

A sieve-regularized image reconstruction algorithm with pose search in transmission tomography

Ryan J. Murphy^a, Donald L. Snyder^a, David G. Politte^b, Joseph A. O’Sullivan^a

^aElectronic Systems & Signals Research Laboratory, Washington University
One Brookings Drive, St. Louis, MO 63130 USA

^bMallinckrodt Institute of Radiology, Washington University, St. Louis, MO 63110 USA

ABSTRACT

We have developed a model for transmission tomography that views the detected data as being Poisson-distributed photon counts. From this model, we derive an alternating minimization (AM) algorithm for the purpose of image reconstruction. This algorithm, which seeks to minimize an objective function (the I-divergence between the measured data and the estimated data), is particularly useful when high-density objects are present in soft tissue and standard image reconstruction algorithms fail. The approach incorporates inequality constraints on the pixel values and seeks to exploit known information about the high-density objects or other priors on the data. Because of the ill-posed nature of this problem, however, the noise and streaking artifacts in the images are not completely mitigated, even under the most ideal conditions, and some form of regularization is required. We describe a sieve-based approach, which constrains the image estimate to reside in a subset of the image space in which all images have been smoothed with a Gaussian kernel. The kernel is spatially varying and does not smooth across known boundaries in the image. Preliminary results show effective reduction of the noise and streak artifacts, but indicate that more work is needed to suppress edge overshoots.

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

1. INTRODUCTION

Computerized tomographic (CT) image reconstruction of soft tissue containing high-density objects has been studied extensively. More recently, a similar application has been to use CT for nondestructive testing of materials or fabricated components, such as automotive or aeronautical components and ceramics. It is well-known that high-density objects embedded in lower density regions cause severe streaking artifacts in the resulting images when conventional algorithms such as filtered back projection (FBP) are used.

FBP, which produces very useful images in most clinical applications, is based on a linear model that assumes the set of CT measurements is the Radon transform of the linear attenuation coefficients of the object being scanned. When the object’s components are similar in attenuation value, this assumption is quite harmless. In the presence of high-density objects (such as brachytherapy applicators, hip prostheses, and dental fillings), however, the nonlinear effects of noisy data, spectral hardening, and scatter (among others) become more pronounced and the assumption fails, resulting in images littered with fine streaks and heavy dark bands connecting the high-density regions. Another example of streaking is the presence of edge overshoots, which appear as alternating light and dark rings around the perimeters of objects. This problem of metal artifact reduction (MAR) has been studied for a long while.^{1–7}

The goal of our study is to develop a more rigorous physical model for transmission tomography that accounts for the nonlinear effects, and then to use this model, along with any prior knowledge of the object being scanned, to form an image reconstruction algorithm. Specifically, we aim to exploit knowledge of the high-density components—the size, composition, orientation, and attenuation characteristics—to improve the initial conditions of the iterative image formation process.

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

This work builds from a 1984 paper by Lange and Carson,⁸ in which they derive an expectation-maximization (EM) algorithm for transmission tomography to form the maximum-likelihood image estimate using a suitable model. Our approach restates the maximum likelihood estimation as a double minimization of an I-divergence, or negative Poisson likelihood, and we derived an alternating minimization (AM) algorithm to form the image.^{9,10} The first minimization enforces a constraint involving the measured CT data, and the second minimization seeks to minimize the I-divergence between the data and our estimate of the data. The image formation algorithm is presented in the following sections.

2. THE DATA MODEL

In transmission tomography, photons (of varying energies) are transmitted from a collimated source through an object, and the nonabsorbed photons (or their energies) are then collected by a detector array. At this phase of the research, we model the detectors as photon-counting devices, and therefore assume that the measured data (labeled $d(y)$) are Poisson-distributed, with means defined by

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

We are trying to estimate the 512×512 image of energy-dependent attenuation coefficients, $\mu(x, E)$, whose pixels are indexed by the set X . The source-detector pairs are indexed by the set Y , and $h(y|x)$ is a nonnegative kernel (determined by the scanner geometry), viewed as the average path length through pixel x on the path defined by source-detector pair y . Thus, the sum in the exponent (i.e., the forward projection) is a discrete approximation of the line integral of attenuation along each path y , and the exponential term can be thought of as the survival probability of the photons at energy E traveling on that path. I_0 is the mean number of incident photons from the source over a spectrum of energies in the absence of attenuating objects. The final term, $\beta(y)$, is the mean of the background events and allows us to model the photon scatter.

The image of attenuation coefficients is modeled as:

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

where $\mu_i(E)$ is the attenuation spectrum (in mm^{-1}) for the i th constituent (e.g. constituents could be water, bone, metal, composites) and $c_i(x)$ is a unitless function (e.g., specific gravity) at pixel x for the i th constituent. For our simulations, we use only the attenuation spectrum for water (i.e., $I = 1$), so we are estimating the image $c(x)$, which is the ratio of the pixel attenuation coefficient (at mean energy E_0) to that of water.

3. ALTERNATING MINIMIZATION (AM)

To form the image estimate, we seek to minimize the I-divergence (equivalent to maximizing likelihood) between the measured data, $d(y)$, and the estimated data means, $g(y)$, where the I-divergence is a discrepancy measure defined by

$$I(d||g) = \sum_{y \in Y} d(y) \ln \frac{d(y)}{g(y)} - d(y) + g(y). \quad (3)$$

O'Sullivan⁹ shows how to rewrite this as a double minimization:

$$\min_q \min_p I(p||q), \quad (4)$$

where p belongs to a linear family that enforces data constraints, and q belongs to an exponential family that makes q a mean of the data for a given energy:

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

$$\mathcal{E}(I_0, h) \equiv \left\{ \mathbf{q} : q(y, E) = I_0(y, E) \exp \left(- \sum_{x \in X} h(y|x) \mu(x, E) \right) \right\}. \quad (6)$$

The resulting AM algorithm alternates between finding p and q , and the final result is shown in O’Sullivan and Benac¹⁰ to be as follows:

- Choose an initial guess of attenuation coefficients, $\hat{c}_i^{(0)}(x)$.
- Using the k th estimate of attenuation coefficients, compute:

$$\hat{q}^{(k)}(y, E) = I_0(y, E) \exp \left(- \sum_{x \in X} \sum_{i=1}^I \mu_i(E) h(y|x) \hat{c}_i^{(k)}(x) \right). \quad (7)$$

- Using the k th estimate of q , compute:

$$\hat{p}^{(k)}(y, E) = \hat{q}^{(k)}(y, E) \frac{d(y)}{\sum_{E'} \hat{q}^{(k)}(y, E') + \beta(y)}. \quad (8)$$

- Compute the *back projections* of the estimates:

$$\tilde{b}_i^{(k)}(x) = \sum_{y \in Y} \sum_E \mu_i(E) h(y|x) \hat{p}^{(k)}(y, E), \quad (9)$$

$$\hat{b}_i^{(k)}(x) = \sum_{y \in Y} \sum_E \mu_i(E) h(y|x) \hat{q}^{(k)}(y, E). \quad (10)$$

- Form estimate $k + 1$ of the attenuation coefficients:

$$\hat{c}_i^{(k+1)}(x) = \hat{c}_i^{(k)}(x) - \frac{1}{Z_i(x)} \ln \left(\frac{\tilde{b}_i^{(k)}(x)}{\hat{b}_i^{(k)}(x)} \right), \quad (11)$$

where $Z_i(x)$ is an appropriate scaling function.

Some results from this algorithm are shown in Figure 1. This was a simplified experiment of synthetically “scanning” a phantom, in the presence of noise and a monoenergetic spectrum, of a water bath with a plastic core containing four rods of varying densities: steel (12:00 position), aluminum (3:00), brass (6:00), and Teflon (10:30). Using this simulated data, we reconstructed the image, and the figure shows that convergence is a very slow process. The presence of the high-density rods (steel and brass) leads to significant streaks connecting the rods in the early iterations, and as these diminish many iterations later, the noise effects become more pronounced in the form of fine starburst patterns. The next section examines a technique of softening the earlier, thicker artifacts, and the following section focuses on the later, thin streaks.

4. ALTERNATING MINIMIZATION WITH POSE SEARCH (AMPS)

The use of the phantom above gives us a simpler model for two distinct applications for this algorithm—1) nondestructive testing (NDT) of materials and components and 2) brachytherapy treatment for advanced cervical cancer. In brachytherapy, the physician implants an intracavitary high-density applicator (containing radioactive sources) adjacent to the tumor. In each of these examples, we will likely know the geometry, composition, and attenuation characteristics of some of the objects present (e.g., the brachytherapy applicator, metal fasteners, rods) in the CT scan, but we would not know the location and orientation (collectively called the pose) of the high-density object(s) in relation to the overall object being scanned.

We seek to adapt the above AM algorithm to search for the correct pose of these objects, set the pixels at this pose to their known values, and then reconstruct the rest of the image as before. The net effect of this is an improvement of the “initial guess” in the AM algorithm, causing less severe errors in the early iterations and thus lessening the thick streaks stemming from these errors.




Figure 1. (a) The phantom: Lucite outer shell and core, water bath, and four rods (clockwise from top): steel, aluminum, brass, and Teflon. With noisy, synthetic data, the image is reconstructed: (b) using filtered back projection; using the AM algorithm at approximately (c) 440 iterations and (d) 220,000 iterations.

4.1. Objects at “known” pose

Let us assume for the moment that the pose is “known”, either by some prior information or by using the estimate of the pose at the k th iteration as the “known”, but temporary, value. In two dimensions, the pose, θ , has three coordinates—two for position (x, y) and one for orientation (ϕ) . A reference image, $c_a(x : \mathbf{0})$, is stored in memory, with the known objects centered at the reference pose, $\mathbf{0}$. This reference image is translated and rotated by θ , creating a known function, $c_a(x : \theta)$, that is zero everywhere, except for the relatively few pixels containing the known objects. The core of our problem statement is outlined by the following constraint: *every pixel in a candidate image for iteration k must be greater than or equal to $c_a(x : \hat{\theta}^{(k)})$, where $\hat{\theta}^{(k)}$ is the k th estimate of the pose.*

This not only enforces a nonnegativity constraint on the image (since $c_a(x : \theta)$ is always nonnegative) but it also allows us to set the pixels that are entirely covered to their exact value. For the pixels that are only partially covered (by a fraction α) by the known function, we model the pixel value to be a linear combination of the components of the pixel: $\alpha \cdot c_{known}(x) + (1 - \alpha) \cdot c_{unknown}(x)$. Since we define $c_a(x : \theta)$ to be equal to

$\alpha \cdot c_{known}(x)$, the inequality constraint is logical for this case as well.

The AM algorithm is modified as follows:

- Define $c^{AM}(x)$ to be the expression for the $(k + 1)$ -st image estimate if we had no pose search:

$$c^{AM}(x) = \hat{c}^{(k)}(x) - \frac{1}{Z(x)} \ln \frac{\tilde{b}^{(k)}(x)}{\hat{b}^{(k)}(x)}. \quad (12)$$

- For the pose θ , the Kuhn-Tucker conditions yield

$$\hat{c}^{(k+1)}(x : \theta) = \begin{cases} \max[c^{AM}(x), c_a(x : \theta)], & \alpha \neq 1 \\ c_a(x : \theta), & \alpha = 1 \end{cases}. \quad (13)$$

4.2. Estimating the pose

There are assuredly many possible methods for choosing candidate pose values and testing them against one another, but that will not be the focus here. We simply choose test points for the pose in the three-dimensional lattice surrounding the current pose, test the I-divergence using that test point, and select the best one. This fits very well with the above algorithm by adding the following additional steps:

- Select a candidate pose value that is “close” (in the lattice) to the current one. Form $\hat{c}^{(k+1)}(x : \theta)$.
- Forward project $\hat{c}^{(k+1)}(x : \theta)$ and form $g(y)$ using equation (1), and then find $I(d||g)$.
- After trying a suitable number of candidate θ values in a lattice of smaller and smaller stepsizes surrounding the current pose, select the θ and the corresponding $\hat{c}^{(k+1)}(x : \theta)$ that yielded the smallest I-divergence; this is $\hat{\theta}^{(k+1)}$ and $\hat{c}^{(k+1)}(x)$, respectively.

It should be noted here that as the pose moves around from one iteration to the next, the pixels that were previously covered by the known high-density objects (but are no longer) can cause some problems. Their values are much higher than they should be, and it takes many iterations for the algorithm to bring these values back down to the low-density levels, and streaks often form in the interim. For this reason, we “reset the initial condition” on these pixels whenever the high-density areas in the image move elsewhere, and set them to plastic or water or whatever is the dominant substance in the region of interest. The pixel is then treated like every other in subsequent iterations.

We have been able to show that our best estimate of the pose is very accurate—within 0.05 mm of the actual pose. We will discuss this more in the next section, but for now, let us look at the best possible scenario, where we get the pose exactly right at each iteration. The result, shown in Figure 2, has dramatically reduced the heavy streaks connecting the rods that were seen in Figure 1(b) (after a similar number of iterations), but they are still present, along with the speckle noise. (Note: the noise is more pronounced because the underlying data here contains higher noise levels than the data used to create Figure 1.) The continuing presence of the streaks and the noise indicates the need for regularization in the image formation process.

5. REGULARIZATION WITH SIEVES

It has become apparent that we need some form of regularization included in our derivation of the AM algorithm. Even in the ideal case of Figure 2, where we impractically used the exact rod locations and attenuation values, there is streaking evident connecting the higher-density objects. The amount of regularization needed is a tradeoff between lower resolution in the reconstructed image and longer computation times versus the severity of undesirable effects in the image caused by noise and the ill-posed nature of the inverse problem.




Figure 2. The result of the AMPS algorithm with the pose fixed at the true pose (the ideal scenario) at each iteration. Noisy, polyenergetic, synthetic data, 500 iterations.

One approach to regularization is the method of sieves, from the 1987 paper on noise and edge artifacts by Snyder, Politte, et al.¹¹ We constrain the estimate of the image, $\lambda(x)$, to reside in the sieve; i.e., the estimate must be in a kernel sieve S defined by:

$$S = \left\{ \lambda : \lambda(x) = \int s(x|z)c(z)dz \right\}, \quad (14)$$

where $c(x)$ would be the output of our AM algorithm. We model the forward projection of our estimate as:

$$m(y) = \int h(y|x)\lambda(x)dx = \int \int h(y|x)s(x|z)dx c(z)dz, \quad (15)$$

where the new forward projection kernel can be thought of as $\int h(y|x)s(x|z)dx$. For our purposes, $s(x|z)$ is a spatially-varying Gaussian smoothing function that we convolve with the image during each projection (forward or back) in the AM algorithm. It is spatially-varying in the sense that no smoothing is performed over the edges or inside of the known objects. See Figure 3 for the result of the AMPS algorithm with the sieves included.

This figure shows some very promising results, coupled with some discouraging ones. The streaks are nearly gone and the speckle noise has been reduced (meeting the original goal of this paper), but the edge overshoots are quite prominent at every edge in the image. These overshoots are predicted in the Snyder paper¹¹ because of the lack of a resolution kernel needed for edge artifact suppression. However, the spatially varying nature of the sieve kernel and a similar constraint on the resolution kernel makes this problem difficult. We are currently pursuing a path to finding this kernel or finding other avenues to regularization.

6. CONCLUSIONS

We described a model for transmission tomography that incorporates nonlinear physical effects such as noise, scatter, and spectral hardening of polyenergetic radiation. From this model, we derived an alternating minimization algorithm to reconstruct images from CT data. When high-density objects are present in soft tissue, the images formed with the AM algorithm are improved (compared to filtered back projection). These images are improved further and converge faster when we incorporate prior knowledge (of the high-density objects, for example) into the algorithm. There are still streaks remaining, however, which are almost completely eliminated when we implement a kernel-sieve method to regularization.




Figure 3. The results from the AMPS algorithm with sieve regularization, starting with the initial pose at $(-1$ mm, $+1$ mm, $0^\circ)$ from the true pose and an initial image of zeros (air). Noisy, polyenergetic, synthetic data, 500 iterations. Final pose: $(0.009$ mm, -0.01 mm, $-0.02^\circ)$ from the true pose.

The presence of edge overshoots with the kernel-sieve approach has revealed the need to continue the search for a better regularization scheme. The lack of streaking in the current image, however, allows us to proceed with confidence that our goal of improving the quality of brachytherapy imaging is attainable.

ACKNOWLEDGMENTS

This work was supported in part by the Graduate Education and Research Partnership with the Boeing Foundation, grant 34566C, and in part by the National Institutes of Health under research grant R01CA75371 from the National Cancer Institute of the NIH, (J. F. Williamson, P. I.).

REFERENCES

1. 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.
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. W. A. Kalender, R. Hebel, and J. Ebersberger, "Reduction of CT artifacts caused by metallic implants," *Radiology* **164**, pp. 576–577, Aug. 1987.
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. B. DeMan, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens, "An iterative maximum-likelihood polychromatic algorithm for CT," *IEEE Trans. Med. Imag.* **20**, pp. 999–1008, Oct. 2001.

8. K. Lange and R. Carson, "EM reconstruction algorithms for emission and transmission tomography," *J. Comp. Assisted Tomo.* **8**, pp. 306–316, Apr. 1984.
9. 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.
10. J. A. O'Sullivan and J. Benac, "Alternating minimization algorithms for transmission tomography," *submitted to IEEE Trans. Med. Imag.*, Jun. 2001.
11. D. L. Snyder, M. I. Miller, L. J. Thomas, Jr., and D. G. Politte, "Noise and edge artifacts in maximum-likelihood reconstructions for emission tomography," *IEEE Trans. Med. Imag.* **6**, pp. 228–238, Sep. 1987.