Incorporating known information into image reconstruction algorithms for transmission tomography

Ryan J. Murphy[a], Joseph A. O'Sullivan[a], Jasenka Benac[a], Donald L. Snyder[a],
Bruce R. Whiting[b], David G. Politte[b], Jeffrey F. Williamson[b]

[a]Electronic Systems & Signals Research Laboratory, Washington University
One Brookings Drive, St. Louis, MO 63130 USA
[b]Mallinckrodt Institute of Radiology, Washington University, St. Louis, MO 63110 USA

ABSTRACT

We propose an alternating minimization (AM) image estimation algorithm for iteratively reconstructing transmission tomography images. The algorithm is based on a model that accounts for much of the underlying physics, including Poisson noise in the measured data, beam hardening of polyenergetic radiation, energy dependence of the attenuation coefficients and scatter. It is well-known that these nonlinear phenomena can cause severe artifacts throughout the image when high-density objects are present in soft tissue, especially when using the conventional technique of filtered back projection (FBP). If we assume no prior knowledge of the high-density object(s), our proposed algorithm yields much improved images in comparison to FBP, but retains significant streaking between the high-density regions. When we incorporate the knowledge of the attenuation and pose parameters of the high-density objects into the algorithm, our simulations yield images with greatly reduced artifacts. To accomplish this, we adapted the algorithm to perform a search at each iteration (or after every n iterations) to find the optimal pose of the object before updating the image. The final iteration returns pose values within 0.1 millimeters and 0.01 degrees of the actual location of the high-density structures.

Keywords: metal artifact reduction, transmission tomography, alternating minimization, image reconstruction, maximum likelihood

1. INTRODUCTION

1.1. Filtered Back Projection

The problem of image reconstruction using computerized tomography (CT) data has been the subject of widespread research. The standard reconstruction technique, filtered back projection (FBP), has been shown to produce high-quality images with little noise in a short amount of time. The FBP algorithm is based on the assumption that the set of CT measurements is the Radon transform of the linear attenuation coefficients of the scanned object; thus, the inverse transform results in an image of the attenuation coefficients of the anatomical structures.

This filtered back projection algorithm works very well when the object's attenuation coefficients are slowly varying, as is the case for most scans of human anatomy. When high-density objects (such as surgical clips, metal rods, hip prostheses, and dental fillings) are present amongst the soft tissue, however, the assumptions necessary for FBP are no longer applicable. Due to many nonlinear effects in the underlying physics, such as noisy projection data, spectrum hardening of polyenergetic radiation, scatter, energy dependence of the attenuation coefficients, partial volume effects, and exponential edge-gradient effects (EEGE), the linear model

Further author information: (Send correspondence to RJM): E-mail: rm3@essrl.wustl.edu, Telephone: (314) 935-7547.

upon which FBP is based degrades rapidly. In this paper, we will present a more rigorous physical model that leads to a class of reconstruction algorithms that accounts for these nonlinearities.

These nonlinearities are often manifested in the reconstructed image as fine white streaks emanating in all directions from the metal object; or, as heavy dark bands that connect a pair of metal objects. (White corresponds to attenuation coefficients that are higher than the surrounding areas, and black signifies lower coefficients). Other examples of streaking include white or black streaks connecting the edges of high-density bodies (EEGE) and light or dark rings around the perimeters of the objects (e.g. Gibbs overshoots). These streaks are collectively referred to as "metal artifacts", and the problem of metal artifact reduction (MAR) has been extensively researched for quite some time.[1–4]

1.2. Previous Work in MAR

Much of the earlier work in MAR has assumed that the metal objects are completely opaque, thus creating regions of "missing" data in the projections, and then has applied some correction to the data. For instance, Kalender *et al.*[1] replaced the missing data with a linear interpolation from the adjacent raw data not "hidden" by the metal. Others, such as Glover and Pelc,[2] implemented a "completion" of the missing data by substituting in calculated values consistent with a projection of an object whose density is the average of the surrounding areas. These algorithms worked with varying degrees of success, seemingly depending on the complexity of the structure of the objects.

1.3. Recent Iterative Techniques

More recently, however, the focus in MAR research has shifted to iterative approaches. These techniques have rarely been used thus far because of their high computational complexity and their need for large amounts of computer storage, but the advances in computer technology have made iterative techniques feasible. Wang, Snyder, *et al.*,[5] in a precursor to our current method, used a version of the expectation-maximization (EM) algorithm to iteratively reconstruct an image using *a priori* knowledge of a synthetic metal object. The object, however, was constrained to be opaque and convex and also to have a known location, thus limiting its usefulness.

These ideas were extended further by Snyder, O'Sullivan, *et al.*,[6] by using a more model-based approach that removed the limitations on the object. In this approach, any knowledge of the object's nonuniform attenuation characteristics, shape, and location could be directly incorporated into the image formation algorithm. The model, however, was based on an EM algorithm for emission tomography, in which it is assumed that the CT detectors are measuring the line integral of the attenuation coefficients in the X-ray path; i.e., the sum of the attenuation values of each pixel on a ray connecting the source and a specific detector. This assumption leads to a further assumption that the overall attenuation of the X-ray beam is a linear function. Thus, this model, though superior to FBP for noiseless data, especially when we directly used the knowledge of the metal objects, was severely degraded when we added nonlinearities (e.g. noise, polyenergetic X-ray sources) into the simulated data.

1.4. A Model for Transmission Tomography

Our work with the emission tomography model demonstrated quite clearly that a model accounting for nonlinear physical phenomena is crucial for an iterative approach to metal artifact reduction. It also showed, however, that there is merit in the idea that we can use *a priori* knowledge of high-density objects while forming the image, and produce a higher-quality image as a result.

Transmission tomography much more accurately describes the data acquisition process of an X-ray CT scanner. Photons (of varying energies) from a collimated source are transmitted through an object, and the nonabsorbed photons (or their energies) are then recorded by an array of detectors. Standard reconstruction algorithms then use the logarithm of this data as an estimate of the projections in attenuation space.

In 1984, Lange and Carson[7] derived an EM algorithm for transmission tomography to obtain the maximum likelihood image estimate using a model based on the actual transmission measurements and photon behavior. The M-step of their algorithm, however, does not yield a closed-form solution, and many researchers have been deriving increasingly improved algorithms ever since.

Our approach described herein will restate the maximum likelihood estimation as a double minimization of an I-divergence and then use an iterative alternating minimization (AM) algorithm to perform the optimization. The first minimization enforces a constraint using the measured data, while the second minimization then seeks to minimize the I-divergence between the measured data and the data estimated by our model. The image which yields the best data estimate (i.e., has the smallest I-divergence) is selected.

2. IMAGE FORMATION ALGORITHM

2.1. The Data Model

We can model the X-ray source as emitting either monoenergetic or polyenergetic radiation, and we will assume for the remainder of this paper that the detectors are photon-counting devices. This allows us to assume that the measured data (labeled $d(y)$) are Poisson-distributed, with means defined by

$$g(y) = \sum_{E} I_0(y, E) \exp\left(- \sum_{x \in X} h(y|x)\mu(x, E) \right) + \beta(y). \tag{1}$$

I_0 is the mean number of incident photons from the source over a spectrum of energies (we look at 1-keV steps from 19 to 120 keV); $h(y|x)$ is a nonnegative kernel determined by the scanner geometry–we view it as the average path length through image pixel x on the path defined by source-detector pair y; $\mu(x, E)$ is the energy-dependent attenuation coefficient (in units of inverse length) of pixel x. More specifically, we define it as

$$\mu(x, E) = \sum_{i=1}^{I} \mu_i(E) c_i(x), \tag{2}$$

where $\mu_i(E)$ is the attenuation spectrum (in mm^{-1}) for the ith constituent (e.g. constituents could be water, bone, metals) and $c_i(x)$ is a unitless specific gravity function of the ith constituent for pixel x. The *forward projection* term,

$$\sum_{x \in X} h(y|x)\mu(x, E),$$

approximates the line integral of attenuation through the object, and the exponential of this term is the survival probability for the incoming photons. Finally, the term $\beta(y)$ is the mean of the background events and we use it to incorporate photon scatter. We usually assume a known constant $\beta(y)$ (setting it to zero indicates that we are ignoring scatter), but our algorithm can be extended to estimate the scatter parameters as part of the alternating minimization procedure.

To simulate monoenergetic data, equation (1) is adjusted slightly so that all photons are assumed to be at the average energy, $E_0 \approx 75$ keV, so I_0 is the mean number of *total* photons, and the attenuation coefficients are determined by their values at E_0.

Thus, the model can simulate data that accounts for the polyenergetic nature of X-ray sources, Poisson noise, and scatter, and any combination of these effects can be "turned off" in the model at any time.

2.2. Alternating Minimization (AM)

In the image formation problem, we are trying to reconstruct the attenuation coefficients of the objects, $\mu(x, E)$. To do so, we wish to maximize the likelihood between the measured data, $d(y)$, and the data estimated by our model, $g(y)$. Equivalently, this problem can be written as minimizing the I-divergence between d and g:

$$\min_{\mu(x,E)} I(d\|g). \tag{3}$$

Using constrained minimization techniques described by O'Sullivan,[8] this can be rewritten as the following double minimization:

$$\min_{q} \min_{p} I(p\|q). \tag{4}$$

Here, p is defined to belong to a linear family, $\mathcal{L}(d)$, and is used to enforce constraints imposed by the measured data:

$$\mathcal{L}(d) \equiv \{p(y,E) \geq 0 : \sum_E p(y,E) = d(y)\}. \tag{5}$$

The q is defined as being a mean of the data for a given energy:

$$q(y,E) = I_0(y,E) \exp\left(-\sum_{x \in X} h(y|x)\mu(x,E)\right), \tag{6}$$

so comparing this with (1) yields:

$$g(y) = \sum_E q(y,E) + \beta(y) \tag{7}$$

The resulting algorithm alternates between estimating the two variables, p and q, and thus is classified as an alternating minimization (AM) algorithm. The algorithm, as discussed in O'Sullivan and Benac,[9] is as follows:

- Choose an initial guess of attenuation coefficients, $\hat{c}_i^{(0)}(x)$. (Recall from (2) that the coefficients, $\mu(x,E)$, are defined as
$$\mu(x,E) = \sum_{i=1}^{I} \mu_i(E) c_i(x).$$
- Using the kth estimate of attenuation coefficients, compute:
$$\hat{q}^{(k)}(y,E) = I_0(y,E) \exp\left(-\sum_{x \in X} \sum_{i=1}^{I} \mu_i(E) h(y|x) \hat{c}_i^{(k)}(x)\right). \tag{8}$$
- Using the kth estimate of q, compute:
$$\hat{p}^{(k)}(y,E) = \hat{q}^{(k)}(y,E) \frac{d(y)}{\sum_{E'} \hat{q}^{(k)}(y,E') + \beta(y)}. \tag{9}$$
- Compute the *backprojections* of the estimates:
$$\tilde{b}_i^{(k)}(x) = \sum_{y \in Y} \sum_E \mu_i(E) h(y|x) \hat{p}^{(k)}(y,E), \tag{10}$$
$$\hat{b}_i^{(k)}(x) = \sum_{y \in Y} \sum_E \mu_i(E) h(y|x) \hat{q}^{(k)}(y,E). \tag{11}$$
- Form estimate $k+1$ of the attenuation coefficients:
$$\hat{c}_i^{(k+1)}(x) = \hat{c}_i^{(k)}(x) - \frac{1}{Z_i(x)} \ln\left(\frac{\tilde{b}_i^{(k)}(x)}{\hat{b}_i^{(k)}(x)}\right), \tag{12}$$
where $Z_i(x)$ is a scaling function that must satisfy the constraint,
$$\sum_{x \in X} \sum_{i=1}^{I} \mu_i(E) h(y|x) \frac{1}{Z_i(x)} \leq 1. \tag{13}$$

2.3. Incorporating Prior Knowledge

Often, there are regions of the image about which we have prior knowledge. For example, if a patient had a hip prosthesis, we might know the object's shape, composition (likely of high-density metal), and attenuation characteristics. Usually, however, we will not know the object's exact location in the patient or in the scanner's field of view.

2.3.1. Objects at known pose

Let us assume briefly that the object's location and orientation (collectively called the "pose") is known to us as well. The above algorithm then needs some modification. First, we divide the set of image pixels, X, into two regions: $X_m(\theta)$, the support pixels of the known objects at pose θ, and $X_s(\theta) = X - X_m(\theta)$, the surrounding pixels in the image. Also, let $C(\theta)$ be the set of images that are zero in $X_m(\theta)$; i.e.,

$$C(\theta) = \{c_s(x),\ x \in X : c_s(x) = 0,\ x \in X_m(\theta)\}.$$

Thus, an admissible image containing the known object at pose θ is of the form:

$$c(x:\theta) = c_s(x) + c_m(x:\theta) = \begin{cases} c_s(x), & x \in X_s(\theta),\ c_s(\cdot) \in C(\theta) \\ c_m(x:\theta), & x \in X_m(\theta) \end{cases} . \tag{14}$$

The resulting algorithm, for each iteration, then becomes: 1) update the image according to equation (12), 2) set the pixels in $X_m(\theta)$ to 0, and 3) add this image to $c_m(x:\theta)$, which contains all zeros except in $X_m(\theta)$, where it contains the attenuation coefficients of the known objects. This procedure is equivalent to replacing the pixels in $X_m(\theta)$ with $c_m(x:\theta)$.

2.3.2. Objects at known pose, with partial volumes

In the last section, we assumed that a pixel either contained the known object(s) or it did not. This yields very nice results for synthetic models, as shown in Section 4, but usually an object will only partially cover some of the pixels. We label this subset of edge pixels $X_e(\theta)$. We model the "true" value of a pixel in $X_e(\theta)$ to be $\alpha \cdot c(known) + (1-\alpha) \cdot c(adjacent\ material)$, where α is the fraction of the pixel x covered by the known object. In the image reconstruction iterations, the approach becomes: 1) update the image according to equation (12), 2) replace all the pixel values of $X_m(\theta)$ with $c_m(x:\theta)$, and 3) make sure the pixel values of $X_e(\theta)$ are *at least* $c_m(x:\theta) = \alpha \cdot c(known)$. Therefore, we allow the values of the edge pixels to converge, rather than force them to some value that is unknown to us.

2.3.3. Objects at unknown pose, with partial volumes

Next, we examine the more realistic case, in which the object's pose is unknown to us. In this case, we use the same method described above to place the known object into the image, but we first must adapt the algorithm to perform a pose search.

3. POSE SEARCH

We assume that the known objects have a rigid structure; that is, each component has the same orientation with respect to the other components. In the future, we will add more degrees of freedom, but for now, we have four metal cylindrical rods, each of radius 6.35 mm, stored in an initial configuration (see Figure 1) of (clockwise from top): steel, aluminum, brass, and teflon. The rods are translated and rotated in this fixed configuraton. Note that the steel and brass rods are much more attenuating than the aluminum and teflon, which have attenuation coefficients on the same order as water, plastic, human tissue, and bone.

3.1. Finding the Next (or First) Pose Estimate

The problem is to find the location and orientation (pose) at which to place this known configuration of rods. One approach could be to: 1) form an image estimate using the above algorithm, 2) select a test pose and replace the rods into the new image, 3) forward project this image and form the estimate of $g(y)$ as in equation (1), 4) find the value of $I(d\|g)$, 5) repeat for all pose guesses and select the one which minimizes I-divergence.

This method is costly because the forward projection step, $\sum_x h(y|x)c(x:\theta)$, requires too much computation to implement for each pose guess. Therefore, we use the following technique:

- Form an image estimate, $\hat{c}_i^{(k)}(x)$. For finding the initial pose estimate (k=0), this image could be a filtered backprojection (with streaks), or an image formed from a few iterations of the alternating minimization (AM) algorithm (also with streaks, since we are not yet accounting for the known high-density areas). Either of these images will show approximately where the rods are located (which will provide us with a coarse estimate of the pose values). For subsequent pose estimates, $\hat{c}_i^{(k)}(x)$ is formed by equation (12).

- It can be shown (see O'Sullivan and Benac[9]) that we can rewrite the original objective function, $I(p\|q)$, as the following decoupled objective function:

$$\sum_i \sum_x \left[\tilde{b}_i^{(k)}(x) c_i(x) + \frac{1}{Z_i(x)} \hat{b}_i^{(k)}(x) \exp(Z_i(x)[\hat{c}_i^{(k)}(x) - c_i(x)]) \right], \tag{15}$$

  which can be minimized directly over $c_i(x)$ to form equation (12).

  We can, however, rewrite this so that we *maximize* the following objective function:

$$\sum_i \sum_x \left[\tilde{b}_i^{(k)}(x)[\hat{c}_i^{(k)}(x) - c_i(x:\theta)] - \frac{1}{Z_i(x)} \hat{b}_i^{(k)}(x) \exp(Z_i(x)[\hat{c}_i^{(k)}(x) - c_i(x:\theta)]) \right]. \tag{16}$$

  Now, by letting $c_i(x:\theta)$ be equal to $\hat{c}_i^{(k+1)}(x)$ as defined in (12) on $X_s(\theta)$, and replacing in the known values of the high-density objects (the metal rods) in $X_m(\theta)$ and $X_e(\theta)$ (the edge pixels) as described in Section 2.3.2, we choose the θ_{opt} which maximizes this objective function.

- We need an efficient method of selecting the subset of pose guesses (θ) over which we maximize the objective function. For finding the initial pose estimate, we use the coarse estimate provided by the initial image as the starting point, and in subsequent iterations, we use the previous pose estimate as the starting point.

  Our approach thus far has been to set up a 3-dimensional lattice, whose axes are the x and y coordinates of location and the rotation angle ϕ. Then, we test all points on the lattice at a fixed distance (e.g., 1 pixel and 1 degree) from the starting pose and move to the one with the highest value of the objective function. When a maximum is found, we reduce the step size by a factor of 10 and repeat. This has led to very accurate estimates of the pose (within 0.1 mm and 0.01 degrees).

  Undoubtedly, there are more efficient methods of searching over the pose domain, and we are currently exploring those as we further the research.

- The resulting image, $\hat{c}_i^{(k+1)}(x)$, is then $c_i(x:\theta_{opt})$.

4. RESULTS

In the experiments that follow, we use synthetic data, meaning that the data $d(y)$ were formed by operations on an image of a simulated plastic phantom containing 4 metal rods in a water bath. This phantom is seen in Figure 1a. The data are formed by a polyenergetic source, with Poisson noise and a constant scatter level. The images, $c(x)$, are displayed with a viewing window of $\pm 20\%$ the attenuation of water (-200 to 200 Hounsfield units (HU)), and have 1 mm $\times$ 1 mm pixels. Note that for the experiments described herein, we use only the attenuation spectrum ($\mu(E)$) of water (thus $I = 1$ in equation (2)) and we define $c(x)$ to be the ratio of attenuation coefficients of the object and water at the average energy, E_0.

When we reconstruct this image using the standard filtered backprojection (FBP) technique, we see pronounced streaking emanating from the rods (see Figure 1b). This image is in good agreement with the FBP image from an actual CT scan of a similar phantom, indicating that our synthetic data has accounted for many of the nonlinear physical phenomena present in real data.

4.1. AM Algorithm Using No Prior Knowledge

When we use the AM algorithm without accounting for our knowledge of the attenuation coefficients of the rods, we get an image that contains significant streaking (see Figure 1c), though it is much improved from the FBP algorithm and from our experiments with emission tomography (as discussed in Section 1.3–not shown).

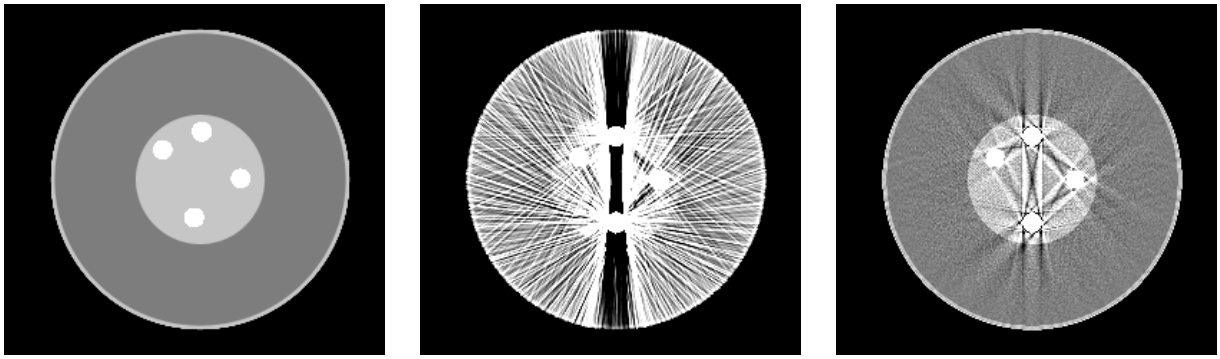

(a) (b) (c)

Figure 1. (a) The simulated phantom used in the experiments–a Lucite phantom in a water bath with four metal rods (clockwise from top): steel, aluminum, brass, and teflon. (b) The image formed from filtered backprojection. (c) The image formed with 500 iterations of the AM algorithm, using no prior knowledge of the metal objects.

4.2. AM Algorithm with Objects at a Known Pose

For this test, our "true" image had the four rods translated by 3 pixels to the left of and 8 pixels down from the origin, with zero degrees of rotation (pose = (-3, -8, 0)). Thus, every pixel in the image is either covered completely by the rods or it is not, so there is no partial volume effect. Thus, assuming the pose is known to us, we simply replace the rod pixels in every iteration with their true attenuation coefficients. The streak-free result is shown in Figure 2a.

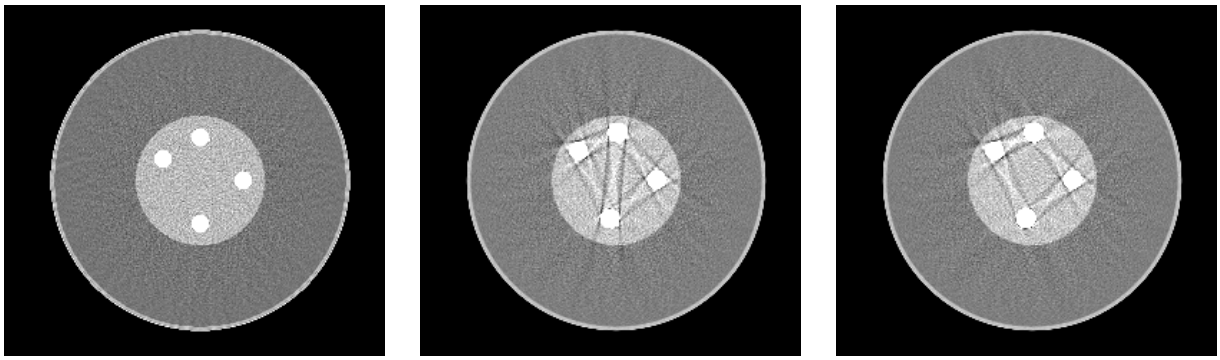

(a) (b) (c)

Figure 2. 500 iterations of the AM algorithm using full prior knowledge of the metal rods and the pose, with the true pose at: (a) (-3, -8, 0) (b) (-1.754, 3.328, 5.22). (c) 500 AM iterations with pose search incorporated, with the true pose at (-1.754, 3.328, 5.22) and the estimated pose at (-1.757, 3.333, 5.205).

4.3. AM Algorithm with Objects at a Known Pose (with Partial Volumes)

For this test, our "true" image had the four rods with a pose of (-1.754, 3.328, 5.22). Thus, some of the pixels may be only partially covered by the rods, so we use the replacement method outlined in Section 2.3.2. The resulting image, when we assume the pose is known, is shown in Figure 2b.

The fact that this image has streaks when compared to Figure 2a illustrates the importance of understanding and accounting for this so-called "partial volume effect".

4.4. AM Algorithm with Objects at a Unknown Pose (with Partial Volumes)

For this test, our "true" image had the four rods with a pose of (-1.754, 3.328, 5.22), so we would expect some partial volume effects. This time, however, the pose is unknown to us, so we concurrently estimate the pose and update the image, as discussed in Section 3. The resulting image is shown in Figure 2c, and the final pose estimate was (-1.757, 3.333, 5.205).

Note that this image is very comparable to the image in (b), the case in which we had full knowledge of the pose all along. This method of estimating the pose closely approximates the ideal situation of being given the exact coordinates *a priori*.

5. CONCLUSIONS

We have described a model for transmission tomography that uses an alternating minimization algorithm to reconstruct images from CT data. The model incorporates nonlinear physical effects such as noise, scatter, and spectrum hardening of polyenergetic radiation. In the presence of high-density objects, the images formed via this algorithm are much improved when compared to the standard technique of filtered backprojection, which assumes that attenuation is a linear function. The images are improved further, however, when we incorporate knowledge of the high-density objects into the algorithm. This algorithm can be modified to perform a search for the pose of these objects, and then we insert the attenuation values of the objects into their appropriate location in the image.

This initial investigation has demonstrated that there is still work left to be done on this problem. Even when we know precisely the pose of the objects, the resulting image still has some residual streaks from partial volume effects. Also, for real data (where we will not know the precise pose) we need to implement the version of the algorithm that incorporates the attenuation spectra of all the components in the image, such as water, brass, steel, and aluminum. This will result in more unknown variables, and will require some form of regularization.

Additionally, we are exploring alternative methods for performing the pose search. Currently, for two dimensional imaging, we are testing many points around the current estimate to find a subsequent pose estimate that improves the objective function. Other methods, such as one based on a gradient search, are potentially important as our research moves towards three dimensions and objects that have articulating parts. Our goal in this research is to improve the treatment of patients who have advanced cervical cancer and who are receiving a treatment protocol that uses brachytherapy applicators containing high density metals.

ACKNOWLEDGMENTS

This work was supported in part by Washington University and in part by the National Institutes of Health under research grant R01CA75371 from the National Cancer Institute of the National Institutes of Health, (J. F. Williamson, P. I.).

REFERENCES

1. W. A. Kalender, R. Hebel, and J. Ebersberger, "Reduction of CT artifacts caused by metallic implants," *Radiology* **164**, pp. 576–577, Aug. 1987.
2. G. H. Glover and N. J. Pelc, "An algorithm for the reduction of metal clip artifacts in CT reconstructions," *Med. Phys.* **8**, pp. 799–807, 1981.
3. B. P. Medoff, W. R. Brody, M. Nassi, and A. Macovski, "Iterative convolution backprojection algorithms for image reconstruction from limited data," *J. Opt. Soc. Am.* **73**(11), pp. 1493–1500, 1983.
4. B. E. Oppenheim, "Reconstruction tomography from incomplete projections," in *Reconstruction Tomography in Diagnostic Radiology and Nuclear Medicine*, M. Ter-Pogossian, ed., pp. 155–183, Baltimore, MD: University Park, 1977.
5. G. Wang, D. L. Snyder, J. A. O'Sullivan, and M. W. Vannier, "Iterative deblurring for CT metal artifact reduction," *IEEE Trans. Med. Imag.* **15**, pp. 657–664, Oct. 1996.

6. D. L. Snyder, J. A. O'Sullivan, B. R. Whiting, R. J. Murphy, J. Benac, J. A. Cataldo, D. G. Politte, and J. F. Williamson, "Deblurring subject to nonnegativity constraints when known functions are present, with application to object-constrained computerized tomography," *IEEE Trans. Med. Imag.* **20**, pp. 1009–1017, Oct. 2001.
7. K. Lange and R. Carson, "EM reconstruction algorithms for emission and transmission tomography," *J. Comp. Assisted Tomo.* **8**, pp. 306–316, Apr. 1984.
8. J. A. O'Sullivan, "Alternating minimization algorithms: from Blahut-Arimoto to expectation-maximization," in *Codes, Curves, and Signals–Common Threads in Communications*, A. Vardy, ed., pp. 173–192, Norwell, MA: Kluwer Academic, 1998.
9. J. A. O'Sullivan and J. Benac, "Alternating minimization algorithms for transmission tomography," *submitted to IEEE Trans. Med. Imag.* , Jun. 2001.