

Quantitative Object Reconstruction using Abel Transform X-Ray Tomography and Mixed Variable Optimization

Mark A. Abramson ^{*} Thomas J. Asaki [†] J. E. Dennis, Jr. [‡]
Kevin R. O'Reilly [§] Rachael L. Pingel [¶]

July 1, 2007

Abstract

This paper introduces a new approach to the problem of quantitative reconstruction of an object from few radiographic views. A mixed variable programming problem is formulated in which the variables of interest are the number and types of materials and geometric parameters. To demonstrate the technique, we considered the problem of reconstructing cylindrically symmetric objects of multiple layers from a single radiograph. The mixed variable pattern search (MVPS) algorithm for linearly constrained problems was applied by means of the NOMADm MATLAB[®] software package. Numerical results are presented for several test configurations and show that, while there are difficulties yet to be overcome, the method is promising for solving this class of problems.

Key words: Mixed variable optimization, derivative-free optimization, mesh adaptive direct search (MADS) algorithms, x-ray tomography, Abel transforms

AMS 65K05, 65R32, 65R10, 90C30, 90C56

1 Introduction

Tomography refers to the “cross-sectional imaging of an object from either transmission or reflection of data collected by illuminating the object from many different directions” [19]. While the majority of research in computerized tomographic techniques has been focused on diagnostic medicine (*e.g.*, radioisotopes, ultrasound, magnetic resonance), there has also been significant

^{*}Air Force Institute of Technology, Department of Mathematics and Statistics, 2950 Hobson Way, Bldg 641, Wright Patterson AFB, Ohio 45433 USA, Mark.Abramson@afit.edu, <http://www.afit.edu/en/ENC/Faculty/-MAbramson/abramson.html>. Support for this author was provided by LANL.

[†]Los Alamos National Laboratory, MS D413, Los Alamos, New Mexico 87545 USA, asaki@lanl.gov

[‡]Rice University, Department of Computational and Applied Mathematics, 8419 42nd Avenue SW, Seattle, WA 98136-2360 USA, dennis@rice.edu, <http://www.rice.edu/~dennis>. Support for this author was provided by LANL, AFORS F49620-01-1-0013, The Boeing Company, and ExxonMobil Upstream Research Company.

[§]United States Air Force

[¶]Brigham Young University, Department of Mathematics, 292 TMCB, Provo, Utah 84602 USA, rachaelp@byu.net

research in other areas, such as mapping of underground resources, nondestructive testing of engineered parts such as rocket engines, brightness distribution determination of a celestial sphere, and three-dimensional imaging using an electron microscope [19]. Our research is motivated by applications, such as cargo screening, in which we seek to recover the parameterized geometry of an object and make exact material identifications using high-energy polychromatic sources.

At its most fundamental level, particle-beam tomography is represented by a measurement operator Φ that maps an object description μ into the data space of radiographic measureables d , i.e., $\Phi\mu = d$. The operator Φ is a general description of the (forward) experimental process. The (inverse) reconstruction process is to determine the object μ that generates radiographic data d . The solution can be sought by two distinct methods. First, one can consider constructing an inverse measurement operator such that $\mu = \Phi^{-1}d$. In the applications we target, the approximations required to render Φ invertible are usually too severe for obtaining sufficiently accurate solutions [19]. The second method is to pose an optimization problem that minimizes an objective function $\|\Phi\mu - d\|$ in the object description variable space. This approach can provide accurate results, but the operation $\Phi\mu$ usually requires the implementation of a time-expensive computational code.

In this paper, we consider an optimization approach that provides a fully quantitative and accurate object description, even when using a quickly computed approximate forward measurement. Our goal is to provide an object description that includes complete geometric details and material compositions. This is accomplished for the class of problems in which the object has a low-dimensional parameterized description. As a concrete, common and important example, we consider the tomographic reconstruction of a cylindrically symmetric object from a single x-ray radiograph. The simplicity of the object should not be confused with a limitation of the method to simple geometries. The key idea is that the object description can be parameterized. We have also chosen to consider x-ray tomography because of its widespread use in applications and because of the highly nonlinear interpretation of Φ that must be considered, due to both experimental and physical complexity. This scenario provides a rich environment for testing our new methods.

This introduction continues with discussions on x-ray imaging and current methods used for reconstructing cylindrically symmetric objects. Section 2 gives a formal description of our proposed tomographic reconstruction method. Section 3 describes the generalized pattern search algorithm for linearly constrained mixed variable optimization problems. Section 4 describes the experimental conditions for numerical testing, and Section 5 presents computational results for several test configurations. Finally, Section 6 offers some concluding remarks.

1.1 X-Ray Imaging

The modeling of Φ , the imaging process of x-ray tomography, requires a consideration of three issues: x-ray source characterization, photon-matter interaction physics, and the detector response function. X-ray sources can provide a variety of photon beam characteristics characterized by spatial and temporal variations in energy and intensity. Physical limitations dictate the types of photon beams that can be easily generated. The most common sources, x-ray tubes and betatrons, produce polychromatic (multi-energetic) divergent photon beams. These photons then interact with the material placed between source and detector. The principle photon-

matter interactions that produce alterations in the beam are photoelectric absorption, Compton scattering, and pair-production. Interactions are material and energy dependent and result in photon absorption, energy loss, and redirection. This fundamentally changed photon beam then interacts with a detector. Detectors are characterized by energy and dose dependent efficiency, spatial blur, and darkfield noise and background.

A reasonably general expression for a radiograph is

$$I(y) = \int_0^\infty R \circ \left(S(E) \exp \left[- \int_{C(y)} \mu(x, E) d\ell \right] \right) dE, \quad (1)$$

where μ , as a function of location $x \in \mathbb{R}^3$ and energy E , is a photon attenuation coefficient description of the object of interest, $d\ell$ indicates a line integral along the unique beam path $C(y)$ incident at detector location $y \in \mathbb{R}^2$ (or \mathbb{R}^3), $S(E)$ is the energy-dependent incident photon number density, R is the detector response function, and $I(y)$ is the photon number (or total energy) recorded by the detector. Even this complex expression assumes that photon energies remain constant and that scattered photons do not interact with the detector. We can normalize (1) to the response obtained in the absence of any objects,

$$I_0(y) = \int_0^\infty R \circ S(E) dE, \quad (2)$$

and define our data $d = I/I_0$ as this normalized transmission radiograph. If we further impose spectral monochromaticity, $S(E) = \delta(E)$, we obtain the simpler form

$$d(y) = \exp \left[- \int_{C(y)} \mu(x) d\ell \right]. \quad (3)$$

If we consider objects of cylindrical symmetry, then attenuation $\mu(r, E)$ is expressed with respect to radial distance r from the center of the object, and the line integral in (1) is replaced by the Abel transform operator P :

$$I(y) = \int_0^\infty R \circ (S(E) \exp [-P(\mu(r, E))]) dE. \quad (4)$$

In imaging applications, P is a discrete matrix operation that projects a discrete object representation onto an image plane. Finally, we note the simplest expression for the imaging of a cylindrically symmetric object by a monochromatic beam:

$$d(y) = \exp [-P(\mu(r))]. \quad (5)$$

Since most medical and industrial applications implement polychromatic sources, quality solutions to (1) are of great value. The principal solution method is to eliminate the energy dependence on μ by defining an effective attenuation,

$$\mu_{\text{eff}}(x) \equiv \frac{\int_0^\infty \mu(x, E) S(E) dE}{\int_0^\infty S(E) dE}, \quad (6)$$

where $S(E)$ can be either the incident or detected photon number density, depending on the application. The simplified (and energy-linearized) radiograph equation now becomes

$$d(y) = \exp \left[- \int_{C(y)} \mu_{\text{eff}}(x) d\ell \right] \quad (7)$$

in the continuous case, and

$$d(y) = \exp [-P(\mu_{\text{eff}}(x))] \quad (8)$$

in the discretized case. Additional pre- and post-processing techniques are often applied for reducing both polychromatic and scattering effects (see [19], for example), but this is beyond the scope of this paper.

1.2 Abel Transform Tomography

In 1917, Radon discovered a way to mathematically reconstruct any function from its projections, and in 1972 Hounsfield invented the first x-ray computed tomographic scanner, which used the Radon transform to reconstruct an object from its x-ray projections [19]. Mathematical inversion of cylindrically symmetric objects was originally solved analytically by Abel in 1826 [2]. The Abel transform is the special case of the Radon transform in which all projections are identical. In this case, a single projection is sufficient for an exact object reconstruction from a pure projection.

The continuous parallel-projection Abel transform,

$$\hat{f}(z) = 2 \int_{|z|}^{\infty} \frac{r f(r)}{\sqrt{r^2 - z^2}} dr, \quad (9)$$

is the one-dimensional projection, $\hat{f}(z)$, of the two-dimensional radial function $f(r)$. It has the well-defined inverse,

$$f(r) = -\frac{1}{\pi r} \frac{d}{dr} \int_r^{\infty} \frac{z \hat{f}(z)}{\sqrt{z^2 - r^2}} dz. \quad (10)$$

Equation (9) can also be represented as a line integral equation along the projection direction:

$$\hat{f}(z) = 2 \int_0^{\infty} f(\sqrt{z^2 + \ell^2}) d\ell. \quad (11)$$

This correspondence can be seen in Figure 1, which shows a radially symmetric function $f(r)$ as shades of gray. Projection is indicated by the dashed arrows.

The inverse transform is sensitive to noise. This can be seen by defining a noisy transform $\hat{f}_n(z) \equiv \hat{f}(z) + \eta(z)$, where $\eta(z)$ denotes the noise, and applying the inverse transform. The result is

$$f_n(r) = f(r) - \frac{1}{\pi r} \frac{d}{dr} \int_r^{\infty} \frac{z \eta(z)}{\sqrt{z^2 - r^2}} dz, \quad (12)$$

which shows arbitrarily large noise amplification as $r \rightarrow 0$.

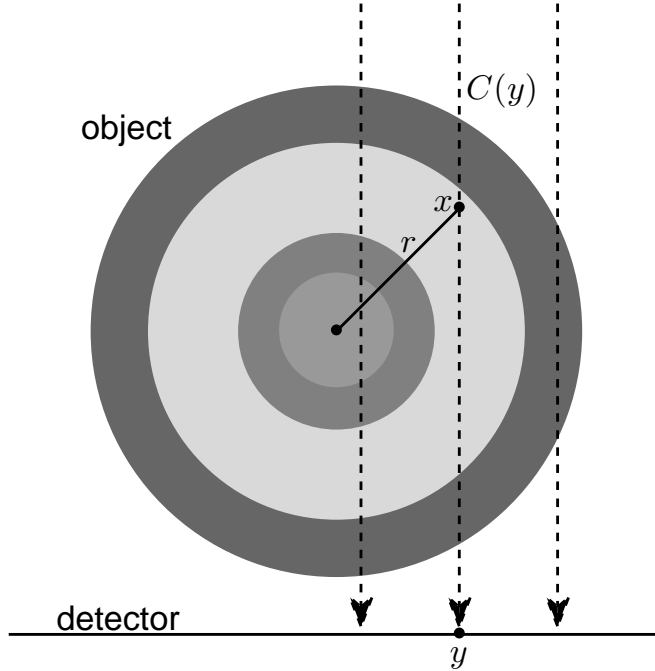


Figure 1: Cylindrically Spherical Object being x-rayed.

A variety of methods have been employed to overcome the difficulties associated with the inverse transform. Basic filtered and unfiltered inversions are described by Dasch [13]. Polynomial interpolation methods have been successful in describing objects of smoothly varying material properties [12, 25]. Basis function expansions are another popular area – see [17, 24] for some examples. Transform techniques are demonstrated by Smith et al. [27].

Asaki et al. [8] suggest discretizing the Abel transform and applying an appropriate regularization. Specifically, they pose the optimization problem,

$$\min_{\mu} F(\mu) = \|K\mu - \ln(d)\|_2 + \alpha R(\mu), \quad (13)$$

where K is a measurement operator that utilizes a discrete geometric projection (Abel transform), $R(\mu)$ is a regularization energy functional, and α is a regularization weight scalar. For a monochromatic source K reduces to a non-sparse projection matrix P . While discretization avoids the integral singularity in (10), it may still produce poor results for several reasons [7, 8]. First, K represents a significantly simplified physical model. Second, P is ill-conditioned, so that significant noise amplification can still occur at small r . Third, regularization methods are unable to incorporate significant prior knowledge about the object. Fourth, regularization methods do not directly provide a means for quantitatively determining the object's material composition except in the simplest cases [7].

For objects of piecewise constant material properties, the most interesting solution method is that of Deutsch et al. [15, 16], who introduce a rapidly converging algorithm for determining

the attenuation coefficients of a layered object. The method requires no prior knowledge of attenuation values, but does assume known and fixed layer boundary positions.

2 Mixed Variable Formulation

We formulate the problem as a mixed variable optimization problem and apply the Audet-Dennis mixed variable pattern search algorithm to solve it. Specifically, we consider objects with $n \in \{n_{\min}, n_{\min} + 1, \dots, n_{\max}\}$ concentric material layers (where $0 < n_{\min} < n_{\max}$), each with a material type $m_i \in M$ and outer edge location $x_i > 0$, $i = 1, 2, \dots, n$, that comes from a library (or list) of materials M and has a minimum thickness of $\delta > 0$. We assume a radiograph of fixed length, measured from the object center $x_0 = 0$ to the edge of the radiograph, which we denote by x_{n+1} for convenience. For radiograph data d of size n_d and measurement operator Φ , we consider the optimization problem,

$$\min_{n, m, x} f(n, m, x) = \frac{1}{\sqrt{n_d}} \|\Phi\mu(n, m, x) - d\|_2 \quad (14)$$

$$\text{s. t. } \delta \leq x_i - x_{i-1} \leq x_{n+1} - n\delta, \quad i = 1, \dots, n, \quad (15)$$

$$n\delta \leq x_n \leq x_{n+1} - \delta, \quad (16)$$

where $\mu(n, m, x)$ is a vector of material properties that depend on the number of material layers n , the material types $m \in M^n$, and material edge locations $x \in \mathbb{R}^n$ (we omit x_{n+1} , since it is not a variable). The linear constraints in (15)–(16) enforce the minimum thickness requirement on each material layer. This mixed variable problem is particularly challenging and interesting because the categorical variables m are nonnumeric and the problem dimension n is itself a variable in the problem.

Our formulation as a mixed variable optimization problem differs from the approach of Deutsch et al. [15, 16] in several ways. First, we allow both the number and position of material boundaries to vary. This allows for the very real possibility in applications that these quantities are unknown or only approximately known. Second, our proposed method does not rely on a linearized transformation between the object description and the radiograph data. This is important for applications in which the imaging particles have significant nonlinear interaction with the object being imaged. Third, our method provides an unambiguous object material description, rather than a scalar attenuation coefficient (or scalar mass density) from which a material must be inferred.

As a demonstration of our reconstruction method we make use of both (4) and (7), though our method can be as easily applied to the most general implementation $d = \Phi\mu$.

3 Mixed Variable Pattern Search

Audet and Dennis [9] introduced the class of mixed variable pattern search (MVPS) algorithms as an extension of the original pattern search algorithms [21, 22, 29] and proved convergence to appropriately defined [23] stationary points. A hierarchy of convergence results, depending on the smoothness of the objective function is given in [4]. Kokkolaras, Audet and Dennis [20]

applied the MVPS algorithm to the design of a thermal insulation system and showed a 65% reduction in the objective function value over previous results that optimized only with respect to the continuous variables. MVPS has also been extended to problems with stochastic noise in the objective function [28] and to problems with nonlinear constraints [4, 5]. A more general framework for derivative-free mixed variable optimization is described in [23].

In describing MVPS, we use the notation that is common to much of the previous work (see [4, 5, 28]), which is different from the rest of the paper. Unfortunately, this requires us to redefine certain variables (such as x) for only this section.

The class of MVPS algorithms is designed to numerically solve optimization problems in which each variable $x = (x^c, x^d)$ is partitioned into its continuous and categorical parts, $x^c \in X^c \subseteq \mathbb{R}^{n^c}$ and $x^d \in X^d \subseteq \mathbb{Z}^{n^d}$, respectively, where n^c and n^d denote the *maximum* respective dimensions of these variables. We adopt the convention of ignoring unused variables. The general MVP problem can now be expressed as

$$\min_{x \in X} f(x), \tag{17}$$

where $f : X \rightarrow \mathbb{R} \cup \{\infty\}$, and the domain X is the union of continuous domains across possible categorical variable values; *i.e.*,

$$X = \bigcup_{x^d \in X^d} (X^c(x^d) \times \{x^d\}),$$

with the convention that $X = X^c$ if $n^d = 0$. Furthermore, X^c is defined by a finite set of bound and linear constraints, dependent on the values of x^d . That is,

$$X^c(x^d) = \{x^c \in \mathbb{R}^{n^c} : \ell(x^d) \leq A(x^d)x^c \leq u(x^d)\},$$

where $A(x^d) \in \mathbb{R}^{m \times n^c}$ is a real matrix, $\ell(x^d), u(x^d) \in (\mathbb{R} \cup \{\pm\infty\})^m$, and $\ell(x^d) \leq u(x^d)$ for all values of x^d . Note that this formulation is a generalization of the standard NLP problem, in that it reduces to a standard NLP if $n^d = 0$, in which case, ℓ , A , and u do not change.

In solving an MVP problem, we note that continuous relaxation of the categorical variables is not possible, which precludes the use of branch and bound methods. Furthermore, we assume that the categorical variable space is sufficiently large such that optimizing over every possible combination of categorical variable values is intractable. This means that the best we can hope for is a locally optimal solution. However, a notion of local optimality for MVP problems is not well-defined because there is no topology associated with the nonnumeric categorical variables. Local optimality must therefore be defined in terms of a problem-specific set of discrete neighbors. This is done by defining a continuous set-valued function $\mathcal{N} : X \rightarrow 2^X$, where 2^X is the power set of X (*i.e.*, the set of all possible subsets of X). Thus a point $y \in X$ is a discrete neighbor of X if $y \in \mathcal{N}(x)$. By convention, $x \in \mathcal{N}(x)$.

For a point $x \in X$ to be locally optimal, we mean that it is locally optimal in the traditional sense when the categorical variables are fixed, that $f(x) \leq f(y)$ for all $y \in \mathcal{N}(x)$, and that, for any $y \in \mathcal{N}(x)$ if $f(x) = f(y)$, then y is also locally optimal (in the traditional sense) when categorical variables are fixed. A formal definition is given in [9].

As an example, Kokkolaras, Audet, and Dennis [20] studied the design of a thermal insulation system of fixed length consisting of a certain number of insulators of various material types and

thicknesses, and a corresponding set of heat intercepts placed between each pair of insulators. The heat intercepts were set at specified temperatures, and the objective was to minimize the power to keep one of the ends at a fixed temperature. Given a specific design, a discrete neighbor was defined to be any design obtained by replacing any single insulator with one of a different type, adding an insulator and corresponding heat intercepts at any location, or removing any insulator with its adjacent heat intercept.

Pattern search is an iterative method that generates a sequence of feasible points with non-increasing function values. At each iteration, the objective function f is evaluated at points lying on a mesh in an attempt to find a point with a lower function value than that of the incumbent. The mesh is defined by a set of directions that form a positive spanning set [14] (i.e., a set of directions such that any vector in the space can be expressed as a nonnegative linear combination of these directions).

At each iteration, the mesh is the direct product of the union of a finite number of lattices in \mathbb{R}^{n^c} with the integer space \mathbb{Z}^{n^d} , as follows. For each combination $i = 1, 2, \dots, i_{\max}$ of values that the categorical variables may take on, a set of positive spanning directions $D^i = G_i Z_i$ is formed, where $G_i \in \mathbb{R}^{n^c \times n^c}$ is a nonsingular generating matrix and $Z_i \in \mathbb{Z}^{n^c \times |D^i|}$. The mesh is then the direct product of X^d with a union of a finite number of lattices in X^c centered at the continuous part of the current iterate:

$$M_k = \bigcup_{i=1}^{i_{\max}} M_k^i \times X^d, \quad M_k^i = \bigcup_{x \in S_k} \{x^c + \Delta_k D^i z : z \in \mathbb{Z}_+^{|D^i|}\} \subset \mathbb{R}^{n^c}, \quad (18)$$

where $\Delta_k > 0$ is the mesh size parameter, and S_k is the set of all previously evaluated trial points S_k (with S_0 as the set of initial points). The set of discrete neighbors is assumed to be constructed such that every neighbor lies on the current mesh.

The evaluation of trial points on the mesh is performed in three steps: an optional SEARCH step and a local POLL step, and an EXTENDED POLL step. In the SEARCH step, the function is evaluated at a finite number of mesh points S_k . The user may specify virtually any finite strategy (including none) for identifying these points. Typically, this step is more global in nature and can make use of a favorite heuristic, a set of space-filling points, or specialized knowledge the user may have. If f is computationally expensive to evaluate, one common approach is to construct and optimize an inexpensive surrogate function [11] on the mesh at each SEARCH step.

If the SEARCH is unsuccessful in finding an *improved mesh point* (i.e., one that has a lower function value than the incumbent), the POLL step is performed, in which mesh points that are adjacent to the incumbent are evaluated, both in the continuous sense (while holding the categorical variable values constant) and in the discrete sense (the current set of discrete neighbors).

Polling with respect to continuous variables requires use of positive spanning sets in \mathbb{R}^{n^c} . Let $D_k^i \subseteq D^i$ denote the set of poll directions corresponding to the i -th set of categorical variable values for each iteration k . The *poll set* $P_k(x)$, centered at a point $x \in X$, is the set of neighboring mesh points in the directions D_k^i , while holding the categorical variables fixed; i.e.,

$$P_k(x) = \{x\} \cup \{(x^c + \Delta_k d, x^d) \in X : d \in D_k^i\} \subset M_k. \quad (19)$$

If polling with respect to the continuous variables fails to find an improved mesh point in $P_k(x_k)$, polling is performed on the current set of discrete neighbors $\mathcal{N}(x_k)$. The objective

function f is evaluated at each of the points in $P_k(x_k) \cup \mathcal{N}(x_k)$ until a lower objective function is found or until all these points have been evaluated.

If the both the SEARCH and POLL steps fail to find an improved mesh point, then the EXTENDED POLL step is performed. In this step, additional polling is performed around each discrete neighbor point $y \in \mathcal{N}(x_k)$, whose objective function value was only a small amount greater than that of the incumbent. This is done with the idea that, while y was not an improved mesh point, it is a promising region of the space in which to search for one; thus polling around y may produce a better objective function value than the incumbent. The EXTENDED POLL step is performed whenever $f(y_k) < f(x_k) + \xi_k$, where the *extended poll trigger* ξ_k satisfies $\xi_k \geq \xi$ for a fixed $\xi > 0$. The values for ξ_k are often set as a percentage of the objective function value at the current iterate [4]. Larger values of ξ will result in more extended polling, thus more computational cost per iteration, but may result in a better final solution.

For each iteration k and for each point $y_k \in \mathcal{N}_k^\xi = \{y \in \mathcal{N}(x_k) : f(x_k) \leq f(y) \leq f(x_k) + \xi_k\}$, the EXTENDED POLL step begins by performing a poll around the discrete neighbor $y_k = y_k^1$ and then continues polling around a finite sequence points $\{y_k^j\}_{j=1}^{J_k}$ satisfying $f(y_k^j) < f(y_k^{j-1})$ for all $j = 2, 3, \dots, J_k$. The index J_k occurs when either $f(y) < f(x_k)$ for some $y \in P(y_k^{J_k})$ or until $f(y) \geq f(y_k^{J_k})$ for all $y \in P(y_k^{J_k})$. The set of EXTENDED POLL points is therefore expressed as

$$\mathcal{X}_k(\xi_k) = \bigcup_{y_k \in \mathcal{N}_k^{\xi_k}} \left(\bigcup_{j=1}^{J_k} P_k(y_k^j) \right). \quad (20)$$

If the SEARCH, POLL, or EXTENDED POLL step is successful in finding an improved mesh point, that point becomes the new incumbent x_{k+1} and the mesh is retained or coarsened. If all three steps fail to find an improved mesh point, then the incumbent is retained (*i.e.* $x_{k+1} = x_k$) and declared to be a *mesh local optimizer*, and the mesh is refined. Refining or coarsening the mesh is done by updating the mesh size parameter according the rule,

$$\Delta_{k+1} = \tau^{w_k} \Delta_k, \quad w_k \in \begin{cases} \{0, 1, \dots, w^+\} & f(x_{k+1}) < f(x_k), \\ \{w^-, w^- + 1, \dots, -1\} & \text{otherwise.} \end{cases} \quad (21)$$

where $\tau > 1$ is rational, and $w^- \leq -1$ and $w^+ \geq 0$ are integers. Mesh coarsening does not affect theoretical convergence properties. In practice, it can slow convergence, but it can also cause the algorithm to skip over a local minimum and find a better one [5].

The full description of the MVPS algorithm is given in Figure 2. Convergence of the algorithm to suitably defined first-order stationary points was proved in [4] and [9].

For problems with linear constraints, infeasible points are simply discarded without being evaluated by the objective function. To guarantee that theoretical convergence properties still hold, the only additional requirement is that the rule for selecting polling directions must conform to the geometry of the nearby linear constraint boundaries [4, 10]. An algorithm for constructing conforming directions is given in [22] and [6] in the nondegenerate and degenerate cases, respectively.

Mixed Variable Pattern Search (MVPS) Algorithm

Initialization: Let $x_0 \in X$ satisfy $f(x_0) < \infty$. Set $\Delta_0 > 0$, $\xi > 0$, $0 < \varepsilon \ll 1$.

For $k = 0, 1, 2, \dots$, perform the following:

1. Set the extended poll trigger $\xi_k \geq \xi$.
2. **SEARCH step:** Employ some finite strategy seeking an improved mesh point; *i.e.*, $x_{k+1} \in M_k$ such that $f(x_{k+1}) < f(x_k)$.
3. **POLL step:** If the SEARCH step does not find an improved mesh point, evaluate f at points in $P_k(x_k) \cup \mathcal{N}(x_k)$ until an improved mesh point x_{k+1} is found (or until all points have been evaluated).
4. **EXTENDED POLL step:** If the SEARCH and POLL steps do not find an improved mesh point, evaluate f at points in $\mathcal{X}_k(\xi_k)$ until an improved mesh point x_{k+1} is found (or until all points have been evaluated).
5. **Update:** If SEARCH, POLL, or EXTENDED POLL finds an improved mesh point, update x_{k+1} , and set $\Delta_{k+1} \geq \Delta_k$ according to (21); otherwise, set $x_{k+1} = x_k$, and set $\Delta_{k+1} < \Delta_k$ according to (21).
If $\Delta_{k+1} < \varepsilon$, terminate the algorithm.

Figure 2: MVPS Algorithm

4 Implementation

Before presenting numerical results, we first describe how the x-ray tomography problem was setup to be solved, specifically with respect to the chosen set of discrete neighbors, materials considered, data generation, and experimental conditions.

4.1 Discrete Neighbors and Materials

Recall from earlier discussion that the set of discrete neighbors must be defined by the user. Thus, when a local solution to an MVP problem is found, it is always with respect to this set. Choosing all possible discrete neighbors might result in a better final solution, but it would require $\sum_{k=1}^{n_{\max}} |M|^k$ function evaluations at each unsuccessful iteration, where $|M|$ is the number of material types considered. This number grows very large, even for modest values of $|M|$ and n_{\max} . Defining the discrete neighborhood structure is perhaps the most crucial aspect of our approach to solving the x-ray tomography problem. It must take into account inherent properties of the Abel transform, as well as known properties of the material types.

To construct this set, we first introduce the idea of a pair of materials being *adjacent* to

each other if they share a common property defined by the user. In our case, adjacent materials may be thought of as materials that may be confused with one another in the current object configuration. For example, beryllium and aluminum exhibit similar properties with respect to the scattering and attenuation of x-ray photons as they pass through these materials. Since it would be easy to confuse them when analyzing the results of a radiograph, beryllium and aluminum may be considered adjacent. However, since lead is not easily confused with air, they would not be considered adjacent to each other.

The adjacency relationship among materials can be represented as a graph, composed of nodes that represent the materials and edges that represent the adjacency relationship between each pair of materials. The associated adjacency matrix is square, symmetric, and binary, where each row or column represents a material type, and a value of 1 in the $(i, j)^{th}$ element means that material i is adjacent to material j . To prevent redundancy, a material is not considered adjacent to itself, and thus the adjacency matrix will have zero entries along the diagonal.

Table 1 shows the complete list of materials we considered, their abbreviations (which are used throughout this paper), and the material class they represent. Seven of these materials were selected to be representative of a larger group of materials that all possess similar physical attributes and interact with x-ray photons in a similar manner. Copper (Cu) and Iron (Fe) were deliberately chosen to represent very similar material groups in order to test the algorithm’s ability to distinguish between two different materials with very similar physical attributes.

The adjacency matrix we used is given in Table 2. It represents a connected graph, so that any two materials can be reached from each other through a finite set of adjacency relationships. This is a key property in the optimization process [1], as it (in theory) allows all possible material configurations to be considered.

Table 1: Materials Considered.

Material	Abbreviation	Representative class
Air	Air	VOIDS and unpressurized gases
Polyethylene	Peth	Light plastics and organic materials
Beryllium	Be	Beryllium
Teflon	Teflon	Dense plastics
Aluminum	Al	Aluminum
Stainless Steel	Fe	Iron-like medium density metals
Copper	Cu	Other medium density metals
Lead	Pb	Heavy metals
Uranium	U	Very heavy metals

Given a specific configuration, we included several different types of discrete neighbors, obtained by doing one of the following:

1. *Swap* any layer with a material of a different type,
2. *Delete* any single material layer,
3. *Insert* a single material layer between any two existing layers or at the object center,
4. *Split* any layer into a pair of adjacent layers of different types,

Table 2: Adjacency Matrix for Materials.

	Air	Peth	Be	Teflon	Al	Fe	Cu	Pb	U
Air	0	1	1	1	1	0	0	0	0
Peth	1	0	1	1	1	0	0	0	0
Be	1	1	0	1	1	0	0	0	0
Teflon	1	1	1	0	1	1	1	0	0
Al	1	1	1	1	0	1	1	0	0
Fe	0	0	0	1	1	0	1	1	1
Cu	0	0	0	1	1	1	0	1	1
Pb	0	0	0	0	0	1	1	0	1
U	0	0	0	0	0	1	1	1	0

5. *Merge* any two layers together,
6. *Remove* simultaneously all layers with less than a specified minimum thickness,
7. *Combine* all physically adjacent layers of the same material type.

We limit this set of neighbors using the adjacency matrix, along with some other reasonable restrictions. In describing each in greater detail, we index the layers $i = 1, 2, \dots, n$, ordered from interior to exterior. The transmission layer $n + 1$ has fixed material type m_{n+1} and outer edge location x_{n+1} , which is the edge of the radiograph.

A swap of materials in a layer involves replacing the material of one layer with a different type from the material library M while holding all layer thicknesses fixed. The new material must be adjacent to the old material (as defined by the adjacency matrix), and $m_n \neq m_{n+1}$ must always hold.

Deletion of any layer i (with its material m_i and edge location x_i) may be performed as long as the number of layers n does not fall below n_{\min} . When layer i is deleted, either layer $i - 1$ or $i + 1$ can be chosen to take on the thickness of the deleted layer and to become layer i . When $i = n \geq n_{\min} + 2$ (i.e., the outermost layer) is considered for deletion, and layer $n - 1$ is of the same material type as the transmission layer (i.e., $m_{n-1} = m_{n+1}$), then layers n and $n - 1$ are both deleted at once.

Insertion of an additional layer between any two layers i and $i + 1$ may be done with any material that is adjacent (as defined by the adjacency matrix) to either m_i or m_{i+1} , provided that n does not exceed n_{\max} after insertion and provided layers i and $i + 1$ are sufficiently thick. In this case, the new edge location is set between x_i and x_{i+1} : either x_i is reduced enough to allow the new layer's edge location to take the old value of x_i (the inner layer's thickness is reduced), or else the new edge location is set to be far enough between x_i and x_{i+1} so that x_i can remain constant (the outer layer's thickness is reduced).

Splitting a layer not only divides the layer exactly in half, but also changes each material type to an adjacent one (actually, the next adjacent one in the list of adjacent materials). For example, Al would split into Teflon and Fe, in either order.

Merging two layers occurs when two layers are combined and the material type of the new layer is chosen to be the weighted average of the material types, relative to the indexing of the

adjacency matrix. For example, Teflon and Peth would be merged into Be.

For neighbors having all thin layers removed simultaneously, either the interior or exterior physically adjacent layer is expanded to compensate.

Finally, since the order in which points are evaluated can have a significant effect on the solution, two ordering strategies were employed. First, neighbors that result in fewer layers (deletion, merging, etc.) were evaluated first, while those that increase the number of layers (insertion, splitting, etc.) were evaluated last. Second, because the outside layers affect the x-rays more than the inside layers do, neighbors with changes to outside layers are always evaluated before those with changes to internal layers.

4.2 Experimental Conditions and Data Generation

The x-ray beam was modeled as produced by a 2.4 MeV Cygnus source [26] because the energy spectrum is well characterized and representative of more common polychromatic sources. We approximated this continuous source as 23 discrete monoenergetic sources of weighted intensities. Energies varied from 100 keV to 2.3 MeV in 100 keV intervals. The beam geometry for all experiments was plane parallel with uniform cross-sectional intensity.

The object symmetry axis and detector plane were oriented perpendicular to the incident beam propagation direction. In this paper we do not model scattering, so the distances between source, object, and detector have no effect on the results. The detector is given a resolution of 200 μm , and has 100% energy collection efficiency.

Simulated radiographs d are produced by constructing a model of the forward projection operator Φ . We used two distinct methods, both which utilize the Abel transform P . A linear polychromatic (LP) approximation is the superposition of linear attenuation radiographs at each of the 23 distinct energy levels of the Cygnus source. A linear monochromatic (LM) approximation uses a single attenuation scalar for each material that is the source-intensity weighted average of attenuation coefficients for each material. The attenuation coefficients were interpolated from National Institute of Standards and Technology data [18]. Both methods partially account for the energy dependence of the attenuation coefficients but in different ways. Simulated Poisson noise was added to the radiograph data assuming 30,000 monochromatic photons per pixel (unattenuated). This procedure provides a realistic noise signature, though not a technically accurate one. A more careful model would require a probabilistic analysis of energy loss in the photon-matter interaction processes. This procedure is beyond the needs of this study.

5 Numerical Testing

To demonstrate the effectiveness of the MVPS algorithm in quantitatively reconstructing objects, we applied the NOMADm MATLAB[®] software package [3] to three test sets that differ in their construction, with each set consisting of 6-10 test object configurations. Following a brief description of test conditions and parameter settings, we describe each test set in turn and present numerical results.

MVPS parameters were set as follows. The extended poll trigger was set at 0.99. At this value, extended polling occurred whenever a neighboring object's configuration resulted in an

objective function value that was less than 99% higher than that of the current best configuration. The initial mesh size was set at 0.1 cm and doubled ($\tau = 2, w_k = 1$) whenever an improved mesh point was found, provided the mesh size did not exceed 0.1 cm. If not found, the mesh size was divided in half ($w_k = -1$). Termination occurred the first time the mesh size fell below 10^{-3} . No SEARCH step was used and the poll step used the standard coordinate directions and their negatives.

To keep computational time reasonable, the minimum and maximum number of allowable material layers were set to $n_{\min} = 1$ and $n_{\max} = 6$, respectively, and the minimum layer thickness was set at $\delta = 0.08$ cm. We assume that we can independently obtain a good estimate of the object radius r , and we set the initial guess to be a cylinder of Al, but represented as two adjacent layers of equal thickness; i.e., $x = [r/2, r]$ and $m = \{\text{Al}, \text{Al}\}$. The representation of two layers of the same material is done to help the algorithm avoid early termination due to limitations of the discrete neighbor set. Aluminum was chosen as the initial material guess because it represents the “center” of the material library.

5.1 Test Scenarios

Scenario 1. The first test set consists of six objects with 3 cm radii, each having three layers whose material types are the six possible orderings of Fe, Be, and Peth, and whose thickness are all equal to 1 cm (See Table 3). Radiographs were computed using the LM method. The material library used in the MVP algorithm included all the materials from Table 1 except Cu. The adjacency matrix was that of Table 2 with the row and column for Cu removed. The cost function computes a relative merit value using the same LM approximation used in radiograph construction.

This experiment is ideal for initial testing because the material types are distinct in their attenuation characteristics and because the cost function and simulation utilize the same forward measurement operator, allowing for the possibility of an exact solution, modulo the effects of data noise.

Table 3: Test Set 1: Object Configurations

Object	Materials	Edge Locations (cm)
1a	[Fe, Be, Peth]	[1.0, 2.0, 3.0]
1b	[Fe, Peth, Be]	[1.0, 2.0, 3.0]
1c	[Be, Fe, Peth]	[1.0, 2.0, 3.0]
1d	[Be, Peth, Fe]	[1.0, 2.0, 3.0]
1e	[Peth, Fe, Be]	[1.0, 2.0, 3.0]
1f	[Peth, Be, Fe]	[1.0, 2.0, 3.0]

Scenario 2. The second test set consists of 10 randomly generated configurations of 1–4 layers with an enforced minimum layer thickness of 0.1 cm. The possible object materials were limited to a subset of the full material library: Air, Peth, Be, Teflon, Al, and Fe. These objects are described in Table 4. The radiographs were computed using the LP method. The material

library, adjacency matrix, and cost function used by the MVP algorithm were the same as Scenario 1.

This experiment was used to test the ability of the algorithm to identify objects of greater variety of materials and layer thicknesses. It was also used to examine the effects of using a fully linearized measurement operator (LM) in the cost function for examining radiographs produced by considering approximate polychromatic effects (LP).

Table 4: Test Set 2: Object Configurations.

Object	Materials	Edge Locations (cm)
2a	[Fe, Teflon, Fe]	[1.3973, 1.7028, 2.7225]
2b	[Be]	[2.4927]
2c	[Be, Fe]	[1.7136, 2.4441]
2d	[Air, Al]	[3.6109, 3.8152]
2e	[Be, Air, Teflon, Be]	[0.2379, 0.8433, 1.9587, 3.4136]
2f	[Al, Teflon, Be]	[0.2151, 2.6260, 3.7330]
2g	[Fe, Be]	[1.4035, 2.8271]
2h	[Peth, Al, Air, Fe]	[1.7489, 2.4271, 3.0819, 3.6769]
2i	[Air, Al, Teflon, Al]	[0.8161, 1.8739, 2.0518, 3.7236]
2j	[Be, Peth, Fe]	[1.3628, 1.9025, 3.6278]

Scenario 3. The final scenario differs from Scenario in only two respects. First, the material library uses the full set of materials from Table 1 (including Cu). Second, and more importantly, the cost function uses the LP method. This means that, similar to Scenario 1, the cost function and simulation apply the same measurement operator, allowing for the possibility of an exact solution, modulo the effects of data noise.

5.2 Results

For each object, a run was made in which all iterates were recorded, and the best point for each setting of the categorical variable values was extracted, with the restriction that any solution with any layer width less than 0.1 cm was discarded. This process led to a small set of meaningful candidate solutions. Since we actually know the true solution for each object, we performed an additional run from a different initial point whenever we failed to find a point whose function value came within 5% of that of the true solution. The second initial point is the same as the first, except the inner material is changed to Air; i.e., $x = [r/2, r]$, $m = \{\text{Air, Al}\}$

In describing the results that follow, we denote by $nSol$ the total number of solutions found whose function value was within 5% of that of the true solution, and we denote by $rank$ the rank of the “true” solution in the list of best points found. For each object we attach a rating of H (Hit), I (Inconclusive), or M (Miss), which has the following interpretation:

1. H ($nSol > 0$ and $rank \leq nSol$): This means that there is a set of solutions and the true solution is among them. Ideally, we want $nSol = rank = 1$, but we recognize that a small set of ranked solutions is often a more meaningful result.

2. I ($nSol = 0$): This is the restart indicator; none of the solutions seem to adequately describe the data (objective function value within 5% of the true solution), from which we conclude that we are stuck at a local minimizer. If $rank > 0$, then the best solutions had configurations that came close with respect to materials and edge locations but not objective function value.
3. M ($nSol > 0$ and the true solution is unranked or not present): This is the worst situation because there appear to be good solutions, so that a restart is not indicated, but the true solution is not in the solution set.

Our motivation for reporting results in this manner is as follows. Recall that the objective function value of a solution is a measure of how accurately the model and solution together describe the data (see (14)). An objective function value of zero would indicate a perfect match, but even the true solution does not achieve this because of the noise in the data. Because of both model deficiencies and data noise, it makes sense to consider candidate solutions with roughly the same objective function value as the true solution; i.e., we do not want to arbitrarily dismiss good local solutions. The choice of 5% in the discussion above was somewhat arbitrary, but the results are robust to the particular value. When the model is reasonably accurate, a true solution in categorical variables should attain an objective function value well within the 5% condition.

Tables 5–7 summarize the results for the three Test Sets, along with reruns for the last two scenarios. Solutions were obtained in a matter of minutes on a single processor PC, which was consistent with our goal of achieving solutions quickly, especially for time-sensitive applications. The median and mean solution times were 10 and 12 minutes, respectively, and the maximum solution time was 33 minutes for the initial run of Problem 2h.

Table 5: Scenario 1 results show an excellent 6-0-0 (Hit-Inconclusive-Miss) score.

Object	$nSol$	$rank$	Rating
1a	1	1	H
1b	1	1	H
1c	5	1	H
1d	4	2	H
1e	6	2	H
1f	1	1	H

Table 5 shows excellent performance for the scenario 1 set of test problems. In all cases, the initial point led to a set of solutions that contained the true solution. The largest solution set contained six candidates. In no case was the true solution ranked worse than second. Figure 3 illustrates the rating of H (hit) for object 1d. The true object (a) and the first three candidate solutions (b-d) are shown along with the corresponding objective function values f . All of the objects in the reconstruction set have similar objective function values. The optimization method returns very good solutions, but the objective function itself has difficulty distinguishing objects differing in categorical variables. All of the reconstructions shown provide equally likely explanations of the data.

Table 6: Scenario 2 results score a 4-5-1 (Hit-Inconclusive-Miss), but improve to 6-3-1 when we include results obtained from the second starting point.

Object	Initial run			Restart		
	<i>nSol</i>	<i>rank</i>	Rating	<i>nSol</i>	<i>rank</i>	Rating
2a	5	1	H	5	1	H
2b	1	1	H	1	1	H
2c	2	-	M	2	-	M
2d	0	2	I	1	1	H
2e	5	1	H	5	1	H
2f	2	1	H	2	1	H
2g	0	1	I	0	1	I
2h	0	1	I	1	1	H
2i	0	-	I	0	-	I
2j	0	-	I	0	-	I

Table 7: Scenario 3 results score a 6-4-0 (Hit-Inconclusive-Miss), but improve to 9-1-0 when we include results obtained from the second starting point.

Object	Initial run			Restart		
	<i>nSol</i>	<i>rank</i>	Rating	<i>nSol</i>	<i>rank</i>	Rating
3a	4	2	H	4	2	H
3b	1	1	H	1	1	H
3c	2	2	H	2	2	H
3d	0	2	I	1	1	H
3e	0	-	I	7	4	H
3f	2	1	H	2	1	H
3g	1	1	H	1	1	H
3h	1	1	H	1	1	H
3i	0	-	I	0	-	I
3j	0	1	I	1	1	H

In Table 6 we note decent performance for the scenario 2 set of test problems. Figure 4 illustrates the rating of M (miss) for object 2c. The true object (a) and the first three candidate solutions (b-d) are shown along with the corresponding objective function values f . The first two reconstructions attain objective function values with the 5% criterion and the third is similar, yet the true solution (in categorical variables) is not present. For the entire test set, we obtain an initial H-I-M result of 4-5-1, which improves to 6-3-1 with the addition of the five restart results. However, some problems could not be solved adequately. We attribute this mediocre performance to model inconsistency. The cost function uses a monochromatic beam approximation while the simulated radiographs were computed using a linear polychromatic approximation. In spite of this, we failed to generate a solution ($nSol > 0$) for only one test problem.

In Table 7 we note excellent performance for the Scenario 3 set of test problems. We are able to solve 9 of 10 problems (using both initial points), and the remaining problem (3i) still indicates a restart. No missed solutions are indicated. Figure 5 illustrates the rating of I




Figure 3: H (hit) rating example for test 1d showing the true object (a), the best three reconstructions (b-d), and corresponding objective function values f . Radial positions of layer boundaries are given in centimeters and material identifications are provided within or near each layer.




Figure 4: M (miss) rating example for test 2c showing the true object (a), the best three reconstructions (b-d), and corresponding objective function values f . Radial positions of layer boundaries are given in centimeters and material identifications are provided within or near each layer.

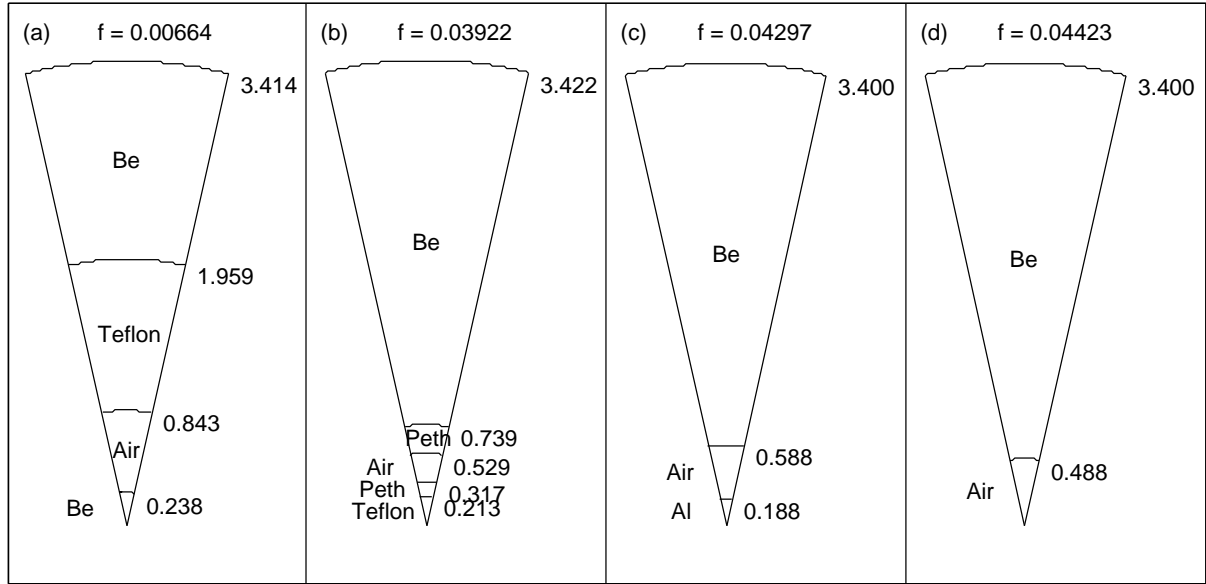


Figure 5: I (inconclusive) rating example for test 3e showing the true object (a), the best three reconstructions (b-d), and corresponding objective function values f . Radial positions of layer boundaries are given in centimeters and material identifications are provided within or near each layer.

(inconclusive) for object 3e before restart. The true object (a) and the first three candidate solutions (b-d) are shown along with the corresponding objective function values f . No solution provides a reasonable objective function value. This indicates that these solutions reside in very non-optimal local minima.

A comparison of the current results with those of previous methods would be appropriate. However, we are unaware of any other quantitative methods of recovering a parameterized object description from radiographs taken with a high-energy polychromatic source. One could apply regularized inverse projection techniques (see [7, 8], for example), but these methods assume a linearized projection and an energy-independent effective material attenuation, as seen in (7). The result is a profile of effective attenuation that still requires two significant postprocessing steps to achieve the geometric and material object description that our method provides. First, one must decide where material edges occur. Second, each effective attenuation value must be assigned to a material. Not only are these decisions affected by the linearization of the problem, but they are also dependent upon the choice of regularization parameter.

6 Conclusions

Quantitative reconstruction of cylindrically symmetrical objects from x-ray data normally requires an excessive amount of time to accomplish, and fast methods, such as regularized inversions fail to produce a quantitative description. This work represents the first real success in

this area, in which the result is a quantitative description of the object computed in a reasonable amount of time. MVPS is able to rapidly reconstruct an unknown object's composition from x-ray radiographs because of its ability to handle mixed variable problems and incorporate prior knowledge about the object.

Clearly, our approach is not foolproof. Some objects are very challenging to reconstruct because trial points are sensitive to either small changes in edge locations or the initial point. Including more discrete neighbors does, in fact, tend to improve solutions, but it also drives up the computational cost. Developing strategies to improve this process remains an open area of research. The success of our approach relies on the existence of a material library that is reasonably short and which contains all materials likely to be found in the object of interest.

This approach to x-ray tomography is not limited to cylindrically symmetric objects. Our formulation is, in a generalized sense, an optimization problem in a space of categorical variables (materials to consider) and continuous variables (object geometric description). The application potential for parameterized object identification from radiographs is clear. This approach is also not limited to particular assumptions on problem linearity. The measurement operator Φ can be any appropriate forward projection and can include such diverse information as source spectrum, beam geometry, nonlinear photon/matter interaction, detector response, and data postprocessing.

References

- [1] E. Aarts and J. K. Lenstra, editors. *Local Search in Combinatorial Optimization*. Princeton University Press, New York, 1991.
- [2] N. H. Abel. Auflösung einer mechanischen aufgabe. *J. reine angew. Math.*, 1:153–157, 1826.
- [3] M. A. Abramson. NOMADm optimization software. <http://www.afit.edu/en/ENC/Faculty/MAbramson/NOMADm.html>.
- [4] M. A. Abramson. *Pattern Search Algorithms for Mixed Variable General Constrained Optimization Problems*. PhD thesis, Department of Computational and Applied Mathematics, Rice University, August 2002.
- [5] M. A. Abramson, C. Audet, and J. E. Dennis, Jr. Filter pattern search algorithms for mixed variable constrained optimization problems. *Pacific Journal of Optimization*, to appear. Also appears as Technical Report TR04-09, Department of Computational and Applied Mathematics, Rice University, Houston, Texas, 2004.
- [6] M. A. Abramson, O. A. Brezhneva, and J.E. Dennis Jr. Pattern search in the presence of degenerate linear constraints. *Optim. Meth. Softw.*, to appear. Also appears as Technical Report TR03-09, Department of Computational and Applied Mathematics, Rice University, Houston, Texas, 2003.
- [7] T. J. Asaki, P. R. Campbell, R. Chartrand, C. E. Powell, K. R. Vixie, and B. E. Wohlberg. Abel inversion using total variation regularization: applications. *Inverse Prob. Sci. Eng.*, 14(8):873–885, 2006.

- [8] T. J. Asaki, R. Chartrand, K. R. Vixie, and B. Wohlerg. Abel inversion using total-variation regularization. *Inverse Problems*, 21:1895–1903, 2005.
- [9] C. Audet and J. E. Dennis, Jr. Pattern search algorithms for mixed variable programming. *SIAM J. Optim.*, 11(3):573–594, 2000.
- [10] C. Audet and J. E. Dennis, Jr. Analysis of generalized pattern searches. *SIAM J. Optim.*, 13(3):889–903, 2003.
- [11] A. J. Booker, J. E. Dennis, Jr., P. D. Frank, D. B. Serafini, V. Torczon, and M. W. Trosset. A rigorous framework for optimization of expensive functions by surrogates. *Struct. Optim.*, 17(1):1–13, February 1999.
- [12] C. J. Cremers and R. C. Birkebak. Application of the Abel integral equation to spectrographic data. *Appl. Optics*, 5(6):1057–1063, 1966.
- [13] C. J. Dasch. One-dimensional tomography: a comparison of Abel, onion-peeling, and filtered backprojection methods. *Appl. Optics*, 31(8):1146–1152, 1992.
- [14] O. L. Davies. *The Design and Analysis of Industrial Experiments*. Hafner Publishing Company, New York, 1954.
- [15] M. Deutsch, A. Notea, and D. Pal. Abel reconstruction of piecewise constant radial density profiles from x-ray radiographs. *Appl. Optics*, 28(15):3183–3186, 1989.
- [16] M. Deutsch, A. Notea, and D. Pal. Inversion of Abel’s integral equation and its application to NDT by x-ray radiography. *NDT International*, 23(1):32–38, 1990.
- [17] V. Dribinski, A. Ossadtchi, V. A. Mandelshtam, and H. Reisler. Reconstruction of Abel-transformable images: The Gaussian basis-set expansion Abel transform method. *Rev. Sci. Instr.*, 73(7):2634–2642, 2002.
- [18] J.H. Hubbell and S.M. Seltzer. Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients (version 1.4). <http://physics.nist.gov/xaamdi>, 2004.
- [19] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, Inc., New York, 1991.
- [20] M. Kokkolaras, C. Audet, and J. E. Dennis, Jr. Mixed variable optimization of the number and composition of heat intercepts in a thermal insulation system. *Optim. Eng.*, 2(1):5–29, 2001.
- [21] R. M. Lewis and V. Torczon. Pattern search algorithms for bound constrained minimization. *SIAM J. Optim.*, 9(4):1082–1099, 1999.
- [22] R. M. Lewis and V. Torczon. Pattern search methods for linearly constrained minimization. *SIAM J. Optim.*, 10(3):917–941, 2000.
- [23] S. Lucidi, V. Piccialli, and M. Sciandrone. An algorithm model for mixed variable programming. *SIAM J. Optim.*, 15(4):1057–1084, 2005.

- [24] C. D. Maldonado, A. P. Caron, and H. N. Olsen. New method for obtaining emission coefficients from emitted intensities: Circularly symmetric light sources. *J. Optic Soc. Amer.*, 55:1247–1254, 1965.
- [25] R. Piche. Noise-filtering properties of numerical methods for the inverse Abel transform. *IEEE Trans. Instr. Meas.*, 41(4):517–522, 1992.
- [26] J. R. Smith, R. Carlson, R. D. Fulton, R. Altes, V. Carboni, J. R. Chavez, P. Corcoran, W. L. Coulter, J. Douglas, D. Droemer, W. A. Gibson, T. B. Helvin, D. J. Henderson, D. L. Johnson, J. E. Maenchen, C. V. Mitton, I. Molina, H. Nishimoto, E. C. Ormond, P. A. Ortega, R. J. Quicksilver, R. N. Ridlon, E. A. Rose, D. W. Scholfield, I. Smith, A. R. Valerio, and R. White. Performance of the cygnus x-ray source. In T. A. Mehlhorn and M. A. Sweeney, editors, *Beams 2002*, volume 650 of *American Institute of Physics Conference Series*, pages 135–138, December 2002.
- [27] L. M. Smith, D. R. Keefer, and S. I. Sudharsanan. Abel inversion using transform techniques. *J. Quant. Spectroscopy Radiation Transfer*, 39:367–373, 1988.
- [28] T. A. Sriver, J. W. Chrissis, and M. A. Abramson. Pattern search ranking and selection algorithms for mixed variable stochastic optimization, 2004. Preprint.
- [29] V. Torczon. On the convergence of pattern search algorithms. *SIAM J. Optim.*, 7(1):1–25, February 1997.