Regularization of Ill-Posed Problems by Envelope Guided Conjugate Gradients

Linda KAUFMAN and Arnold NEUMAIER

We propose a new way to iteratively solve large scale ill-posed problems by exploiting the relation between Tikhonov regularization and multiobjective optimization to obtain, iteratively, approximations to the Tikhonov L-curve and its corner. Monitoring the change of the approximate L-curves allows us to adjust the regularization parameter adaptively during a preconditioned conjugate gradient iteration, so that the desired solution can be reconstructed with a low number of iterations. We apply the technique to an idealized image reconstruction problem in positron emission tomography.

Key Words: Envelope; Ill-posed; L-curve; Multiobjective optimization; Preconditioned conjugate gradients; Tikhonov regularization.

1. INTRODUCTION

Many problems in applied mathematics lead to models of the form

$$F(x) = y + \epsilon,$$

where x is an unknown vector of parameters, often restricted to a subset $\Omega \subset \mathbb{R}^n$ (e.g., by nonnegativity constraints $x \geq 0$) to be determined from a data vector y. The term ϵ is a noise term, often of unknown (but small) magnitude, which accounts for systematic model errors and random stochastic errors.

Frequently, the solution of $F(x) = y$ depends very sensitively on y, or is not even uniquely determined by it. This is particularly noticeable in discretized versions of ill-posed problems that often arise from so-called "inverse" problems for which standard existence and uniqueness arguments break down. To get reliable and robust estimates for x from the measured data y, one needs to exploit the availability of qualitative information about x, which usually is given in the form of a vague statement that some measure of smoothness of x is not large.

Sources of ill-posed problems include image smoothing, deconvolution (Wiener filtering), shape from shading, computer-assisted tomography (CAT, PET), indirect measurement, nondestructive testing, inverse scattering, seismic analysis, parameter identifi-

Linda Kaufman is a Member of Technical Staff, Bell Laboratories, Rm. 2c-461, 700 Mountain Avenue, Murray Hill, NJ 07974-0636; email: lck@research.bell-labs.com. Arnold Neumaier is Professor, Institut für Mathematik, Universität Wien, Strudlhofgasse 4, A-1090 Wien, Austria; email: neum@cma.univie.ac.at.

Journal of Computational and Graphical Statistics, Volume 6, Number 4, Pages 451–463

cation in dynamical systems, analytic continuation, inverse Laplace transform, relaxation spectra, and partial differential equations of mixed type (e.g., multiphase flow) or with nonstandard boundary conditions (e.g., backward heat equation). A recent survey by Engl (1993) gives details and references for a number of such problems, and an overview of the analytical results available. Hanke and Hansen (1993) surveyed the numerical analysis of methods for approximating the solution of ill-posed problems.

For example, image reconstruction is the problem of finding a piecewise smooth image x from data derived from this image (or rather the corresponding real life original). The data may be a noisy version of the image, an incomplete version of such a noisy image, an image degraded by an approximately known blurring operator, or projection data of various sorts. Depending on the modeling details, the data given are related to the desired information in a linear way,

$$Ax \approx b, \tag{1.1}$$

and, typical of inverse problems, the coefficient matrix is very ill-conditioned.

The problem of finding a smoothed version x of a noisy image b is treated by taking A as the identity matrix, that of completing an incomplete noisy image b with missing pixels to a full smoothed image x by taking for A a (0,1)-matrix with ones in the positions of the known pixels, and that of deblurring an image b degraded by a known linear blurring operator that defines A again has this form.

In all applications, the right-hand side is contaminated with noise from modeling errors and measurement errors, and the solution is only determined up to additive corrections by near-null vectors of the matrix. Near-null vectors may have quite large components, generally of an oscillating nature, and the oscillations picked up in this way are the cause of the "snowy" contribution to the reconstructed picture.

To suppress these terms—in standard terminology, to *regularize* the solution—one must exploit qualitative features of the expected solution, in particular the (piecewise) continuity of the solution, and a typical approach is to demand that some linear transform of the solution (suitable differences measuring "continuity" or "smoothness") must be kept reasonably small. This gives rise to a qualitative condition of the form

$$Jx \quad \text{well-scaled and not too large.} \tag{1.2}$$

In this article we introduce a new aspect to the analysis of ill-posed problems by looking at it from the point of view of multiobjective optimization. Indeed, the standard regularization techniques can be viewed as compromises designed to make small both a measure $r(x)$ for the lack of fit and a measure $q(x)$ for the lack of smoothness. This new point of view results in new numerical techniques for the robust iterative solution of high-dimensional ill-posed problems, which arise especially in problems where two- or three-dimensional functions or images need to be determined.

We demonstrate the effectiveness of the new algorithm to a problem with projection data from positron emission tomography (PET) with several different penalty functions.

2. REGULARIZATION AS A MULTIOBJECTIVE OPTIMIZATION PROBLEM

In this section, we discuss general linear regularization problems, where, with suitable matrices A and J, the informal problem (1.1)–(1.2) is formalized by requiring that both the residual norm

$$r(x) := \|Ax - b\|_2^2, \tag{2.1}$$

and the smoothness indicator (usually referred to as *penalty term*)

$$q(x) := \|Jx\|_2^2 \tag{2.2}$$

be minimized over a convex set of admissible vectors x. However, proofs are written in such a way that they also hold for convex nonquadratic q and, qualitatively, the results tend to be true also for more general nonlinear q and r. (Suitable nonquadratic (smooth) penalty terms are discussed, e.g., in Geman and McClure (1985), Hebert and Leahy (1990), Green (1990), Vogel and Oman (1996), and their relative merits are discussed in Lalush and Tsui (1992).)

In general, (2.1) and (2.2) are conflicting requirements, and one can minimize one only at the expense of the other. Therefore one settles for unimprovable vectors x (also called *efficient* or *Pareto optimal*, see, e.g., Sawaragi, Nakayama, and Tanino 1985), which have the property that for all admissible vectors y, either $r(y) \geq r(x)$ or $q(y) \geq q(x)$. If we write $q_0 = q(x)$, this is equivalent to the requirement that x minimizes $r(x)$ under the constraint $q(x) \leq q_0$, and by introducing a Lagrange multiplier λ, one sees that x must be a stationary point of

$$f_\lambda(x) := r(x) + \lambda q(x) \tag{2.3}$$

for some $\lambda \geq 0$. Because r and q are convex quadratics, x is in fact a minimizer of $f_\lambda(x)$. Conversely, it is easy to see that all minimizers of (2.3) with $\lambda \geq 0$ are Pareto optimal.

Alternatively, if one minimizes a differentiable compromise function of the form $\omega(q(x), r(x))$, where ω is strictly monotone in both arguments, the gradient

$$\nabla_q \omega(q(x), r(x))q'(x) + \nabla_r \omega(q(x), r(x))r'(x)$$

must vanish, whence, again, x is a stationary point of (2.3) with

$$\lambda = \nabla_q \omega(q(x), r(x))/\nabla_r \omega(q(x), r(x)) \geq 0.$$

Similarly, the optimization of $r(x)$ under the constraint $q(x) \leq \gamma$, or of $q(x)$ under the constraint $r(x) \leq c$, or more generally, of a compromise function $\omega(q(x), r(x))$ under a constraint $c(q(x), r(x)) \leq \gamma$ can again be shown to lead to a stationary point of (2.3).

Thus, (2.3) is the most general compromise function that needs to be considered, and it is the simplest since it is a convex quadratic function of x. The use of (2.3) to get an acceptable compromise goes back to Tikhonov (1963a, 1963b) for the case of solving ill-posed integral equations, and is generally referred to as *Tikhonov regularization*. See,

for example, Morozov (1984) and Engl (1993). In the context of maximum likelihood methods, this technique is derived from a Markov random field approach and known as the *maximum a posteriori* or MAP procedure, and the penalty term is then referred to as the *Gibbs prior* (see, e.g., Lalush and Tsui 1992). Statisticians usually refer to the special case where

$$q(x) = \|x\|_2^2, \tag{2.4}$$

as *ridge regression*.

However, the choice of λ is more an art than a science; usually, solutions are computed for a large number of λ's and one of them is selected by suitable heuristics; see Hansen (1992), Hanke and Hansen (1993), and Honerkamp and Weese (1990). Moreover, for many problems, the matrices are large and their sparsity structure is not sufficiently regular, which makes such an approach (using singular value decompositions) very expensive.

In the following, we take a closer look at the distribution of the set of all pairs $(q(x), r(x))$ and prove some general results that allow a flexible iterative approach to the regularization problem.

Definition 1. $Q = \{(q, r) | q, r \geq 0\}$ *denotes the positive quadrant in* $\mathbb{R}^2$. *We denote by* $\mathrm{conv}(S)$ *the* **convex hull** *of a subset* S *of* Q*—that is, the set of all convex linear combinations* $(\sum \alpha_k q_k, \sum \alpha_k r_k), \alpha_k \geq 0, \sum \alpha_k = 1$ *of points* $(q_k, r_k) \in S$. *We also define the* **envelope** *of* S *to be the set* $\mathrm{env}(S)$ *consisting of all pairs* $(q', r') \in Q$ *such that* $q' \geq q, r' \geq r$ *for some* $(q, r) \in \mathrm{conv}(S)$.

It is easy to see that the envelope of a nonempty set $S \in Q$ is unbounded, and its boundary is given by some L-shaped curve, which we call the *L-curve* of S. The definition of an envelope immediately implies that the full envelope consists of the pairs (q, r) with $r \geq \psi(q)$, where $\psi : \mathbb{R}_+ \to \mathbb{R}_+ \cup \{\infty\}$ is a convex and monotone decreasing function whose graph is the L-curve. When S is finite, the L-curve is piecewise linear and bends at vertices that are essentially the points of S lying on the L-curve. (A few may be missing if there are collinear triples of points.)

Theorem 1. *Suppose the rows of the matrix* (A^T, J^T) *are linearly independent. Then:*

1. *For each* $\lambda > 0$, *the optimization problem* (2.3) *has a unique solution* x_λ.
2. *Let* $q_\lambda := q(x_\lambda)$ *and* $r_\lambda := r(x_\lambda)$. *The curve* Γ *consisting of all pairs* (q_λ, r_λ) *is strictly monotone decreasing and convex, and its envelope contains the envelope of any set of pairs* $(q(x), r(x))$. *In particular,* Γ *is the L-curve of the envelope of the set of all pairs* $(q(x), r(x))$.

Proof: We give the proof in such a way that one can see that the theorem is also valid in the presence of convex constraints, and for more general convex functionals in place of q and r such that (2.3) is convex and has a unique finite minimizer.

1. Because the rows of the matrix (A^T, J^T) are linearly independent, the Hessian $2(A^T A + \lambda J^T J)$ of (2.3) is positive definite when $\lambda > 0$. Hence (2.3) is uniformly convex and has a unique minimizer x_λ. Because the Hessian is nonsingular, the implicit function theorem implies that x_λ depends continuously differentiable on λ.

2. For any $x^0 \in \mathbb{R}^n$, the vector $\hat{x}$ minimizing $r(x)$ under the constraint $q(x) \le q(x^0)$ has the form $\hat{x} = x_\lambda$ for some Lagrange multiplier λ and satisfies $r(\hat{x}) \le r(x^0)$. Hence any $(q(x^0), r(x^0))$ is in the envelope of Γ.

If we restrict the optimization to the curve $\{x_\mu | \mu \ge 0\}$ we see that $r_\mu + \lambda q_\mu$ takes its minimum at $\mu = \lambda$. Therefore the derivative with respect to μ vanishes at $\mu = \lambda$, giving

$$r'_\lambda + \lambda q'_\lambda = 0. \tag{2.5}$$

Because $\lambda \ge 0$, this shows that r_λ is increasing when q_λ is decreasing, and conversely. Since r_λ takes its global minimum at $\lambda = 0$, it must be increasing for small λ. To show a global strict monotone decrease of the curve Γ, it therefore suffices to show that Γ contains only a single point (q, r) with fixed value of either q or r.

To see this, note that minimality gives

$$r_\lambda + \lambda q_\lambda \le r_\mu + \lambda q_\mu, \quad r_\mu + \mu q_\mu \le r_\lambda + \mu q_\lambda.$$

Now if $q_\lambda = q_\mu$ for $\lambda, \mu > 0$, this implies that $r_\lambda = r_\mu$, and conversely. Thus, the point (q_λ, r_λ) is uniquely determined by each of its components, though λ itself need not be unique. This latter situation seems very pathological, and it is likely that it never occurs. For the linear case without constraints, this can indeed be proved using a generalized singular value decomposition; see Hansen (1992).

To show that Γ is convex, let $0 \le \mu' \le \mu \le \mu''$ and $\mu \le \lambda_0$. By monotonicity, the number $\alpha := (q_\mu - q_{\mu'})/(q_{\mu''} - q_{\mu'})$ lies between 0 and 1. The convexity of q and r implies that the vector $x^0 := x_{\mu'} + \alpha(x_{\mu''} - x_{\mu'})$ satisfies $q(x^0) \le q_{\mu'} + \alpha(q_{\mu''} - q_{\mu'}) = q_\mu$ and $r(x^0) \le r_{\mu'} + \alpha(r_{\mu''} - r_{\mu'})$. Constructing $\hat{x}$ from x^0 as before, we see that $q_\lambda \le q_{\mu'} + \alpha(q_{\mu''} - q_{\mu'}) = q_\mu$ and $r_\lambda \le r_{\mu'} + \alpha(r_{\mu''} - r_{\mu'})$ for some $\lambda > 0$. The first inequality gives us $\lambda \ge \mu$, hence $r_\mu \le r_\lambda \le r_{\mu'} + \alpha(r_{\mu''} - r_{\mu'})$. Thus, (q_μ, r_μ) lies on or below the line segment joining $(q_{\mu'}, r_{\mu'})$ and $(q_{\mu''}, r_{\mu''})$. □

We call the curve Γ in Theorem 1 the *Tikhonov L-curve*. The convention to draw the curves with the residual size on the vertical axis agrees with that used by Lawson and Hanson (1974, chap. 26). For points generated by an iterative method, this convention produces points "going down," which looks natural. Note, however, that the L-curves considered by Hansen (1992) are the transposes of ours, since he displays the residuals on the horizontal axis. Moreover, both Lawson and Hanson (1974) and Hansen (1992) draw L-curves using log-log scales, and in such scales the curves no longer look convex over the full range of λ.

Proposition 1. *Let $(q_0, r_0), (q_1, r_1), \ldots, (q_N, r_N)$ be the vertices of an L-curve of a finite set S, ordered such that*

$$r_0 > r_1 > \cdots > r_N. \tag{2.6}$$

Then

$$q_0 < q_1 < \cdots < q_N, \tag{2.7}$$

and the slopes

$$s_k := (r_{k-1} - r_k)/(q_k - q_{k-1}) \quad (k = 1, \ldots, N) \tag{2.8}$$

satisfy the relations

$$s_1 > s_2 > \cdots > s_N > 0. \tag{2.9}$$

Proof: Because the L-curve is piecewise linear, monotonicity gives (2.7) and convexity gives (2.9); all inequalities are strict since, by definition of a vertex, the derivative is discontinuous at each vertex. □

When solving (2.3) for a particular value of λ with an iterative method, we generate a sequence of points converging to a point on the Tikhonov L-curve. If we repeat this for many values of λ, we generate a set of points whose L-curve closely approximates that of the Tikhonov L-curve. In particular, the vertices of the approximating L-curves determine at any stage of the calculation a set of currently best compromises for the regularization problem. As in the case of Tikhonov regularization, the selection of a particular one of these compromises must rely on heuristics.

Intuitively, the most appropriate compromise seems to be the vertex where the L-curve is "most bent," (see Hansen 1992). However, the measure for assessing how much the curve bends at a vertex must be chosen with care. Indeed, if we simply rescale (1.1) and/or (1.2), we can change the shape of the L-curve drastically; clearly we must find a measure that is invariant under such changes. A natural further requirement is that this measure only depends on the vertex and its two neighboring vertices.

Scaling q and r simultaneously by the same amount leaves the slopes s_k defined in Proposition 1 unchanged, and scaling them by different amounts multiplies all s_k by the same factor. Hence the quotient of two consecutive slopes,

$$c_k := s_k/s_{k-1}, \tag{2.10}$$

is scale invariant, and approaches the minimal value 1 when the slopes are nearly the same—that is, when the L-curve is hardly bent. Thus, a large quotient (2.10) indicates a strong bend.

Another indicator results from looking at curvature terms measuring bentness based on a difference quotient of slopes, with correction factors designed to make this difference quotient scaling invariant. The simplest choice is

$$c_k := q_k(s_k - s_{k+1})/(r_{k-1} - r_{k+1}). \tag{2.11}$$

This expression is always positive, and approaches zero as the bend at the vertex becomes weaker and weaker. Again, a large value of (2.11) indicates a strong bend.

Thus, we may consider the following criterion for deciding on a corner:

Definition 2. *The L-curve is said to be* **most bent** *at the vertex* (q_k, r_k) $(1 \le k \le N-1)$*, where* c_k*, given by* (2.10) **(version 1)** *or* (2.11) **(version 2)***, is largest. This point is called the* **corner** *of the L-curve; it is called* proper *when* $k \ne 1, N-1$.

Remark 1.

1. *In many cases, the two versions define the same corner; but sometimes the decision depends on the version chosen.*

2. *Due to the nature of this test it is impossible to check for a bend at the leftmost and rightmost vertex ($k = 0, N$). Hence we accept a corner as proper only if it is interior enough; that is, if its index lies between 2 and $N - 2$. A corner at $k = 1, N - 1$ is a clear sign that only one branch of the L-curve has been found, and further points should be computed to see whether the corner moves or becomes proper.*

3. AN ENVELOPE GUIDED CONJUGATE GRADIENT ALGORITHM

Hansen and O'Leary (1993) suggested a scheme for determining the point of maximum curvature of the Tikhonov L-curve, by calculating suitable points along the L-curve, using a bracketing scheme to find the point of maximum curvature. Clearly, their algorithm assumes that the calculation of the optimum x of the penalty function (2.3) for a new value of λ is not time consuming. This assumption, unfortunately, is not reasonable for most large-scale problems. This does not mean that the L-curve approach must be abandoned, but it must be suitably modified.

In our iterative approach we move in the (q, r) plane trying to approximate a corner of the Tikhonov L-curve, but instead of finding points on the curve itself, we use the L-curves determined by the envelopes of the set of points generated so far to guide us to the desired destination. We try to get better and better approximations to a point of large curvature of the Tikhonov L-curve by monitoring the corners of the current L-curves.

The sets of points whose envelope defines the current L-curves are generated by the preconditioned conjugate gradient method described in Kaufman (1993), applied to the objective function (2.3) for suitable values of λ. For each particular value of λ, our aim is not so much to optimize (2.3) as to try to find points outside the current envelope that will bring the L-curve closer to the optimal Tikhonov L-curve, especially near the desired corner.

Assuming one has computed points x_k with which one has determined values (q_k, r_k), one can determine the vertices of the envelope as follows:

3.1 ENVELOPE CONSTRUCTION ALGORITHM

1. Sort the points so that $r_k \leq r_{k+1}$, that is, r decreases.
2. Discard any point k for which $q_k \geq q_{k+j}$ for all j. Thus, all points left in the set have been ordered so that r decreases and q increases. Renumber the ordered points with consecutive numbers.
3. Insist on convexity by computing for each point k,

$$s_k = (r_{k-1} - r_k)/(q_k - q_{k-1});$$

discard points k such that $s_k \leq s_{k+1}$. If any points have been discarded, renumber the points and repeat Step 3.

4. The points left in the set are the vertices of the envelope. The corner of the envelope is the vertex with the largest value of (2.10) (version 1) or (2.11) (version 2).

When updating the set of points by adding new points, points already discarded in a previous step cannot become new vertices; hence it is sufficient to store the vectors x_k (and the associated (q_k, r_k)) that belong to current vertices. This is an important consideration, because the set of vertices is typically rather small, often containing only 1 or 2 elements. To keep storage requirements low, we limited the number of vertices retained to 8, overwriting—if necessary—the first or last vertex, whichever is further away from the current corner.

Using the envelope construction algorithm one can consider the following generic scheme for obtaining a regularized solution. (This generic scheme is still somewhat vague, and we discuss specific details afterwards.)

3.2 GENERIC ENVELOPE GUIDED ALGORITHM

1. Set $\lambda = 0$.
2. While the corner is the highest-numbered vertex of the envelope (in particular, while the envelope has a single vertex only) proceed as follows:
   (a) Take a step with your favorite optimizer for (2.3).
   (b) Determine the envelope of all the points seen thus far and find its corner.
3. Eliminate all points before the corner from consideration.
4. While the corner of the envelope is near (in the tail strategy discussed in detail in the following: at or next to) the left endpoint of the envelope proceed as follows:
   (a) Apply a few (we used three) steps of an optimization procedure to the new f in (2.3). For each iterate add the pair (q, r) to the list of points defining the current L-curve, and update the envelope and its corner.
   (b) Adjust λ.

The algorithm terminates when a limit on the number of iterations is reached, and returns the corner of the final L-curve as approximate solution of the regularization problem.

In Step 4(b) one would wish to choose λ to decrease both r and q. Locally, at x_k we are assured that q would decrease if

$$\nabla q(x_k)^T(-\nabla r(x_k) - \lambda \nabla q(x_k)) < 0$$

which suggests that

$$\lambda > \lambda_{\min} = \max\left(\epsilon, -\frac{\nabla q(x_k)^T \nabla r(x_k)}{\nabla q(x_k)^T \nabla q(x_k)}\right), \tag{3.1}$$

where ϵ is the machine precision. Note that if $\lambda_{\min} > \epsilon$, then at x_k, $\lambda_{\min}$ also minimizes

$$g(\lambda) = \|\nabla r(x_k) + \lambda \nabla q(x_k)\|^2. \tag{3.2}$$

Hence, whenever we are close to the Tikhonov L-curve we can expect $(q_{\lambda_{\min}}, r_{\lambda_{\min}})$ to be a point on the Tikhonov L-curve close to (q_k, r_k). Locally, we are assured that r would decrease if

$$\nabla r(x_k)^T(-\nabla r(x_k) - \lambda \nabla q(x_k)) < 0$$

which suggests that

$$\lambda < \lambda_{\max} = -\frac{\nabla r(x_k)^T \nabla r(x_k)}{\nabla q(x_k)^T \nabla r(x_k)}. \tag{3.3}$$

Because the gradients of q and r are already needed to compute the gradient of the objective function, required in most optimization procedures, the cost for computing $\lambda_{\max}$ and $\lambda_{\min}$ is negligible. Note that often $\lambda_{\min}$ is smaller than $\lambda_{\max}$ by many orders of magnitude.

The most robust case of the general scheme, which we call the *tail strategy*, is similar to the adaptive annealing algorithm in Symonds, Han, Santago, and Snyder (1994). One can use the envelope as a gatekeeper and begin every major iteration in Step 4(b) with x from the previous iteration. Since, in many applications, the major features in the solution are defined by the end of Step 2 but noise still slightly obscures the signal, one does not have to be that careful with the choice of λ. At the first encounter of step 4(b) we set λ to the geometric mean of $\lambda_{\min}$ and $\lambda_{\max}$ from (3.1) and (3.3); that is, we set λ to $\hat{\lambda} = (\lambda_{\max}\lambda_{\min})^{1/2}$ and take three steps of the PCG algorithm. At subsequent encounters of Step 4(b), the current value of λ, λ_k, is reset to λ_{k+1} by the following scheme

$$\lambda_{k+1} = \begin{cases} \min(\sigma\lambda_k, .5(\lambda_k + \lambda_{\max})) & \text{if } (q_k, r_k) \text{ is below the corner;} \\ \max(\lambda_k/2, .5(\lambda_k + \lambda_{\min})) & \text{if } (q_k, r_k) \text{ is above the corner;} \\ \lambda_k & \text{if } (q_k, r_k) \text{ is the corner of the envelope.} \end{cases} \tag{3.4}$$

For the tail strategy, σ in (3.4) was set to 4.

A less robust strategy and one more sensitive to the choice of λ begins every major iteration in Step 4(b) with x as the corner of the last envelope if the envelope has more than two points. We call this the *corner* strategy in our graphs. With the corner strategy, the features in the data may not be that well defined at the first iteration of Step 4 and one has to be more careful that λ does not become too large too quickly. Our computational experience suggested that the first nonzero value of λ should be $\lambda_{\min}$ from (3.1). On subsequent encounters of Step 4(b), the strategy of adjusting λ given in (3.4) was used with σ set to 2. Note that the formula for $\lambda_{\min}$ is based on just a first order estimate. We found that for asymptotically sublinear penalty functions, that after changing λ the next point the (q, r) plot was sometimes interior to the envelope because this iterate was a simple steepest descent step rather than a conjugate gradient step and the estimate of $\lambda_{\min}$ was too low. Our experience suggested restarting with the x corresponding to the next point in the envelope with a lower value of r was superior to trying to guess the correct λ. Termination of the whole algorithm for the corner strategy depends on whether the corner position has changed significantly. The corner strategy is much more sensitive to wrong decisions and for this reason the authors prefer the tail strategy.

Step 3 was included in the generic algorithm to eliminate the possibility that the final corner would be at or above the first corner, a problem that had occurred in practice.

The generic algorithm assumes that it will take several iterations to reduce the residual significantly and one would wish to decrease it further once the initial corner is found. For problems in image restoration, where A is the identity matrix, after one iteration with the conjugate gradient technique with $\lambda = 0$, the residual is probably already too small. In this case one may wish to begin with $\lambda = 1$ and to skip Step 3.

4. APPLICATION OF THE ENVELOPE ALGORITHM TO PET

The new algorithm has been applied successfully to a problem with projection data from positron emission tomography (PET), with several different penalty functions.

In positron emission tomography (PET), the patient is given a tagged substance that emits positrons. Each positron annihilates with an electron and emits two photons in opposite directions. The patient is surrounded by a ring of detectors, which are wired so that whenever any pair of detectors senses a photon within a very small time interval, the size of which is system-dependent, the count for that pair is incremented. In a matter of minutes several million photon pairs may be detected. The reconstruction problem in PET is to determine a map of the brain by calculating the distribution of the annihilations from the data gathered by the ring of detectors.

Assume one has imposed a grid of boxes on the affected organ and tries to compute the unknown number x_B of annihilations in the box B. Let b_d represent the number of photon pairs detected in tube d and let the matrix entry $A_{d,B}$ represent the probability that an emission in box B is detected in tube d. Then the desired density x satisfies the approximate relation $Ax \approx b$. The problem also has a nonnegativity constraint.

Because of the size of the A matrix, iterative methods must be used to minimize $r(x) := \|Ax - b\|_2^2$. Unfortunately, these techniques are all plagued by a snowy image problem because of a low signal to error ratio that has been ameliorated using regularization.

In this article we consider two candidates for $q(x)$ namely the (convex, quadratic) function

$$q(x) = \sum_j \left(\frac{1}{8} \sum_{i \sim j} x_i - x_j \right)^2, \tag{4.1}$$

where $i \sim j$ denotes the statement that i is adjacent to j (i.e., is one of the eight neighbors of j on the grid), and the convex, nonquadratic penalty term

$$q(x) = \sum_j \sum_{i \sim j} H(x_i - x_j, \delta), \tag{4.2}$$

where is H is the piecewise quadratic-linear function

$$H(d, \delta) = \begin{cases} d^2 & \text{if } |d| < \delta, \\ 2\delta|d| - \delta^2 & \text{otherwise,} \end{cases} \tag{4.3}$$

introduced by Huber (1981) in the context of robust regression. Additional penalty functions were explored by Kaufman and Neumaier (1994).

We created an idealized problem to show the efficacy of our technique. We generated the matrix A according to the angle of view technique of Shepp and Vardi (1982) using a 128×128 grid and 128 detectors, resulting in $\binom{128}{2}$ approximate equations for 128^2 variables; but roughly a quarter of the latter (the values outside a circle inscribed to the box defining the grid) were fixed at zero. A right-hand side was constructed by simulating 1 million random annihilations, yielding the image $\hat{x}$ given in the top left corner of Figure

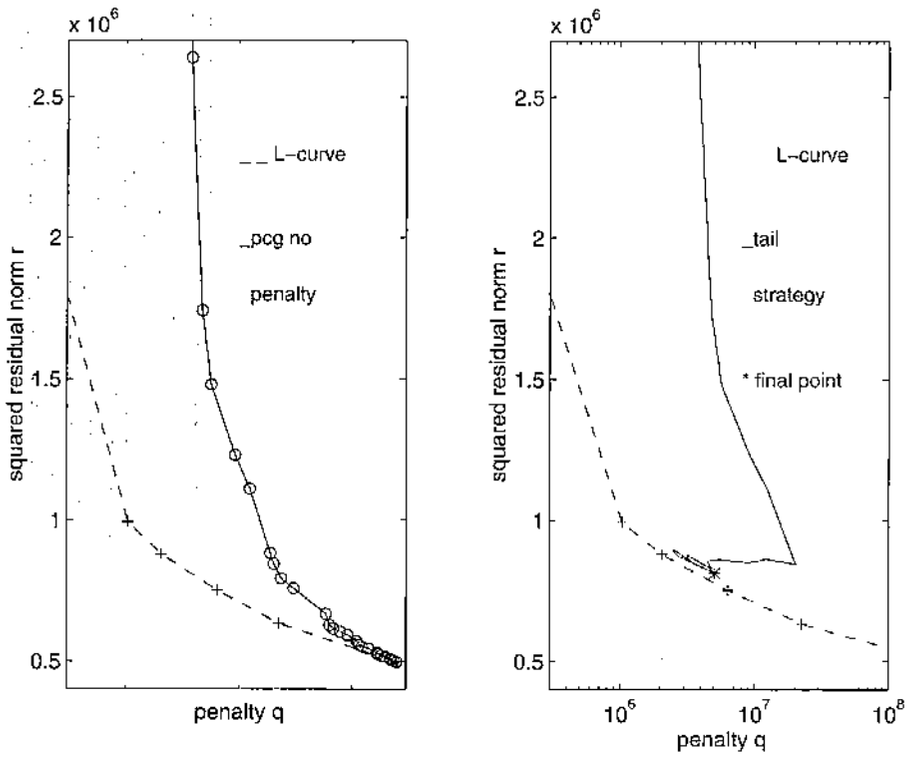

Figure 1. Tikhonov L-curve and Iteration History for the Quadratic Penalty Function with 1 million Photon Pairs.

2. Note that A and b were constructed independently and in fact $||A\hat{x} - b||_2$ was roughly the same order of magnitude as $||b||_1$.

The number of photon pairs used in the simulation determines the noise level; it decreases with an increasing number of recorded photon pairs. The randomness in choosing the direction of the annihilating photons and the inaccuracy resulting from the discretization used to set up the matrix barely allows the reconstruction of the small features (representing idealized tumors) for the 1 million photon case.

For our experiments, the sets of points whose envelope defines the current L-curves were generated by the preconditioned conjugate gradient method (PCG) of Kaufman (1993) that successfully handles nonnegativity constraints. The starting vector for the first optimization was zero outside the inscribed circle of the box defining the grid and uniform on and within that circle. It was normalized so the total sum of entries equal to the number of photon pairs counted.

The experiments reported here all use version 1 as the corner criteria. Comparison with version 2 are given in Kaufman and Neumaier (1994). Figure 1 gives a detailed view of the progress of the unregularized PCG algorithm and the tail strategy for (4.1) by displaying the curves in the (q, r)-plane defined by the iterates and the (approximate) Tikhonov L-curve whose corner is the desired target. We used a maximum of 32 PCG iterations.

First, observe that the Tikhonov curve is approximately L-shaped, so that the main

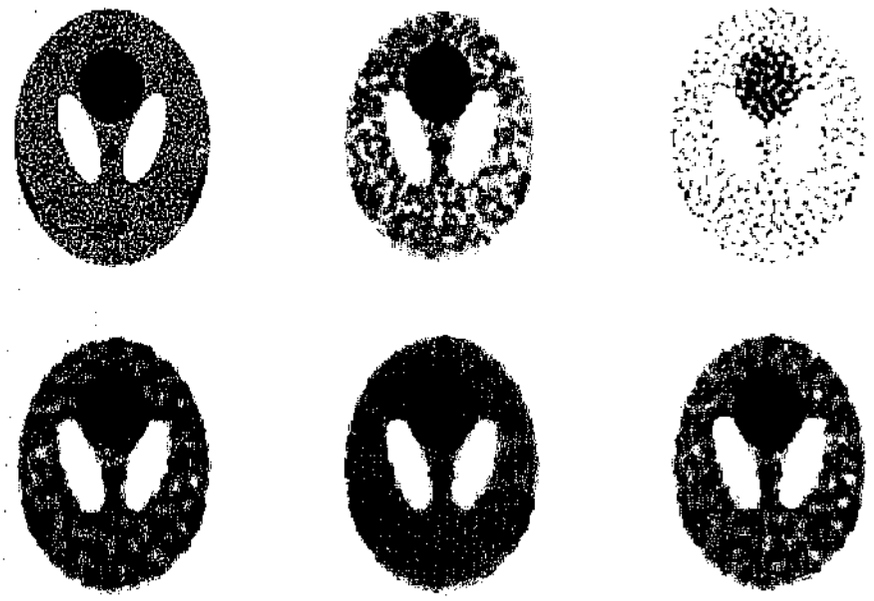

Figure 2. Phantom, and no Penalty After 8 (best) and 33 Iterations (top row); Tail Strategy, Quadratic Penalty, Huber $\delta = 16$, Huber $\delta = 64$ (bottom row).

assumption of the envelope technique holds. Second, observe that the curve defined by the points of the PCG iteration alone is further from the corner of the L-curve than the final corner of the tail strategy.

The effectiveness of the algorithms is shown in Figure 2. To simulate the effect of a posteriori image processing we transformed the components of the final vector linearly so that they span the range $[1, 256]$, and then replaced all pixel values < 100 by 100 and all pixel values > 200 by 200, and rescaled the new gray values to fill the interval [1,256] again.

Figure 2 shows the images computed in various ways for the 1 million phantom case. The top row displays the phantom used (for comparison), the optimal approximation—with minimal Euclidean distance to the phantom—computed by the PCG method without regularization (after 8 iterations), and the deterioration resulting from continuing the unregularized iteration until iteration 33. The second row shows the results of the tail strategy with the quadratic penalty function and the Huber function with two different values of δ. The images for the Huber function were sensitive to the value of δ. Obviously, the tail strategy is an improvement over iterating the PCG method without regularization.

Thus, we have shown that the envelope guided conjugate gradient strategy of Section 3 has been successfully applied to a problem. It is inexpensive, robust, and effective.

ACKNOWLEDGMENTS

This work was done while the second author was at AT&T Bell Laboratories, Murray Hill, NJ. 1991 MSC Classification: primary 65F10, secondary 65R30, 90C29.

[Received December 1995. Revised September 1996.]

REFERENCES

Engl, H.W. (1993), "Regularization Methods for the Stable Solution of Inverse Problems," *Surveys of Mathematics in Industry*, 3, 71–143.

Geman, S., and McClure, D. (1985), "Bayesian Image Analysis: An Application to Single Photon Emission Tomography," in *Proceedings of the Statistical Computing Section*, Washington, DC: American Statistical Association, pp. 12–18.

Green, P. (1990), "Bayesian Reconstructions from Emission Tomography Using a Modified EM Algorithm," *IEEE Transactions on Medical Imaging*, 9, 84–93.

Hanke, M., and Hansen, P.C. (1993), "Regularization Methods for Large-Scale Problems," *Surveys of Mathematics in Industry*, 3, 253–315.

Hansen, P.C. (1992), "Analysis of Discrete Ill-Posed Problems by Means of the L-Curve," *SIAM Review*, 34, 561–580.

Hansen, P.C., and O'Leary, D.P. (1993), "The Use of the L-Curve in the Regularization of Discrete Ill-Posed Problems," *SIAM Journal of Scientific Computing*, 14, 1487–1506.

Hebert, T., and Leahy, R. (1990), "A Generalized EM Algorithm for the 3-D Bayesian Reconstruction from Poisson Data Using Gibbs Priors," *IEEE Transactions on Medical Imaging*, 8, 194–202.

Honerkamp, J., and Weese, J. (1990), "Tikhonov Regularization Methods for Ill-Posed Problems," *Continuum Mechanics Thermodynamics*, 2, 17–30.

Huber, P.J. (1981), *Robust Statistics*, New York: John Wiley.

Kaufman, L. (1993), "Maximum Likelihood, Least Squares, and Penalized Least Squares for PET," *IEEE Transactions on Medical Imaging*, 12, 200–214.

Kaufman, L., and Neumaier, A. (1994), "Image Reconstruction Through Regularization by Envelope Guided Conjugate Gradients," Technical Report CS94/4-14, AT&T Bell Laboratories, Murray Hill, NJ (available through the WWW, http://cm.bell-labs.com/cm/cs/doc/94/4-14.ps.gz, 2MB).

Lalush, D.S., and Tsui, B.M.W. (1992), "Simulation Evaluation of Gibbs Prior Distributions for Use in Maximum A Posteriori SPECT Reconstructions," *IEEE Transactions Medical Imaging*, 11, 267–275.

Lawson, C.L., and Hanson, R.J. (1974), *Solving Least Squares Problems*, Englewood Cliffs, NJ: Prentice-Hall.

Morozov, V.A. (1984), *Methods for Solving Incorrectly Posed Problems*, Berlin: Springer-Verlag.

Sawaragi, Y., Nakayama, H., and Tanino, T. (1985), *Theory of Multiobjective Optimization*, Orlando, FL: Academic Press.

Symonds, M., Han, Y.-S., Santago, P., and Snyder, W. (1994), "Reconstruction of Positron Emission Tomography Images by Using Maximum a Posteriori and Mean Field Annealing Techniques," *Physics of Medical Imaging. Proceedings of SPIE*, Vol. 2613, pp. 212–222.

Shepp, L., and Vardi, Y. (1982), "Maximum Likelihood Reconstruction in Positron Emission Tomography," *IEEE Transactions on Medical Imaging*, 1, 113–122.

Tikhonov, A.N. (1963a), "Solution of Incorrectly Formulated Problems and the Regularization Method," *Soviet Mathematics Doklady*, 4, 1035–1038.

——— (1963b), "Regularization of Incorrectly Posed Problems," *Soviet Mathematics Doklady*, 4, 1624–1627.

Vogel, C.R., and Oman, M.E. (1996), "Iterative Methods for Total Variation Denoising," *SIAM Journal on Statistical Computing 17*, 227–238.