 4(b). Since the dynamic subsets approach eventually uses a single subset, we would expect true convergence with this method. Note the stepwise improvements in function value every four iterations when the number of subsets is changed.

C. Image Quality Experiments

For image quality experiments, we considered all three pedicle screw implant scenarios and three different reconstruction methods. Specifically, we applied 200 iterations of the KCR approach with $P = 10$ and the dynamic subsets approach (with M decreasing through the aforementioned sequence every 20 iterations). For comparison, we also reconstructed the data using a FBP approach and a traditional penalized-likelihood (PL) estimator. For FBP processing, the data was pre-processed by thresholding the data slightly above zero (to eliminate zero count measurements) prior to logtransformation. (An alternate approach, not shown here, that interpolated over these zero values using a median filter was also applied with qualitatively nearly the same results.) We note that these FBP reconstructions do not utilize any specialized processing to mitigate metal artifacts and, as such, do not represent the best possible analytic reconstruction. Rather, they represent a baseline imaging performance when no metal artifact corrections have been applied.

Both the PL and KCR approach used a quadratic penalty with the same value of regularization parameter (β). Since the objectives are very similar, these two approaches

produce images with similar resolution properties¹ when the parameter is matched. We therefore selected β to qualitatively match the spatial resolution of the FBP images.

The results are presented in Figure 5 with the true digital phantom volume shown in the left column. As before, the position of components is indicated using a color overlay (red for the first component, blue for the second component). Such overlays are seen in the true volume images and the KCR images (but not in FBP or traditional PL, since those methods only include an image model and do not include a component registration model). We note that, if desired, one may always show a traditional attenuation image using (4) for KCR without the color overlay – e.g., for the single-component unilateral screw case as shown in Figure 6. Axial and coronal images are shown for each reconstruction technique and each screw placement scenario. The coronal slice was chosen to illustrate placement of the screw through the pedicle. This view is of particular clinical importance in assessing possible breach of the screw through the cortical bone – e.g., lateral breach (sub-optimal screw placement) or medial breach (intrusion on the spinal canal).

At the $b = 10^4$ exposure level, photon starvation artifacts due to the high attenuation of the screw are very apparent in the FBP reconstructions. These artifacts strongly diminish the low and medium contrast details in the surrounding areas in the axial images, and are worst for the bilateral screw placement scenario (which exhibits the greatest amount of photon starvation due to the increased amount of metal in the volume). In comparison, the PL images are significantly improved and eliminate most of the streak artifacts. However, there are still significant residual artifacts, particularly near the boundary of the screws, which obscure much of the vertebra and confound assessment of possible cortical breach in the coronal images. The KCR approach produces images that are nearly artifact-free. While there are some residual artifacts (most notably at the head of the screw), the contrast of the artifacts is extremely low compared to both FBP and PL. Moreover, the coronal images easily demonstrate the lack of cortical breaches in screw placement and agree well with the true volume images. The image quality for the KCR approach remains high even near the boundary of the screw.

As noted previously, the registration estimates for the KCR approach demonstrated sub-voxel errors even in the multiple component cases. Because the forward model uses an additive approach to integrate the attenuation of all components, overlapping components would tend to result in unrealistic attenuation profiles (e.g., twice the attenuation in the overlapping region) that give a poor fit to the measured data. Thus, the model effectively discourages component intersection without the explicit need for constraints based on component boundaries.

D. Regularization Experiments

Regularization for the KCR approach is somewhat unique. Recall from (7) that the roughness penalty is applied to background anatomy, and there is no penalty between the anatomy and the known components. This suggests that KCR is likely to exhibit a different

¹We note that resolution properties for penalized-likelihood approaches (including KCR) are generally space-variant and object-dependent making exact resolution matches difficult.

relationship between image quality and regularization than traditional implementations of penalized-likelihood. This relationship was investigated through two experiments on single-component pedicle screw reconstructions: 1) using a traditional quadratic penalty and a range of regularization parameters, and 2) a scenario comparing quadratic and non-quadratic penalties.

Figure 7 shows the same, previously investigated, unilateral single-component screw case across a range of regularization parameters using a quadratic penalty. Specifically, the same regularization parameter was used as in the previous study along with a value of β ten times greater and a value ten times less. The varying levels of regularization illustrate the expected increase in noise with lower β and decreasing spatial resolution with higher β . However, we note that all of the images are essentially free of the streaks associated with photon starvation due to the screw. This suggests that the streak elimination is not simply due to regularization, and that the integration of prior knowledge regarding the structure and composition of the screw does indeed improve the reconstruction. We note that lower regularization values appear to converge more slowly, and the KCR objective function may be more susceptible to local minima (particularly with respect to registration parameters) when the volume estimate becomes very noisy. The error in the translation estimate was 0.31, 0.13, and 0.05 voxels for the three regularization cases in Figure 7, ordered from lowest to highest β . Similarly, the mean absolute angular error was 0.24, 0.09, and 0.04 degrees. Thus, across this particular range of β , translation estimates appeared to benefit from increased regularization.

Quadratic penalties reduce noise in reconstructed volumes but also tend to blur true edges in an image. In fact, the artifacts and blooming about the implant in the PL reconstructions in Figure 5 may be a result of the quadratic penalty, whereas an edge-preserving penalty might better deal with the component-anatomy boundary. We note that it is straightforward to modify the KCR approach to use non-quadratic, edge-preserving penalties of the type in [16, 38–40] to provide a different set of noise-resolution trade-offs.

For the following experiments, we selected a modified, total-variation type penalty function that uses a quadratic penalty function for differences within a δ -neighborhood of the origin, and is linear (with unity slope) beyond that neighborhood. Again, the strength of this penalty is controlled by a scalar parameter, β . We apply this edge-preserving penalty in a variant of the single-component screw scenario (with the same acquisition technique as previously stated) for both PL and KCR and compare the results to the quadratic penalties.

For PL, selection of an optimal β, δ pair can be difficult and time consuming, particularly for high dynamic range reconstructions like those with metal components. We performed an exhaustive search reconstructing all β, δ pairs and selected the pair that yielded the (qualitatively) optimal result. Specifically, we sought to minimize reconstruction artifacts without creating an image dominated by piecewise constant plateaus in gray level variation². The edge-preserving penalty with optimal β, δ pair was compared with a quadratic penalty

²We note that one can eliminate all streaks in the reconstruction with a β, δ pair that effectively segments the image into five levels (air, water, liver, bone, screw) without gray level variation and considerable loss of fine detail structures.

whose noise and resolution characteristics were qualitatively similar in the anatomical background. The same regularization parameter values were used for both PL and KCR, and the results are summarized in Figure 8.

First, we note that the artifacts in the PL reconstruction may be greatly reduced using such a non-quadratic penalty. However, significant artifacts remain in the PL reconstruction, and the degree of artifact reduction is seen to tradeoff with the noise-resolution properties in the rest of the image. In contrast, the KCR technique is insensitive to regularization choices in terms of artifact strength. Thus, KCR appears robust to regularization choices and potentially simplifies the selection of regularization parameter, allowing for significant artifact reduction across a broad range of noise-resolution parameterization not obtainable from traditional PL approaches.

IV. Discussion

This paper reports a novel tomographic reconstruction approach referred to as KCR that allows for incorporation of known components into the object model. The position and pose of such components are estimated simultaneously with the background attenuation. Application of this technique was demonstrated in images containing metallic pedicle screws within a spine surgery context, but the technique is general for any kind of component for which the shape and composition is known – including objects with heterogeneous composition. In simulated pedicle screw reconstructions, KCR largely eliminated artifacts associated with photon starvation and allowed clear visualization of anatomy right up to the boundary of the implant, whereas traditional FBP and PL approaches exhibited substantial artifacts. As such, we anticipate the KCR method to hold significant promise in cone-beam CT (CBCT)-based imaging for interventional guidance – e.g., in spine procedures [36] as well as a host of other interventions for which the introduction of a surgical device confounds 3D image quality. Moreover, KCR may be an important tool in dose reduction, since it allows for successful operation in a photon starvation regime. Investigation in the context of other applications, such as diagnostic CT, imaging of hip and knee implants, and CT-based biopsy needle guidance will be considered in future work.

We note that only a basic FBP reconstruction was used for comparison and that more sophisticated approaches that include strategies for metal artifact reduction would likely have a less dramatic difference between FBP and KCR. While it is outside of the scope of this paper, future investigations should compare such FBP-based artifact reduction methods [8, 9, 14] with KCR to determine these differences.

The value of incorporation of this kind of prior knowledge has previously been reported with the development of a related reconstruction approach, object-constrained CT (OCCT), in [26–28]. Similar observations were made that artifacts associated with high-density known components could be largely mitigated through the incorporation of prior knowledge. Both KCR and OCCT rely on Poisson objectives and a likelihood-based objective function. However, KCR adopts a penalty term that allows for the advantages of a penalized-

likelihood approach, namely control of the noise-resolution tradeoff and improved convergence rates as well as improved fidelity of registration estimates (as discussed above).

In OCCT, component knowledge is incorporated through a constrained objective where specific voxel values must be equal to or greater than known attenuation coefficients, whereas KCR achieves such incorporation through the masking operation. The latter approach admits a straightforward extension to multiple known components without the need for explicit intersection detection. Both approaches utilize an alternating optimization scheme to solve for the respective objective functions, but employ different underlying algorithms. Specifically, for image updates: paraboloidal surrogates for KCR and so-called alternating minimization (AM) for OCCT; and for registration updates: quasi-Newton updates for KCR and various methods for OCCT. Deeper investigation is required to determine whether there are computational or convergence advantages for one approach over another.

One difference between KCR and OCCT involves the so-called “overshoots” in registration estimates discussed by Murphy et al [REF]. In that case, the algorithm was specifically modified using a dithering approach to avoid these overshoots and related streaking artifacts. Such issues were not observed with KCR. It may be that the penalty function in KCR effectively eliminates this problem; however, additional investigation and analysis is required to ascertain the underlying source.

One limitation of the basic KCR framework reported above is that the forward model employed in this initial implementation does not incorporate a polyenergetic beam and the physics (e.g., scatter) associated with a spectrum of x-rays. The mismatch between this monoenergetic algorithm and polyenergetic data will have an effect on convergence and residual artifacts. While these investigations are beyond the scope of this paper, an investigation of this mismatch and an incorporation of a polyenergetic forward model [41–44] is the subject of ongoing work, with initial results found in [45]. Such modifications will likely be important to application in real clinical data since beam hardening is a major contributor to metal artifacts. Similarly, it is unclear how important the role of scatter will be, but it is straightforward to include a known scatter fraction in the forward model.

One potential limitation of this particular implementation of the KCR framework arises in cases where one has inexact knowledge of certain components – for example: metal rods used in vertebral fixation that are bent during the surgical procedure; knee or hip implants that have experienced significant wear prior to a surgical revision; or needle guidance in which the needle flexes upon insertion. However, we note that the mapping in (9) allows straightforward modification to a deformable registration model to encompass some of these scenarios (particularly when the differences can be expressed in a relatively low-dimensional parameter space). Such modifications to the model are underway and preliminary investigations have been reported in [46].

The results summarized above focused largely on the image quality resulting from KCR; however, there are two outputs from the estimator: the image volume and the registration parameters. The image volume is likely the primary outcome, but the accurate estimation of

the pose values themselves may also have clinical value in assessing the placement of components introduced in the body. For example, the pose output of the KCR method allows comparison of pre-operative planning to the actual device delivery and may offer value in quantifying surgical precision.

Acknowledgments

The authors wish to thank DePuy Spine, Inc. (US and Switzerland) for providing CAD models of pedicle screws for this work. Research supported in part by the National Institutes of Health R01-CA-127444.

VI. Appendix: Derivation of Image Updates

The original derivations by Erdogan and Fessler [32] express (5) as the summation of marginal log-likelihoods, $h_i(l_i)$, such that

$$\begin{aligned} L(\mu, \Lambda; y) &= \sum_i h_i(l_i) \\ h_i(l_i) &= y_i \log[b_i \exp(-l_i)] - b_i \exp(-l_i). \end{aligned} \quad (16)$$

We make the following definitions

$$\begin{aligned} g_i(t_i) &\triangleq h_i(t_i + d_i) = h_i(l_i) \\ t_i &\triangleq \mathbf{B}\mu_* = \mathbf{A} \left(\prod_{n=1}^N \mathbf{D} \{ \mathbf{W}(\lambda^{(n)}) s^{(n)} \} \right) \mu_* \\ d_i &\triangleq \mathbf{A} \left(\sum_{n=1}^N \mathbf{W}(\lambda^{(n)}) \mu_i^{(n)} \right) \end{aligned} \quad (17)$$

and note that $t_i \geq 0$ and $d_i \geq 0$ since \mathbf{A} represents a projection operator with nonnegative elements, since \mathbf{W} was chosen to use B-spline approximation which preserves nonnegativity, and because $s^{(n)}$, μ_* , and $\mu_i^{(n)}$ are also nonnegative. (The matrix \mathbf{B} represents a kind of modified system matrix.) Because $t_i \geq 0$ and $l_i = t_i + d_i \geq 0$, the function $g_i(t_i)$ satisfies the conditions of Theorem 1 in [32] and we write a paraboloidal surrogate for $g_i(t_i)$ with optimal curvature as

$$\begin{aligned} q_i(t_i; t_i^{[k]}) &= g_i(t_i^{[k]}) + \dot{g}_i(t_i^{[k]}) (t - t_i^{[k]}) - \frac{1}{2} c_i^{[k]} (t - t_i^{[k]})^2 \\ c_i^{[k]}(t_i^{[k]}) &= \begin{cases} \left[-2 \frac{g_i(0) - g_i(t_i^{[k]}) + \dot{g}_i(t_i^{[k]}) t_i^{[k]}}{(t_i^{[k]})^2} \right] & t_i^{[k]} > 0 \\ \left[-\ddot{g}_i(t_i^{[k]}) \right]_+ & t_i^{[k]} = 0 \end{cases} \end{aligned} \quad (18)$$

where the superscript $[k]$ denotes the value at the k^{th} iteration, c_i denotes the curvature of the paraboloidal surrogate, and $[\cdot]_+$ denotes a function to ensure nonnegativity (truncation at zero). The first and second derivatives of g , denoted \dot{g} and \ddot{g} , respectively, are straightforward to compute from (16) and (17). While the surrogates in (18) may be used to form an overall surrogate for the log-likelihood, we apply the convexity trick of De Pierro [47] that was also used in [33],

$$q_i \left([\mathbf{B}\mu_*]_i; t_i^{[k]} \right) = q_i \left(\sum_i \alpha_{ij} \left[\frac{b_{ij}}{\alpha_{ij}} (\mu_{*j} - \mu_{*j}^{[k]}) + t_i^{[k]} \right]; t_i^{[k]} \right) \geq \sum_i \alpha_{ij} q_i \left(\frac{b_{ij}}{\alpha_{ij}} (\mu_{*j} - \mu_{*j}^{[k]}) + t_i^{[k]}; t_i^{[k]} \right) \quad (19)$$

where b_{ij} are elements of the modified system matrix \mathbf{B} , to obtain separable surrogates

$$Q_j \left(\mu_{*j}; \mu_*^{[k]} \right) \triangleq \sum_i \alpha_{ij} q_i \left(\frac{b_{ij}}{\alpha_{ij}} (\mu_{*j} - \mu_{*j}^{[k]}) + t_i^{[k]}; t_i^{[k]} \right) \\ L(\mu_*, \Lambda; y) \geq \sum_j Q_j \left(\mu_{*j}; \mu_*^{[k]} \right), \quad (20)$$

where we choose $\alpha_{ij} = b_{ij} / \sum_k b_{ik}$. Thus, for an unregularized objective, we may apply the following updates

$$\mu_{*j}^{[k+1]} = \underset{\mu_{*j} \geq 0}{\operatorname{arg\,max}} Q_j \left(\mu_{*j}; \mu_*^{[k]} \right) \\ = \left[\mu_{*j}^{[k]} + \frac{\sum_i b_{ij} \dot{g}_i^{[k]} \left(t_i^{[k]} \right)}{\sum_i b_{ij} \left(\sum_i b_{ij} \right) c_i^{[k]} \left(t_i^{[k]} \right)} \right]_+ \quad (21)$$

References

1. De Man B, et al. Metal streak artifacts in X-ray computed tomography: A simulation study. *IEEE Trans Nuclear Science*. 1999; 46:691–696.
2. Barrett JF, Keat N. Artifacts in CT: recognition and avoidance. *Radiographics*. Nov-Dec;2004 24:1679–91. [PubMed: 15537976]
3. Stulberg SD, et al. Monitoring pelvic osteolysis following total hip replacement surgery: an algorithm for surveillance. *J Bone Joint Surg Am*. 2002; 84-A(Suppl 2):116–22. [PubMed: 12479349]
4. Holly LT, Foley KT. Three-dimensional fluoroscopy-guided percutaneous thoracolumbar pedicle screw placement. Technical note. *J Neurosurg*. Oct.2003 99:324–9. [PubMed: 14563154]
5. Wang MY, et al. Reliability of three-dimensional fluoroscopy for detecting pedicle screw violations in the thoracic and lumbar spine. *Neurosurgery*. May.2004 54:1138–42. discussion 1142–3. [PubMed: 15113468]
6. Daly B, et al. Percutaneous abdominal and pelvic interventional procedures using CT fluoroscopy guidance. *AJR Am J Roentgenol*. Sep.1999 173:637–44. [PubMed: 10470894]
7. De Man B, et al. Reduction of metal steak artifacts in X-ray computed tomography using a transmission maximum a posteriori algorithm. *IEEE Trans Nuclear Science*. 2000; 47:977–981. 2000.
8. Glover GH, Pelc NJ. An algorithm for the reduction of metal clip artifacts in CT reconstructions. *Med Phys*. Nov-Dec;1981 8:799–807. [PubMed: 7322078]
9. Kalender WA, et al. Reduction of CT artifacts caused by metallic implants. *Radiology*. Aug.1987 164:576–7. [PubMed: 3602406]
10. Li H, et al. Metal artifact suppression from reformatted projections in multislice helical CT using dual-front active contours. *Med Phys*. Oct.2010 37:5155–64. [PubMed: 21089749]
11. Robertson DD, et al. Total hip prosthesis metal-artifact suppression using iterative deblurring reconstruction. *J Comput Assist Tomogr*. Mar-Apr;1997 21:293–8. [PubMed: 9071303]

12. Wang G, et al. Iterative deblurring for CT metal artifact reduction. *IEEE Trans Med Imaging*. 1996; 15:657–64. [PubMed: 18215947]
13. Watzke O, Kalender WA. A pragmatic approach to metal artifact reduction in CT: merging of metal artifact reduced images. *Eur Radiol*. May.2004 14:849–56. [PubMed: 15014974]
14. Medoff BP, et al. Iterative Convolution Backprojection Algorithms for Image-Reconstruction from Limited Data. *Journal of the Optical Society of America*. 1983; 73:1493–1500.
15. Rinkel J, et al. Computed tomographic metal artifact reduction for the detection and quantitation of small features near large metallic implants: a comparison of published methods. *J Comput Assist Tomogr*. Jul-Aug;2008 32:621–9. [PubMed: 18664852]
16. Lange K. Convergence of EM image reconstruction algorithms with Gibbs smoothing. *IEEE Trans Med Imaging*. 1990; 9:439–46. [PubMed: 18222791]
17. Hebert T, Leahy R. A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors. *IEEE Trans Med Imaging*. 1989; 8:194–202. [PubMed: 18230517]
18. Thibault JB, et al. A three-dimensional statistical approach to improved image quality for multislice helical CT. *Med Phys*. Nov.2007 34:4526–44. [PubMed: 18072519]
19. Wang J, et al. Iterative image reconstruction for CBCT using edgepreserving prior. *Med Phys*. Jan. 2009 36:252–60. [PubMed: 19235393]
20. Chen GH, et al. Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets. *Med Phys*. Feb. 2008 35:660–3. [PubMed: 18383687]
21. Stayman J, et al. Penalized-likelihood reconstruction for sparse data acquisitions with unregistered prior images and compressed sensing penalties. *SPIE Medical Imaging*. 2011
22. Fessler JA. Optimization transfer approach to joint registration/reconstruction for motion-compensated image reconstruction. presented at the ISBI. 2010
23. Jacobson M, Fessler JA. Joint estimation of image and deformation parameters in motion-corrected PET. *Proc IEEE Nuc Sci Symp Med Im Conf*. 2003:3290–3294.
24. Chun SY, Fessler JA. Joint image reconstruction and nonrigid motion estimation with a simple penalty that encourages local invertibility. *Proc SPIE 7258, Medical Imaging 2009: Phys Med Im*. 2009:72580U.
25. Stayman JW, et al. Likelihood-based CT Reconstruction of Objects Containing Known Components. *Int'l Mtg Fully 3D Image Recon, Potsdam, Germany*. 2011:254–7.
26. Snyder DL, et al. Deblurring subject to nonnegativity constraints when known functions are present with application to object-constrained computerized tomography. *IEEE Trans Med Imaging*. Oct.2001 20:1009–17. [PubMed: 11686436]
27. Murphy RJ, et al. Pose estimation of known objects during transmission tomographic image reconstruction. *IEEE Trans Med Imaging*. Oct.2006 25:1392–404. [PubMed: 17024842]
28. Williamson JF, et al. Prospects for quantitative computed tomography imaging in the presence of foreign metal bodies using statistical image reconstruction. *Med Phys*. Oct.2002 29:2404–18. [PubMed: 12408315]
29. Thevenaz P, et al. Interpolation revisited. *IEEE Trans Med Imaging*. Jul.2000 19:739–58. [PubMed: 11055789]
30. Luenberger, DG. *Linear and nonlinear programming*. 3rd. New York: Springer; 2007.
31. Liu DC, Nocedal J. On the Limited Memory BFGS Method for Large-Scale Optimization. *Mathematical Programming*. Dec.1989 45:503–528.
32. Erdogan H, Fessler JA. Monotonic algorithms for transmission tomography. *IEEE Trans Med Imaging*. Sep.1999 18:801–14. [PubMed: 10571385]
33. Erdogan H, Fessler JA. Ordered subsets algorithms for transmission tomography. *Phys Med Biol*. Nov.1999 44:2835–51. [PubMed: 10588288]
34. Siddon RL. Fast calculation of the exact radiological path for a three-dimensional CT array. *Med Phys*. Mar-Apr;1985 12:252–5. [PubMed: 4000088]
35. Long Y, et al. 3D forward and back-projection for X-ray CT using separable footprints. *IEEE Trans Med Imaging*. Nov.2010 29:1839–50. [PubMed: 20529732]

36. Siewerdsen JH, et al. Volume CT with a flat-panel detector on a mobile, isocentric C-arm: Pre-clinical investigation in guidance of minimally invasive surgery. *Medical Physics*. Jan.2005 32:241–254. [PubMed: 15719975]
37. Schafer S, et al. Mobile C-arm cone-beam CT for guidance of spine surgery: Image quality, radiation dose, and integration with interventional guidance. *Medical Physics*. Aug.2011 38:4563–4574. [PubMed: 21928628]
38. Vogel CR, Oman ME. Fast, robust total variation-based reconstruction of noisy, blurred images. *IEEE Trans Image Process*. 1998; 7:813–24. [PubMed: 18276295]
39. Pan X, et al. Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction? *Inverse Probl*. Jan 1.2009 25:1230009. [PubMed: 20376330]
40. Sidky EY, Pan X. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Phys Med Biol*. Sep 7.2008 53:4777–807. [PubMed: 18701771]
41. De Man B, et al. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Trans Med Imaging*. Oct.2001 20:999–1008. [PubMed: 11686446]
42. Elbakri IA, Fessler JA. Statistical image reconstruction for polyenergetic X-ray computed tomography. *IEEE Trans Med Imaging*. Feb.2002 21:89–99. [PubMed: 11929108]
43. Elbakri IA, Fessler JA. Segmentation-free statistical image reconstruction for polyenergetic x-ray computed tomography with experimental validation. *Phys Med Biol*. Aug 7.2003 48:2453–77. [PubMed: 12953909]
44. Srivastava S, Fessler JA. Simplified statistical image reconstruction algorithm for polyenergetic X-ray CT. 2005 IEEE Nuclear Science Symposium Conference Record. 2005; 1–5:1551–1555.
45. Zbijewski, W., et al. CT Reconstruction Using Spectral and Morphological Prior Knowledge: Application to Imaging the Prosthetic Knee; The Second International Conference on Image Formation in X-Ray Computed Tomography; Salt Lake City, UT. 2012.
46. Stayman JW, et al. Model-based Reconstruction of Objects with Inexactly Known Components. SPIE Medical Imaging, San Diego, CA. 2012:83131S-1-6.
47. De Pierro AR. On the relation between the ISRA and the EM algorithm for positron emission tomography. *IEEE Trans Med Imaging*. 1993; 12:328–33. [PubMed: 18218422]




Figure 1. Illustration of kernel-based interpolation in two dimensions. (Upper Left) Mapping of the moving image (v) to the transformed image (u). (Upper Right) A single transformed image point is computed based on a neighborhood of values in the moving image. (Lower) Kernels that are a function of the position in the transformed image are applied in succession along each dimension to yield the value in the transformed image at that point.




Figure 2.

(a) CAD model of a pedicle screw. Two components of the polyaxial model are illustrated in red and blue. (b–d) Axial, sagittal, and coronal slices of the digital phantom used as a true representation of the anatomical background.




Figure 3.

A sequence of KCR estimated axial slices and implant poses from initial guess to iterations one through three. The pose estimate of the pedicle screw is represented by the red image overlay on the estimated background volume. Iteration zero represents initialization by an FBP reconstruction for the image background and arbitrary placement of the pedicle screw component model in proximity to the true location. Note the simultaneous nature of the estimation process with both registration updates and image updates between successive iterations.




Figure 4. (a) Convergence plots for KCR using alternating applications of image updates using 60 subsets and registration updates using 4 (∇), 6 (\circ), and 10 (\square), BFGS iterations. (b) Convergence plots for alternating updates of 10 BFGS iterations and ordered subsets based image updates using 1 (∇), 5 (\circ), 10 (\square), 20 (\triangle), 40 (\diamond), and 60 (\star) subsets, as well as an approach that uses a dynamic number of subsets (\triangleright). Objective function differences are plotted after both the registration block (identifying symbol) and after the image update (\odot).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 5.

Illustration of the image quality of KCR compared to FBP and traditional PL estimation. Axial and coronal images are presented for each method and each pedicle screw implant scenario: Top) Unilateral single-component screw; Middle) Bilateral single-component screws; Bottom) Unilateral two-component polyaxial screw. The window and level for all images is 500 and 150 HU, respectively. The color overlay in the true image (left) shows the true pose of the component(s), whereas that in the KCR reconstruction (right) shows the pose as estimated by the simultaneous registration and reconstruction process.




Figure 6. Illustration of the total attenuation (μ) for the single component unilateral screw case: (Left) True phantom, (Right) KCR. Since KCR estimates both the background attenuation (μ_*) and the registration parameters, Λ , one may represent the result as either a traditional attenuation image (Figure 6 (right)) or with color overlay (Figure 5 (right)) – whichever better suits display preferences and dynamic range. The window and level for all images is 800 and 300 HU, respectively.




Figure 7.

The effect of varying the regularization parameter in KCR. While there is a clear noise-resolution trade-off (viz., larger β decreasing the noise and sacrificing spatial resolution), all of the images are largely free of streak artifacts. The window and level for all images is 500 and 150 HU, respectively.




Figure 8. A comparison of quadratic versus edge-preserving penalties for both penalized-likelihood and KCR approaches. Reconstructions of a single pedicle screw implant (truth image near the implant shown in the yellow inset) are shown for A) PL with quadratic penalty, B) PL with edge-preserving penalty, C) KCR with quadratic penalty, and D) KCR with edge-preserving penalty. The regularization parameters for PL and KCR are matched for each choice of penalty, but the edge-preserving penalty has been optimized for PL image quality. While the edge-preserving penalty may be tuned to mitigate artifacts associated with the implant in the PL reconstruction, KCR can provide nearly artifact-free reconstructions across a range of parameter choices.

Table I

Outline of the alternating maximization used by the KCR Algorithm

```

 $\mu_*^{[0, \cdot, 0]}$  = Initial reconstruction (FBP)
 $\Lambda^{[0, 0, \cdot]}$  = Initial guess for registration parameters
 $\mathbf{H}^{[0, 0, \cdot]}$  =  $\alpha \mathbf{I}$ , initial guess for inverse Hessian
for  $k=0$  to max_iterations-1,
  % Registration Update Block
  for  $r=1$  to  $P$  (number of registration updates)

    Compute  $\nabla_{\Lambda} L \left( \mu_*^{[k, \cdot, 0]}, \Lambda^{[k, p-1, \cdot]} \right)$ 
     $\mathbf{H}^{[k, p, \cdot]}$  = BFGS update using new gradient

     $\hat{\gamma}$  = line search in  $\Lambda^{[k, p-1, \cdot]} + \gamma \mathbf{H}^{[k, p, \cdot]} \nabla_{\Lambda} L \left( \mu_*^{[k, \cdot, 0]}, \Lambda^{[k, p-1, \cdot]} \right)$ 

     $\Lambda^{[k, p, \cdot]} = \Lambda^{[k, p-1, \cdot]} + \hat{\gamma} \mathbf{H}^{[k, p, \cdot]} \nabla_{\Lambda} L \left( \mu_*^{[k, \cdot, 0]}, \Lambda^{[k, p-1, \cdot]} \right)$ 

  end

  Compute  $d_i^{[k, p, \cdot]} \left( \Lambda^{[k, p, \cdot]} \right)$ 
  % Image Update Block
  for  $m=1$  to  $M$  (number of subsets)

    Compute curvatures,  $c_i^{[k, \cdot, m]} \left( \mu_*^{[k, \cdot, m-1]}, \Lambda^{[k, p, \cdot]} \right)$ 

    Compute modified line integrals,  $t_i^{[k, \cdot, m]} \left( \mu_*^{[k, \cdot, m-1]}, \Lambda^{[k, P, \cdot]} \right)$ 

    
$$\mu_{*j}^{[k, \cdot, m]} = \left[ \mu_{*j}^{[k, \cdot, m-1]} + \frac{M \sum_{i \in \Omega_m} b_{ij} \delta_i \left( t_i^{[k, \cdot, m-1]}, d_i^{[k, P, \cdot]} \right) - \beta R_j \left( \mu_{*j}^{[k, \cdot, m-1]} \right)}{M \sum_{i \in \Omega_m} b_{ij} \left( b_{ij} c_i^{[k, \cdot, m-1]} + 2\beta R_j \left( \mu_{*j}^{[k, \cdot, m-1]} \right) \right)} \right]_+$$


  end

   $\mu_*^{[k+1, \cdot, 0]} = \mu_*^{[k, \cdot, M]}$ ,  $\Lambda^{[k+1, 0, \cdot]} = \Lambda^{[k, P, \cdot]}$ ,  $\mathbf{H}^{[k+1, 0, \cdot]} = \mathbf{H}^{[k, P, \cdot]}$ 
end

```
