Inverse Problems in Science and Engineering
Vol. 14, No. 8, December 2006, 873–885

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

Abel inversion using total variation regularization: applications

THOMAS J. ASAKI*†, PATRICK R. CAMPBELL‡, RICK CHARTRAND‡, COLLIN E. POWELL†, KEVIN R. VIXIE‡ and BRENDT E. WOHLBERG‡

†Continuum Dynamics (CCS-2), MS D413, Los Alamos National Laboratory, Los Alamos, NM 87545, USA
‡Mathematical Modeling and Analysis (T-7), MS B284, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

(*Received 14 April 2005; in final form 14 June 2006*)

We apply total-variation (TV) regularization methods to Abel inversion tomography. Inversions are performed using the fixed-point iteration method and the regularization parameter is chosen such that the resulting data fidelity approximates the known or estimated statistical character of the noisy data. Five one-dimensional (1D) examples illustrate the favorable characteristics of TV-regularized solutions: noise suppression and density discontinuity preservation. Experimental and simulated examples from X-ray radiography also illustrate limitations due to a linear projection approximation. TV-regularized inversions are shown to be superior to squared gradient (Tikhonov) regularized inversions for objects with density discontinuities. We also introduce an adaptive TV method that utilizes a modified discrete gradient operator acting only apart from data-determined density discontinuities. This method provides improved density level preservation relative to the basic TV method.

Keywords: Abel transform; Inverse problems; Total variation; Tomography; Regularization

2000 Mathematics Subject Classifications: 45E10; 15A29; 49M37; 65F22; 78A46

1. Introduction

The density reconstruction of objects from several radiographic views is a classic and important tomography problem. A large subclass of problems is the interrogation of manufactured items that consist of a small number of different materials. We present here some results of our investigations on the application of total variation (TV) regularization to the object reconstruction process. Our results show that the choice of regularization can have significant impact on the interpretation of

*Corresponding author. Tel.: 505-667-5787. Fax: 505-665-4479. Email: asaki@lanl.gov

Inverse Problems in Science and Engineering
ISSN 1741-5977 print/ISSN 1741-5985 online © 2006 Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/17415970600882549

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

radiographs. This can be especially important in manufacturing, homeland security, and high-energy particle beam applications.

The general inversion is typically formulated using the Radon Transform or related approaches [3] and is regularized using any number of techniques. Vogel [2] presents a good introduction to regularization of inverse problems. For objects with cylindrical symmetry, tomographic applications require only a single viewing angle and the Radon Transform reduces to the Abel Transform. In this study we consider objects of one-dimensional (1D) description $\rho(r)$. We leave discussions of applications to objects of general cylindrical symmetry $\rho(r, z)$ to a future article. The continuous Abel transform is

$$\mathrm{d}(x) = 2\int_{|x|}^{\infty} \frac{r\rho(r)}{\sqrt{r^2 - x^2}}\,\mathrm{d}r, \tag{1}$$

where d is a line-integral density relative to ρ. For example, if ρ is a volumetric density then d is an areal density. Equation (1) has a well-defined inverse:

$$\rho(r) = -\frac{1}{\pi r}\frac{\mathrm{d}}{\mathrm{d}r}\int_{r}^{\infty} \frac{x\mathrm{d}(x)}{\sqrt{x^2 - r^2}}\,\mathrm{d}x. \tag{2}$$

A discrete version of equation (2) is the basis for Abel inversion tomography.

In practice, there are a number of difficulties to address. First, equation (1) is a simplified description of typically very nonlinear experiments. Second, the inverse problem can be ill-posed. Since radiographs are transmission intensity maps (or some equivalent), the corresponding intrinsic material property is an attenuation coefficient. Thus, obtaining an object density requires an additional transformation either in the radiographic space (intensity to areal density) or in the object description (attenuation to density). Both approaches are nontrivial requiring a detailed understanding of the physics of interaction between particles in the interrogating beam and the object materials. The simplest case – reconstructing an object made of a single material – can usually be solved with good accuracy. However, multiple-material situations cause this additional transformation to be nonunique, and significant prior information about the object is necessary before the inverse Abel transform can be utilized for quantitative evaluation.

A third difficulty is the ill-conditioning of the discrete inverse Abel transform. While a given inversion is unique, small perturbations in d lead to large deviations in ρ. This is because equation (2) defines an unbounded operator (see [1] for details). The inversion must be regularized to obtain meaningful results from noisy data.

The discrete Abel transform can be formulated as a matrix P. If we consider the object radial density values as a vector ρ of n elements and the areal density projections as a vector d of m elements, then P is a nonsparse $m \times n$ matrix. This projection is invertible if $n \leq m$, but is poorly conditioned when $n \sim m$. Condition numbers increase linearly in n when $n \gg 1$. For example, a square matrix operator of size $n \times n$ has condition number $\approx 1.75n$.

Inverse Abel transform tomography is formulated as a functional minimization problem:

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

$$\min_{\rho} F(\rho) = \min_{\rho}\left\{\|P\rho - d\|_{df} + \alpha R(\rho)\right\}, \tag{3}$$

where $\|\cdot\|_{df}$ is an appropriate data fidelity norm, and $R(\rho)$ is a regularization term determined by a probability model of the types of objects we expect.

In this article, we focus on the use of regularization to address the important, but ill-conditioned, inverse Abel Transform tomography problem. We demonstrate, by example, the importance of the choice of R for recovering object density profiles ρ from noisy areal density data d. In particular, we show that the use of TV regularization has advantages over Tikhonov regularization in preserving material property discontinuities in reconstructed objects.

2. Methods

We compute P as the parallel planar projection from a 2D object space onto a 1D data space. Noise is treated as stationary Gaussian white noise. In practice, this approach works well, but typical experiments are dominated by signal-dependent Poisson noise. The examples that follow have different noise characteristics, so our treatment of data fidelity can be expected to have mixed results. We consider functionals of the form

$$F(\rho) = \frac{1}{2}\int_0^M \mathrm{d}r\,|P\rho - \mathrm{d}|^2 + 2\pi\alpha\int_0^M r\,\mathrm{d}r\,|D\rho|^p. \tag{4}$$

where ρ is supported on $[0, M]$. In particular, we consider $p \in \{0, 1, 2\}$, corresponding to the following regularization types: none, TV, and H^1, respectively. The minimizing solution of equation (3) depends on the choices of α and p. We select α that leads to a solution with data fidelity norm equal to the known or estimated variance in the data noise. Many problems may also benefit from a more careful approach to data fidelity modeling that assumes correct statistics. For example, it is clear that most radiography applications are governed by Poisson statistics. Such treatments are outside the scope of this study.

Our approach is to use the largest invertible discrete linear projection operator P and regularize the inversion. Given a data vector of length m, we reconstruct the object at the same resolution. Thus, P is an $m \times m$ matrix. The methods outlined in the following sections are not limited to this invertibility condition.

2.1. *Inversion without regularization*

The unregularized inversion is the $p = 0$ case of equation (4), as the second term of the right hand side does not depend on ρ. This lack of regularization is appropriate for situations in which no prior object knowledge is available. The object density reconstruction is given explicitly by the pseudo-inverse $\rho = (P^T P)^{-1} P^T d$. It is expected to produce poor results in real scenarios due to the combination of noisy data and ill-conditioning. In particular, we expect noise in the most stable directions of P (the most unstable of P^{-1}) to be amplified significantly and lead to poor inversions.

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

2.2. *H^1-regularization*

The $p=2$ case of equation (4) is the H^1-regularization minimization. It can be shown to be equivalent to applying the diffusion operator on the unregularized solution. The equivalent diffusion time is inversely related to α. The result is a smooth reconstruction clearly biased against discontinuities in ρ. The same is true for all $p > 1$ solutions. The object density reconstruction is given explicitly by

$$\rho = \left(P^T P - 4\pi\alpha r D^2 - 4\pi\alpha D\right)^{-1} P^T d. \tag{5}$$

For the gradient operator matrix, D, we use a simple forward differencing with Neumann boundary conditions. If the inverse is nearly singular, it is advantageous to use the solution methods outlined in the next subsection.

2.3. *TV regularization*

The $p=1$ case of equation (4) is the TV regularization minimization. This regularization is not biased against density discontinuities; it penalizes such density edges by a reduction in amplitude, but not at the expense of smoothing the edge. Following Vogel [4], we use the lagged-diffusivity fixed point method to find the minimum of $F(\rho)$. We compute the gradient and an approximate Hessian of the discrete cost functional F that define a quasi-Newton step towards the minimum. Taking the Gateaux derivatives of equation (4), we obtain

$$F'(\rho) = P^T(P\rho - d) + \alpha L(\rho)\rho, \tag{6}$$

and

$$F''(\rho) = P^T P + \alpha L(\rho) + \alpha L'(\rho)\rho, \tag{7}$$

where we adopt the notation $R'(\rho) = L(\rho)\rho$. In particular, if $p=2$ then $L(\rho) = 2\pi\left(rD^2 + D\right)$. If $p=1$, the nondifferentiability of R when $D\rho = 0$ is problematic. The difficulty is to overcome by choosing a suitably small parameter β and defining R by

$$R(\rho) = 2\pi \int_0^\infty r\,\mathrm{d}r\sqrt{|D\rho|^2 + \beta^2}. \tag{8}$$

Now R has a well-defined derivative from which $L(\rho)$ is identified:

$$L(\rho) = 2\pi\left[(rD + I) \circ \left(\frac{D}{\sqrt{|D\rho|^2 + \beta^2}}\right)\right]. \tag{9}$$

Since $L(\rho)$ depends explicitly on the solution ρ, we let the solution iterate $\rho_{\nu+1}$ depend on the previous solution ρ_ν. This is the origin of the term lagged-diffusivity. The quasi-Newton iteration is given by

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

$$\begin{aligned}\rho_{\nu+1} &= \rho_\nu + s_\nu \\ &= \rho_\nu - [\mathrm{Hess}\, F(\rho)]^{-1} DF(\rho_\nu) \\ &\approx \rho_\nu - \left[P^T P + \alpha L(\rho_\nu)\right]^{-1}\left[P^T(P\rho_\nu - d) + \alpha L(\rho_\nu)\rho_\nu\right]. \end{aligned} \tag{10}$$

where in the correction term s_ν we use the approximate Hessian which omits the $L'(\rho)\rho$ term. While Newton methods are expected to have quadratic convergence, the use of a lagged-diffusivity and approximate Hessian guarantee only linear convergence [2]. Methods with faster convergence (for example, primal dual) require larger computational overhead. For the small problems of interest here ($n < 10^3$), linear convergence is sufficient. Typically $\sim$20 iterations are sufficient to reach our convergence criterion $\|s_\nu\|_2/\|\rho_\nu\|_2 < 10^{-3}$.

The discretization of equation (10) Newton is straightforward. First, we define the diffusivity function

$$\phi = \frac{1}{\sqrt{|D\rho|^2 + \beta^2}}, \tag{11}$$

so that

$$L(\rho) = \phi D + rD \cdot \phi D. \tag{12}$$

We then use simple forward differencing to construct the discrete gradient operator. The use of second-order methods can present difficulties when attempting to recover discontinuities. See [1] for a detailed discussion of the functional F when $p = 1$.

2.4. *Adaptive TV regularization*

Since most static objects have only piecewise continuous densities, we would like to use a regularization with the edge-preserving quality of TV and absolute amplitude smoothing quality of H^1. To this end, we present a fourth regularization method which we designate as adaptive total variation (aTV). In this approach, TV-regularized solutions are used to identify object radial locations of suspected density discontinuities. Then, an H^1 regularized inversion on the data set apart from these locations completes the analysis.

To accomplish this we mask the discrete gradient operator according to an edge location set that identifies suspected object discontinuities. This edge location set E is represented by a diagonal matrix of diagonal entry 0 if the object has a suspected density discontinuity at the location and entry 1 otherwise. The masked gradient operator is then $D^* = D \cdot E$. The aTV solution is found by the following algorithm:

(i) Determine α_{TV}, the regularization parameter that provides the desired TV reconstruction, and $\Delta\rho = \rho_{\max} - \rho_{\min}$, the density range of the TV reconstruction.
(ii) Set E equal to the identity matrix so that $D^* = D \cdot E = D$.
(iii) Set the discontinuity threshold $t = \Delta\rho$.
(iv) Set $\alpha > \alpha_{\mathrm{TV}}$.
(v) Compute the TV regularized solution ρ using α and D^*.

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

(vi) Compute the masked gradient operator matrix D^*. First set $E = I$, then set $E_{i,i} = 0$ if $\rho_{i+1} - \rho_i > t$. Now $D^* = D \cdot E$.
(vii) Update t.
(viii) Repeat steps 5–7 until convergence criteria are met on either the solution ρ or on the threshold t.
(ix) Compute the H^1 regularized solution ρ using a very large α and the final masked gradient operator matrix D^*.

The choice of t within the inner iteration is somewhat of an art, though certain principles apply. The threshold t is initially large relative to expected density discontinuities and lowered at each iteration down to a value somewhat larger than the estimated or known noise level η in the reconstruction ρ. The regularization parameter α is set large relative to α_{TV}. The large initial parameter values allow only the largest density discontinuities to be recognized at early iterations. Incremental adjustments combined with a masked gradient operator allow discontinuities on smaller and smaller size scales to be identified. The examples in the following section use the following parameters.

$$\alpha = 2\alpha_{\mathrm{TV}}, \tag{13}$$

and

$$t_k = 2[\eta + 2^{-k}(\Delta\rho - 2\eta)], \quad k = 1, 2, \ldots, 10, \tag{14}$$

where k is the aTV iteration index.

3. Examples and discussion

We show results of the four regularized inversion methods applied to both synthetic and real data. In all cases we work with objects of 1D description $\rho(r)$ and corresponding 1D data $d(x)$.

3.1. *Synthetic Example 1*

The first example is the reconstruction of an object of 10 nested varying density rings. The object, data, and reconstructions are shown in figure 1. The object density profile (a) is given by $n = 200$ ring densities. The corresponding projection data (b) is a synthetically generated areal density with added Gaussian noise. The variance of the noise is a uniform 1.5% of the maximum noiseless data value. The noisy data is shown in black and the noiseless data is shown in gray. In each of the remaining subfigures, the actual object is shown in gray and a reconstruction is shown in black: (c) unregularized; (d) H^1; (e) TV; and (f) aTV. As expected, the unregularized reconstruction is overwhelmed by noise amplification through the ill-conditioned inverse projection [4]. A smoothing regularizer does very well at reducing high frequency noise artifacts and even preserves the general character of the object, but the same regularizer is unable to capture density discontinuities, smoothing the edges. The TV regularizer captures many of the discontinuities. The aTV regularization procedure does the best of all. It best identifies all discontinuities and best preserves the actual density levels to within that given by the local

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

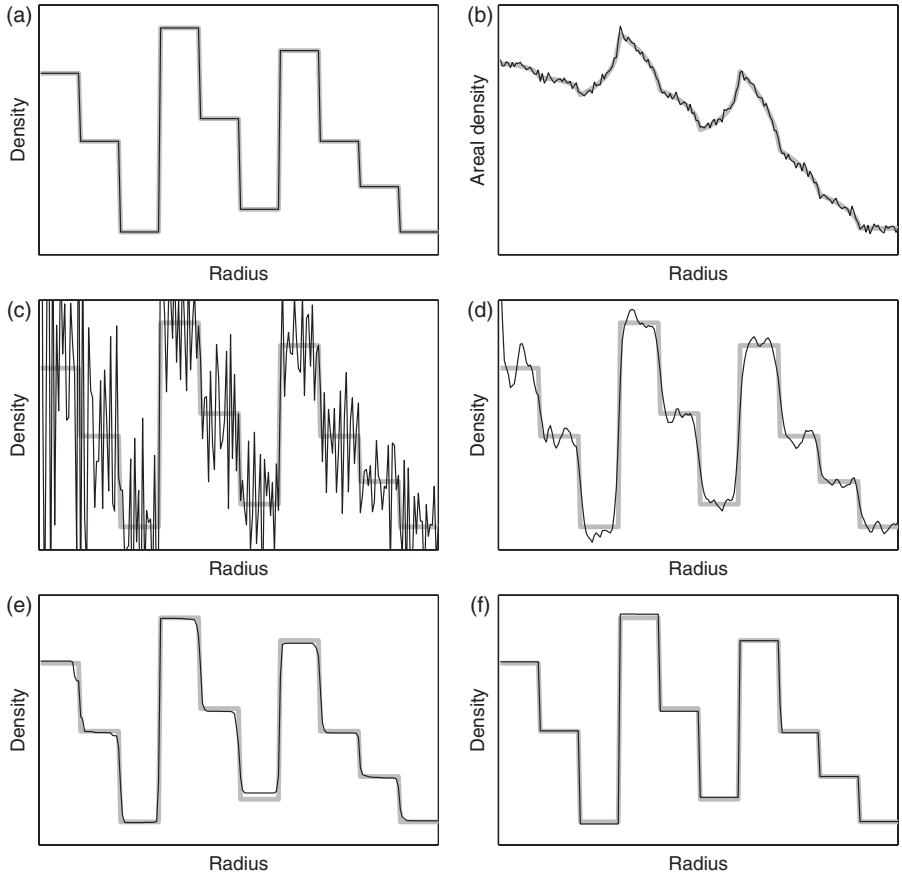

Figure 1. Regularized inversion examples on a synthetic data set generated from a simple nested-ring computational test object. The object density is defined at 200 radial positions and consists of 10 individual density values. The density profile of the object (a) is projected onto one-dimension (b, gray) and Gaussian noise is added to obtain a synthetic 1D radiograph (b, black). The remaining figures are four regularized object density reconstruction examples (black) against the object reference (gray): (c) unregularized; (d) H^1; (e) TV; and (f) aTV.

data. Figure 2 shows the results of several stages in the iterative aTV process along the path from figure 1(e)–(f). Figure 2(a) shows the original TV-regularized reconstruction, and is identical to figure 1(e). Figures (b)–(g) show intermediate reconstructions along with the current catalog of edge locations (vertical dashed lines). At each iteration, as the edge location detection threshold is lowered, the edge location set and gradient operator matrix are updated. Only the first six aTV iteration results are shown, as the remaining computations do not affect the edge location set in this example. Figure (h) shows the results of the final H^1 regularized solution; it is the result shown in figure 1(f).

3.2. *Synthetic Example 2*

The second example is the reconstruction of an object with a mixture of piecewise smooth and piecewise constant density variations. The object, data, and reconstructions

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

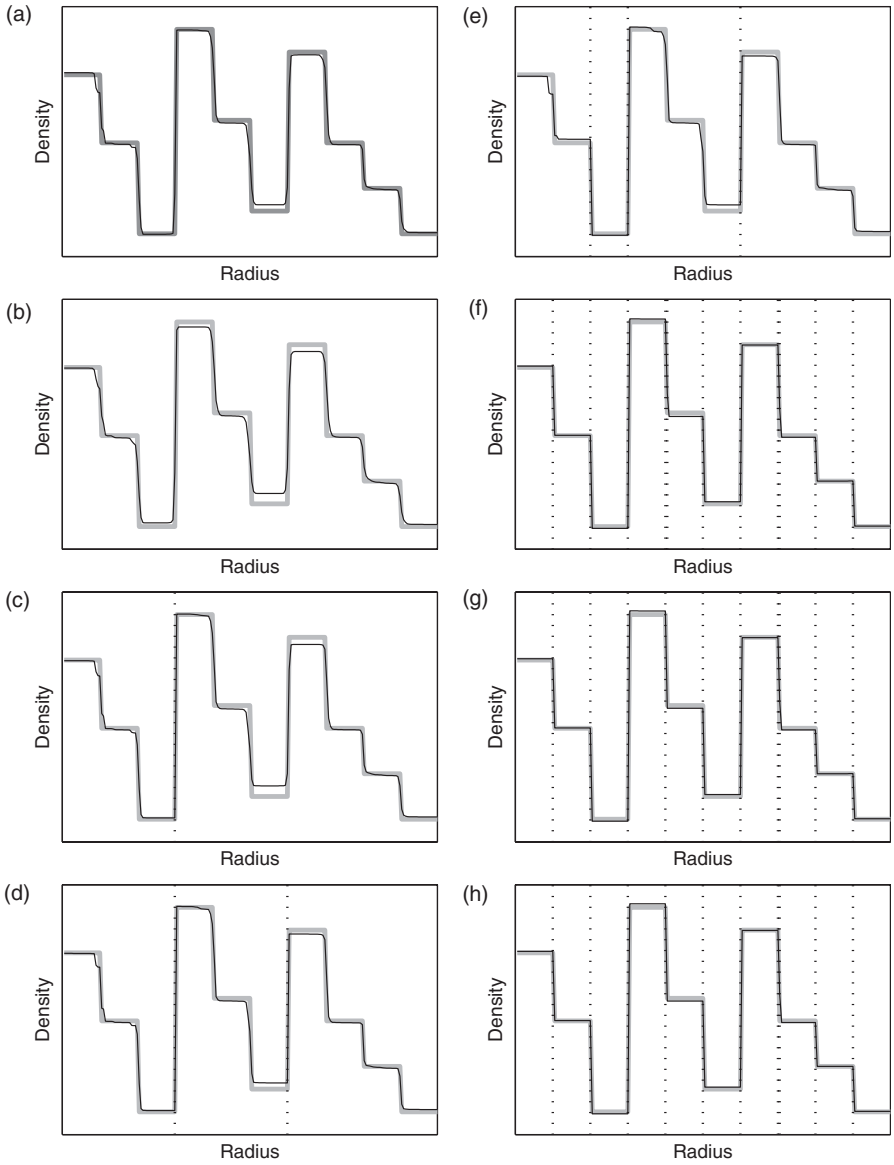

Figure 2. Sample iterative steps during the adaptive TV reconstruction of simulated object #1. The original TV reconstruction (a) is iteratively modified through a procedure that identifies density discontinuity locations. Figures (b)–(g) show intermediate reconstructions and identified edges are shown as vertical dashed lines. The final figure (h) shows the final reconstruction where H^1 regularization has been applied to all regions not considered to be edge locations.

are shown in figure 3. The object is of similar description to the previous object. This object is more complex in that it has some regions of smoothly varying density and density discontinuities are not regularly spaced. The data is synthetic with added Gaussian noise of variance 1.0% of the maximum noiseless data value. The reconstructions are ordered as in the previous example. We note that the TV-based regularizers again perform better. Even the smallest object features are partially recovered. This suggests that

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

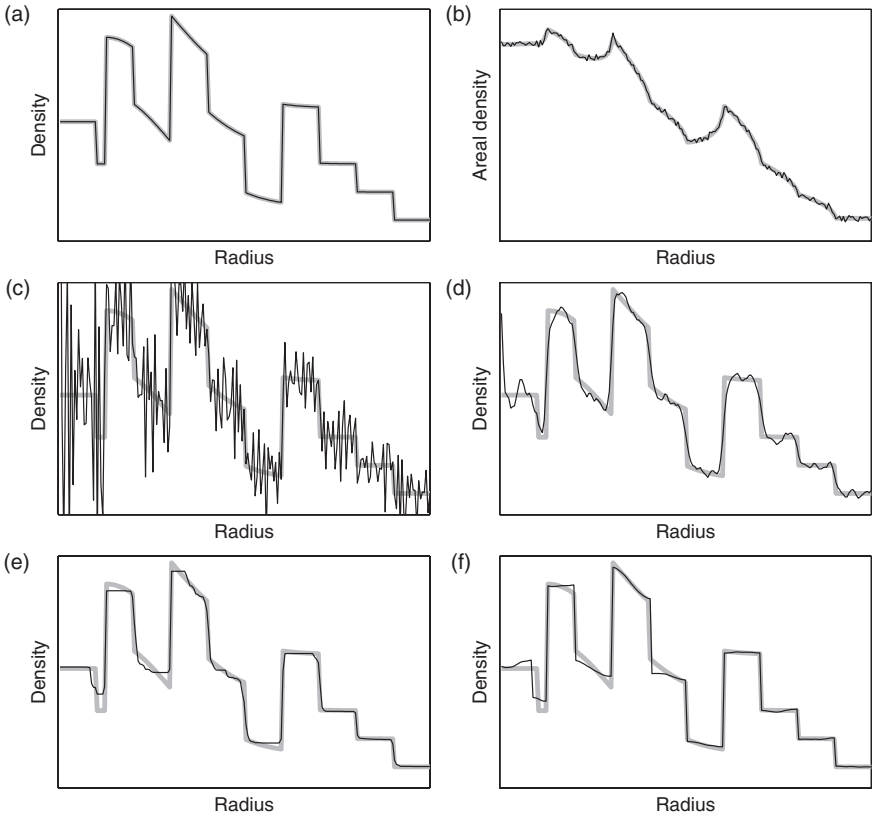

Figure 3. Regularized inversion examples on a synthetic data set generated from a variable density computational test object. The object density is defined at 200 radial positions. The density profile of the object (a) is projected onto 1D (b, gray) and Gaussian noise is added to obtain a synthetic 1D radiograph (b, black). The remaining figures are four regularized object density reconstruction examples (black) against the object reference (gray): (c) unregularized; (d) H^1; (e) TV; and (f) aTV.

our regularizer and data fidelity metric are excellent choices. Figure 3(e) does show staircasing, in which regions of nonzero slope are recovered as a series of piecewise constant values over multiple pixels.

3.3. *Simulated Example 1*

The third example is an edge-detection exercise from a simulated X-ray radiograph of a set of nested spherical shells. Figure 4 shows the radiograph (a), computed radial areal density (b), and reconstructions (c) and (d). Radiographs were numerically simulated using the Monte Carlo code MCNP[1] with X-ray scattering

[1] MCNP is a particle transport code that includes detailed physics treatments of photo-atomic and photo-electron interactions. It correctly treats Bremsstrahlung photons created from recoil or ejected electrons as well as photon scattering. It allows high-fidelity simulations of actual experimental conditions including source particulars and experiment geometry. Our simulated radiographs were provided by Jeff Favorite of the Applied Physics Division at Los Alamos National Laboratory.

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

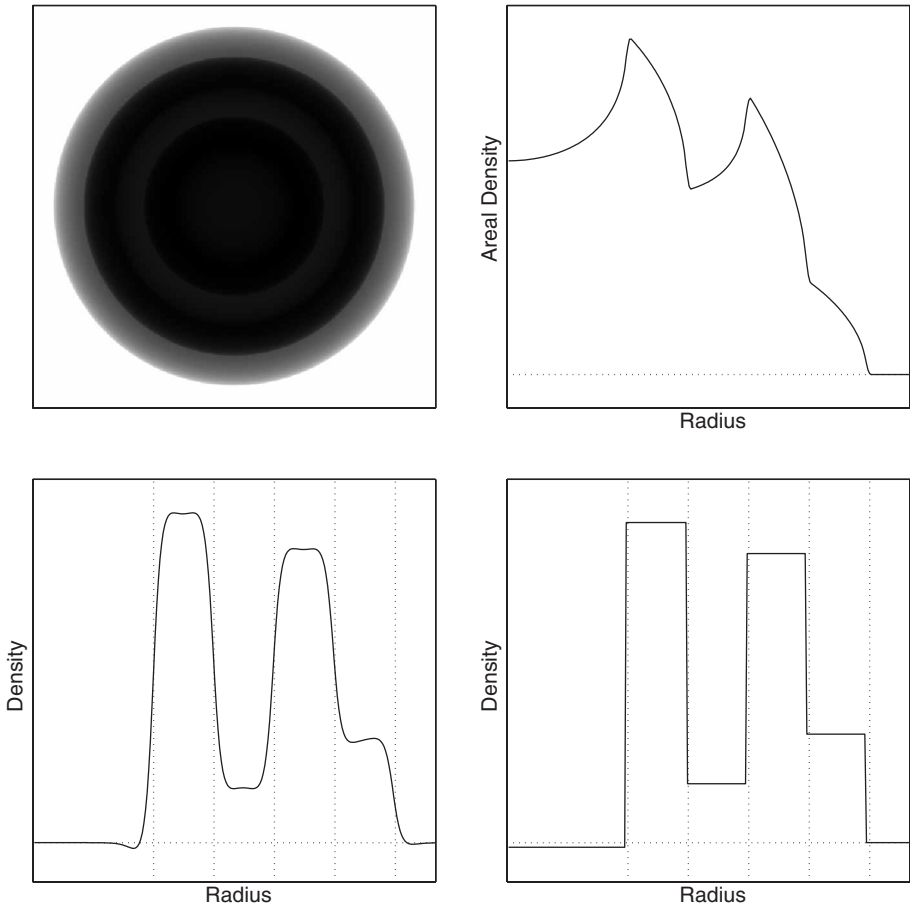

Figure 4. Edge detection example from a simulated X-ray experiment. In this example the simulation does not account for X-ray scattering. The four figures are: (a) false color radiograph; (b) computed radial areal density; (c) H^1 regularized inversion; and (d) aTV inversion. Locations of actual object density discontinuities are shown as vertical gridlines. The horizontal gridline indicates the zero density level.

suppressed. The simulated source is a bi-chromatic cobalt-60 source set at a distance sufficient to adequately approximate a parallel beam. The object is a set of four nested spheres whose boundary locations are indicated by vertical grid lines in the reconstruction figures. In this case, we must convert the radiograph from intensity (I) to areal density. Without explicit prior knowledge of the object, the areal density is not uniquely determined. Instead, we use an exponential attenuation approximation ($d \propto -\ln(I/I_0)$) and invert assuming that the object is a single material. The relative noise level is very small in this simulation with a standard deviation of approximately 0.01% of the maximum areal density value. We find that density edges are accurately reconstructed using a TV prior, but poorly defined using a smoothing regularizer. Because of the absence of calibration data, we do not attempt to recover densities.

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

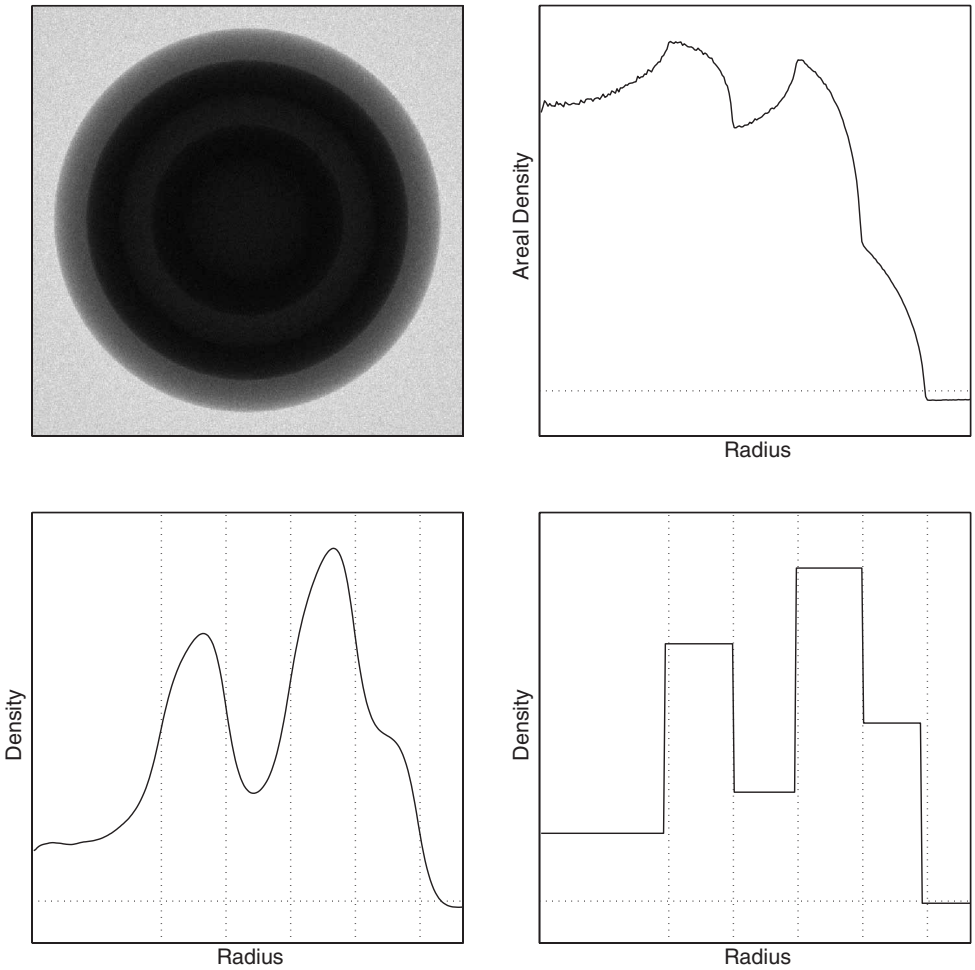

Figure 5. Edge detection example from a simulated X-ray experiment. This example includes the effects of X-ray scattering. The four figures are: (a) false color radiograph; (b) computed radial areal density; (c) H^1 regularized inversion; and (d) aTV inversion. Locations of actual object density discontinuities are shown as vertical gridlines. The horizontal gridline indicates the zero density level.

3.4. *Simulated Example 2*

The fourth example is nearly identical to that of the previous example. The data now includes X-ray scattering effects but the reconstruction process is the same and does nothing to account for the differences. Figure 5 shows (a) a false-color radiograph, (b) computed radial areal density, (c) H^1 regularized reconstruction, and (d) a TV-regularized reconstruction. No attempt was made to correct for scattering effects. We find that locations of material boundaries (density discontinuities) are still very well determined. The main effect of the scattering is to render impossible a quantitative analysis of the individual densities. In this relatively noiseless scenario, the material boundaries could be obtained from derivative information on the H^1 solution. However, the previous examples show that this is not generally the case.

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

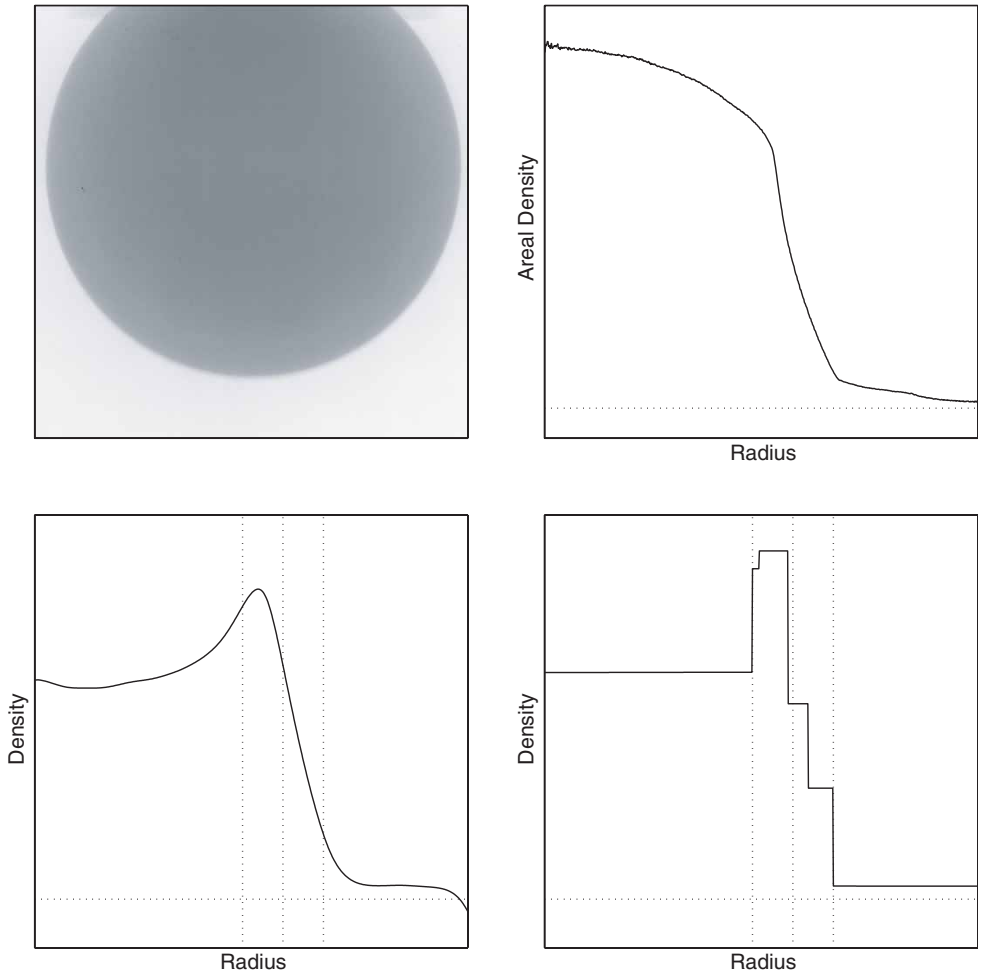

Figure 6. Edge detection example on a spherical nondestructive testing example object examined by an actual X-ray experiment. The four figures are: (a) false color radiograph; (b) computed radial areal density; (c) H^1 regularized inversion; and (d) aTV inversion. Locations of actual object density discontinuities are shown as vertical gridlines. The horizontal gridline indicates the zero density level.

3.5. *Experimental example*

The fifth example is the reconstruction of a similar spherical computational test object from an actual X-ray radiograph.[2] Figure 6 shows (a) a false-color radiograph, (b) radial areal density, (c) H^1 regularized reconstruction, and (d) aTV regularized reconstruction. Locations of the actual spherical shell boundaries are indicated by vertical gridlines in the reconstructions. We find that the H^1 solution has a very

[2] The radiograph was acquired by Hans Snyder and David Miko of the Nuclear Nonproliferation Division at Los Alamos National Laboratory.

Downloaded By: [Chartrand, Rick] At: 02:25 15 March 2007

difficult time distinguishing even qualitative object geometry. The aTV solution is significantly better in identifying all material boundaries. It does, however, retain two intermediate boundaries suggesting a more complex structure than in the real object. We attribute this effect to a combination of the scattering distortion effects, uncertainty in beam intensity correction across the image plane, and measurement noise. The effects of scattering are pronounced in the central region of the reconstructions; the interior and exterior densities of the actual object are both zero.

4. Conclusion

It is expected that any regularization based on prior knowledge of objects to be reconstructed will provide suppression of inversion amplified data noise and visually pleasing results relative to an unregularized solution. The particular choice of regularization, however, has significant influence on the details of the final result. We have shown that TV-based regularizations are much better choices for reconstructing objects of piecewise smooth density, especially if the location of density edges are an important reconstruction feature.

We have presented methods for the applying TV regularization to 1D Abel inversion problems with regularization in the two-dimensional object space. Our approach involves the use of a lagged diffusivity fixed point analysis. We also introduced a simple adaptive gradient operator approach for identifying density discontinuities and applying zero-cost penalty in the edge location set of the regularization cost functional.

Current and future work includes (1) the use of data fidelity that reflects known data characteristics, (2) extensions to 2D cylindrically-symmetric object descriptions $\rho(r, z)$, (3) arbitrary orientation of the object symmetry axis, and (4) arbitrary beam geometry (nonparallel).

Acknowledgment

This work was supported by funding from the Department of Energy in the Radiographic Requirements Program and the Technology Integration Program.

References

[1] Kak, A.C. and Slaney, M., 2001, *Principles of Computerized Tomography* (Philadelphia: Society for Industrial and Applied Mathematics).
[2] Vogel, C.R., 2002, *Computational Methods for Inverse Problems* (Philadelphia: Society for Industrial and Applied Mathematics).
[3] Asaki, T.J., Chartrand, R., Vixie, K.R. and Wohlberg, B., 2005, Abel inversion using total-variation regularization. *Inverse Problems*, **21**, 1895–1903.
[4] Chan, T.F. and Mulet, P., 1999 On the convergence of the lagged diffusivity fixed point method in total variation image restoration. *SIAM Journal of Numerical Analysis*, **36**, 354–367.