

THE STABILIZING PROPERTIES OF NONNEGATIVITY
CONSTRAINTS IN LEAST-SQUARES
IMAGE RECONSTRUCTION

Johnathan M. Bardsley¹ §, Jorma K. Merikoski², Roberto Vio³

¹Department of Mathematical Sciences
University of Montana
Missoula, MT 59812-0864, USA
e-mail: bardsleyj@mso.umt.edu

²Department of Mathematics and Statistics
University of Tampere
Tampere, FI-33014, FINLAND
e-mail: jorma.merikoski@uta.fi

³Chip Computers Consulting s.r.l.
82, Viale Don L. Sturzo 82, S. Liberale di Marcon, Venice, 30020, ITALY
e-mail: robertovio@tin.it

Abstract: The incorporation of nonnegativity constraints in image reconstruction problems is known to have a stabilizing effect on solution methods. In this paper, we both demonstrate and provide an explanation of this phenomena when the image reconstruction problem of interest has least squares form. The benefits of using this natural constraint suggest the importance of incorporating *a priori* knowledge about solutions when possible. In fact, if this prior information is significantly strong, sophisticated likelihood functions and computational methods may not be necessary.

AMS Subject Classification: 65F20

Key Words: nonnegativity constraints, astronomical imaging, least squares

Received: December 10, 2007

© 2008, Academic Publications Ltd.

§Correspondence author

1. Introduction

In this paper we consider real $m \times n$ linear systems of the form

$$\mathbf{Ax} = \mathbf{b}, \tag{1}$$

where $m \geq n$, \mathbf{A} has full column rank, and subject to the nonnegativity constraints

$$x_i \geq 0, \quad i = 1, \dots, n. \tag{2}$$

In the sequel, we will use $\mathbf{x} \geq \mathbf{0}$ to denote (2). Such a model arises naturally in image reconstruction, where \mathbf{x} denotes the measured intensities, which are nonnegative, at various pixels of an object being viewed by an imaging system. In such applications, the matrix \mathbf{A} is typically ill-conditioned.

In practice, the $m \times 1$ vector \mathbf{b} in (1) is contaminated with noise, and hence, is the realization of a random process. We assume \mathbf{b} is formed via the discrete statistical model

$$\mathbf{b} = \mathbf{Ax}_{\text{true}} + \mathbf{N}, \tag{3}$$

where \mathbf{x}_{true} is the *object*, or *true image*, and \mathbf{N} is the $m \times 1$ random vector corresponding to the noise.

A standard approach for estimating \mathbf{x}_{true} from \mathbf{b} is to compute a regularized approximate solution of the least squares minimization problem

$$\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|^2, \tag{4}$$

where $\|\cdot\|$ denotes the standard ℓ^2 norm. Regularized approximate solution methods are typically needed due to the fact that \mathbf{A} is ill-conditioned. We will not discuss regularization methods in the paper; for such a discussion see [2], [9].

Instead, our objective is to analyze the effects of replacing (4) by the non-negatively constrained least squares problem

$$\min_{\mathbf{x} \geq \mathbf{0}} \|\mathbf{Ax} - \mathbf{b}\|^2. \tag{5}$$

In particular, we will argue that one can expect (5) to be a more stable problem than (4). As a result, the exact solution of (5) will typically be of a higher resolution than the exact solution of (4), and hence, though regularization will often still be needed in (5), the regularized approximate solution can be expected to be of a higher resolution.

The paper is organized as follows. We will begin in Section 2 with a derivation of a linear system whose minimum norm solution is the unique solution of (5). We then show in Section 3 that the condition number of this matrix

can be expected to be much smaller than that of the original matrix \mathbf{A} . Taken together, these results will suggest that nonnegativity constraints provide stability in least squares image reconstruction. We demonstrate this in a number of numerical experiments, which are performed in the Section 4. We end with conclusions in Section 5.

2. Nonnegatively Constrained Least Squares

We begin by making a few definitions. First, we say that \mathbf{x} is *feasible* if $\mathbf{x} \geq \mathbf{0}$, and we define the *active set* of any feasible \mathbf{x} by

$$\mathcal{A}(\mathbf{x}) = \{i \mid x_i = 0\}. \quad (6)$$

We can then define the diagonal matrix $\mathbf{D}(\mathbf{x})$ for $\mathbf{x} \geq \mathbf{0}$ by

$$[\mathbf{D}(\mathbf{x})]_{ii} = \begin{cases} 1, & i \notin \mathcal{A}(\mathbf{x}), \\ 0, & i \in \mathcal{A}(\mathbf{x}). \end{cases} \quad (7)$$

We now focus our attention on (5). Suppose that \mathbf{x}^* is a local solution of (5) and define $\mathbf{D}^* = \mathbf{D}(\mathbf{x}^*)$. The Karush-Kuhn-Tucker (KKT) conditions (see [8]) for (5) have the form

$$[\mathbf{A}^T \mathbf{A} \mathbf{x}^* - \mathbf{A}^T \mathbf{b}]_i \geq 0, \quad i = 1, \dots, n, \quad (8)$$

$$x_i^* \cdot [\mathbf{A}^T \mathbf{A} \mathbf{x}^* - \mathbf{A}^T \mathbf{b}]_i = 0, \quad i = 1, \dots, n. \quad (9)$$

We have from (9) that for all i such that $x_i^* > 0$,

$$[\mathbf{A}^T \mathbf{A} \mathbf{x}^* - \mathbf{A}^T \mathbf{b}]_i = 0. \quad (10)$$

This implies, using the fact that $\mathbf{D}^* \mathbf{x}^* = \mathbf{x}^*$, that

$$\mathbf{D}^* \mathbf{A}^T \mathbf{A} \mathbf{D}^* \mathbf{x}^* - \mathbf{D}^* \mathbf{A}^T \mathbf{b} = \mathbf{0}, \quad (11)$$

which are the normal equations for

$$\min_{\mathbf{x}} \|\mathbf{A} \mathbf{D}^* \mathbf{x} - \mathbf{b}\|^2. \quad (12)$$

Thus \mathbf{x}^* is the minimum norm solution of (12); that is,

$$\mathbf{x}^* = (\mathbf{A} \mathbf{D}^*)^\dagger \mathbf{b}, \quad (13)$$

where “ \dagger ” denotes psuedo-inverse, which we define in the next section.

In practice one does not solve (12) or compute (13) since \mathbf{D}^* is not known *a priori*. Nonetheless, a careful consideration of (12) and the associated linear system

$$\mathbf{A} \mathbf{D}^* \mathbf{x} = \mathbf{b} \quad (14)$$

will be a key component our discussion.

3. The Condition Numbers of \mathbf{A} and \mathbf{AD}^* .

Our objective in this section is to investigate and compare the stability of linear systems (1) and (14). For this we need to define the *singular value decomposition* (SVD) of a matrix. The SVD of \mathbf{A} (see, e.g., [3], [4]) is given by

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad \mathbf{\Sigma} = \begin{pmatrix} \mathbf{\Sigma}_1 \\ \mathbf{0} \end{pmatrix}, \quad (15)$$

where \mathbf{U} is an $m \times m$ orthogonal matrix, \mathbf{V} is an $n \times n$ orthogonal matrix, and $\mathbf{\Sigma}$ is $m \times n$ with $\mathbf{\Sigma}_1 = \text{diag}(s_1, s_2, \dots, s_n)$ and $s_1 \geq s_2 \geq \dots \geq s_n$. We call the s_i 's the *singular values* of \mathbf{A} and note that $s_n > 0$ since \mathbf{A} has full column rank.

With the SVD in hand, we can now make two important definitions. The *pseudo-inverse* of \mathbf{A} is given by

$$\mathbf{A}^\dagger = \mathbf{V}\mathbf{\Sigma}^\dagger\mathbf{U}^T, \quad \mathbf{\Sigma}^\dagger = \begin{pmatrix} \mathbf{\Sigma}_1^{-1} & \mathbf{0} \end{pmatrix}.$$

The *condition number* of \mathbf{A} is defined

$$\text{cond } \mathbf{A} = s_1/s_n. \quad (16)$$

We say that \mathbf{A} is *ill-conditioned* if $\text{cond } \mathbf{A}$ is extremely large, which occurs, we assume, due to the fact that $0 < s_n \ll 1$.

3.1. An Inequality Involving $\text{cond } \mathbf{A}$ and $\text{cond } \mathbf{AD}^*$

We begin with a theoretical result to this effect. For this we will need to survey some basic facts about singular values (for a thorough exposition, see, e.g., [3], [4]). Let \mathbf{B} be an arbitrary real $m \times n$ matrix with full column rank, and let $s_i(\mathbf{B})$ denote the i -th largest singular value of \mathbf{B} . Then

$$\max_{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|} = s_1(\mathbf{B}), \quad \min_{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^n} \frac{\|\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|} = s_n(\mathbf{B}),$$

and hence, (see [1], Exercise 6.15)

$$\|\mathbf{B}\| = s_1(\mathbf{B}), \quad \text{and} \quad \|\mathbf{B}^\dagger\| = 1/s_n(\mathbf{B}).$$

Thus, consistent with (16), we have

$$\text{cond } \mathbf{B} = s_1(\mathbf{B})/s_n(\mathbf{B}) = \|\mathbf{B}\| \|\mathbf{B}^\dagger\|.$$

Let us return to our original task. There exists a permutation matrix \mathbf{Q} such that the first r diagonal entries of the diagonal matrix $\mathbf{D}' = \mathbf{Q}^T \mathbf{D}^* \mathbf{Q}$ are one, with the remaining diagonal entries zero. We define $\mathbf{A}' = \mathbf{A}\mathbf{Q}$. Then the singular values of $\mathbf{A}'\mathbf{D}'$ are equal to the singular values of \mathbf{AD}^* . Therefore,

without loss of generality, we may assume that

$$\mathbf{D}^* = \begin{pmatrix} \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix},$$

where \mathbf{I}_r is the $r \times r$ identity matrix.

Consider the partition $\mathbf{A} = (\mathbf{A}_1 \ \mathbf{A}_2)$, where the first part consists of the first r columns and the second part of the remaining columns of \mathbf{A} . Then $\mathbf{AD}^* = (\mathbf{A}_1 \ \mathbf{0})$. Since \mathbf{A} has full column rank, so does \mathbf{A}_1 , and hence \mathbf{A}_1^\dagger and $\text{cond } \mathbf{A}_1$ are defined. This allows us to define both the psuedo-inverse and condition number of \mathbf{AD}^* ; namely,

$$(\mathbf{AD}^*)^\dagger \stackrel{\text{def}}{=} \begin{pmatrix} \mathbf{A}_1^\dagger \\ \mathbf{0} \end{pmatrix},$$

and

$$\text{cond } \mathbf{AD}^* \stackrel{\text{def}}{=} \text{cond } \mathbf{A}_1.$$

We will now show that $\text{cond } \mathbf{AD}^* \leq \text{cond } \mathbf{A}$. First note that $r \leq n \leq m$. Then if $\mathbf{x} \in \mathbb{R}^r$, and we denote

$$\tilde{\mathbf{x}} = \begin{pmatrix} \mathbf{x} \\ \mathbf{0} \end{pmatrix},$$

we have

$$s_1(\mathbf{A}_1) = \max_{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^r} \frac{\|\mathbf{A}_1 \mathbf{x}\|}{\|\mathbf{x}\|} = \max_{\mathbf{0} \neq \tilde{\mathbf{x}} \in \mathbb{R}^n} \frac{\|\mathbf{A} \tilde{\mathbf{x}}\|}{\|\tilde{\mathbf{x}}\|} \leq s_1(\mathbf{A}), \tag{17}$$

and

$$s_r(\mathbf{A}_1) = \min_{\mathbf{0} \neq \mathbf{x} \in \mathbb{R}^r} \frac{\|\mathbf{A}_1 \mathbf{x}\|}{\|\mathbf{x}\|} = \min_{\mathbf{0} \neq \tilde{\mathbf{x}} \in \mathbb{R}^n} \frac{\|\mathbf{A} \tilde{\mathbf{x}}\|}{\|\tilde{\mathbf{x}}\|} \geq s_n(\mathbf{A}). \tag{18}$$

Therefore

$$\text{cond } \mathbf{AD}^* = \frac{s_1(\mathbf{A}_1)}{s_r(\mathbf{A}_1)} \leq \frac{s_1(\mathbf{A})}{s_n(\mathbf{A})} = \text{cond } \mathbf{A}. \tag{19}$$

Thus the conditioning of \mathbf{AD}^* can be no worse than that of \mathbf{A} .

In astronomical imaging examples, a strong case can be made that an inequality stronger than (19) will often hold. Due to the fact that there is typically a substantial black background in the astronomical objects being viewed, r is often much smaller than n . In such cases, the vectors $\tilde{\mathbf{x}}$ will cover a small subspace of \mathbb{R}^n . Since the maximum (respectively minimum) of a function over a big set is often times much larger (respectively smaller) than that over a small subset, it is not unusual for $s_1(\mathbf{A}_1) \ll s_1(\mathbf{A})$ in (17) and/or $s_r(\mathbf{A}_1) \gg s_n(\mathbf{A})$ in (18), which yields $s_1(\mathbf{A}_1)/s_r(\mathbf{A}_1) \ll s_1(\mathbf{A})/s_n(\mathbf{A})$, and hence, (19) can be

replaced by

$$\text{cond } \mathbf{A}\mathbf{D}^* \ll \text{cond } \mathbf{A}. \quad (20)$$

When (20) holds, one can expect that the solutions of (5) will be much more stable with respect to errors in \mathbf{b} than the solutions of (1). The fact that (20) does in fact hold in practical examples is confirmed in our numerical experiments. First, however, we discuss the connection between condition number and stability.

3.2. The Condition Number and Stability

Recall that we have assumed that for \mathbf{A} to have a large condition number, it must be that $0 < s_n \ll 1$. Furthermore, for our discussion, and without loss of generality, we assume that $s_1 = 1$. And finally, we make the assumption that the distribution of the singular values between s_1 and s_n is continuous, i.e. without jumps. We note that these assumptions are all reasonable in astronomical imaging.

We want to determine what the condition number can tell us about the sensitivity of $\mathbf{A}^\dagger \mathbf{b}$ to perturbations in the noise vector \mathbf{N} . To do this, we first assume that our data \mathbf{b} satisfies (3) exactly. Then

$$\mathbf{A}^\dagger \mathbf{b} = \mathbf{x}_{\text{true}} + \mathbf{A}^\dagger \mathbf{N}.$$

Using the SVD of \mathbf{A} given above we obtain

$$\mathbf{A}^\dagger \mathbf{N} = \sum_{i=1}^n \left(\frac{\mathbf{u}_i^T \mathbf{N}}{s_i} \right) \mathbf{v}_i, \quad (21)$$

where \mathbf{u}_i and \mathbf{v}_i are the i -th column of \mathbf{U} and \mathbf{V} respectively. We note that since \mathbf{U} and \mathbf{V} are unitary, the sets $\{\mathbf{u}_i\}_{i=1}^m$ and $\{\mathbf{v}_i\}_{i=1}^n$ are orthonormal. We note that \mathbf{N} can be expressed as

$$\mathbf{N} = (\mathbf{u}_1^T \mathbf{N}) \mathbf{u}_1 + (\mathbf{u}_2^T \mathbf{N}) \mathbf{u}_2 + \cdots + (\mathbf{u}_m^T \mathbf{N}) \mathbf{u}_m,$$

and since it contains random noise, $\mathbf{u}_i^T \mathbf{N}$ will not decay to zero as $i, n \rightarrow \infty$. An appeal to (21) then immediately tells us, given our assumptions, that the values of $\{s_i\}_{i=1}^n$ will determine the sensitivity $\mathbf{A}^\dagger \mathbf{b}$ to perturbations in \mathbf{N} . Given our assumption that the s_i 's are smoothly distributed, a large value of $\text{cond} \mathbf{A}$ will likely suggest a large instability in $\mathbf{A}^\dagger \mathbf{b}$ to perturbations in \mathbf{N} .

4. Numerical Results

4.1. One Dimensional Test Problems

We consider the one-dimensional (1D) version of a model that occurs in two-dimensional (2D) astronomical imaging given by

$$b(s) = \int_0^1 A(s-s')x(s')ds'. \quad (22)$$

Here x represents light intensity as a function of spatial position, and b represents image intensity. The *point spread function* (PSF) A characterizes the optical blurring effects that occur during image formation.

A PSF that models the long-time average effects of atmospheric turbulence on light propagation is the Gaussian, see [9]. Its one dimensional version is

$$A(s) = C \exp(-s^2/2\gamma^2),$$

where C and γ are positive parameters.

Discretizing (22) yields a linear system of the form (1). If midpoint quadrature is used in the s' variable, the resulting matrix \mathbf{A} has the form

$$[\mathbf{A}]_{ij} = hC \exp\left(-\frac{((i-j)h)^2}{2\gamma^2}\right), \quad 1 \leq i, j \leq n, \quad (23)$$

where $h = 1/n$. Here, \mathbf{A} has full column rank.

We now consider some specific examples. We build our matrix \mathbf{A} using (23) with $n = 80$ and $\gamma^2 = 0.00125$. We generated our data vector \mathbf{b} using model (3), where \mathbf{x}_{true} is the object, or true image, and \mathbf{N} is a Gaussian random vector with mean 0 and standard deviation chosen so that the signal-to-noise ration (SNR) is 30.

We consider several different choices for \mathbf{x}_{true} since \mathbf{D}^* , and hence the condition number of $\mathbf{A}\mathbf{D}^*$, depends upon the object. We begin by considering objects that have the form

$$\mathbf{x}_{\text{true}} = \begin{cases} 1 & \frac{1}{2} - \rho \leq [\mathbf{x}_{\text{true}}]_i \leq \frac{1}{2} + \rho, \\ 0 & \text{otherwise,} \end{cases} \quad (24)$$

where $0 \leq \rho \leq \frac{1}{2}$. We perform experiments for four separate values of ρ : $\rho = 0.01$, $\rho = 0.04$, $\rho = 0.07$, and $\rho = 0.1$. The plots of these objects and of the corresponding blurred, noisy images are given in Figure 1. An accurate estimate of \mathbf{D}^* for each example can be obtained using the projected Newton (PN) method, see [5]. Here we compute 200 projected Newton iterations in

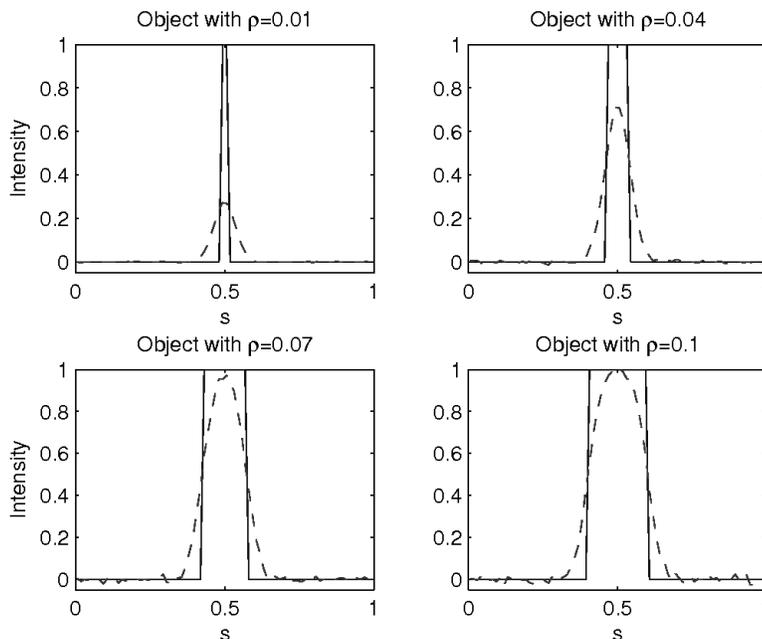


Figure 1: Plots of objects and blurred, noisy images. The plot of the discrete object \mathbf{x}_{true} has the solid line and is defined by (24). The plot of the object corresponding to $\rho = 0.01$ is given in the upper left, to $\rho = 0.04$ in the upper right, to $\rho = 0.07$ in the lower left, and to $\rho = 0.1$ in the lower right. The plots of the corresponding blurred, noisy images have the dashed line.

order to obtain our estimate for \mathbf{D}^* ; that is,

$$\mathbf{D}^* = \mathbf{D}(\mathbf{x}_{200}^{\text{PN}}),$$

where $\mathbf{x}_{200}^{\text{PN}}$ is the 200-th PN iterate and $\mathbf{D}(\mathbf{x})$ is defined in (7). From this, $\text{cond} \mathbf{A}\mathbf{D}^* = s_1/s_r$ can be computed. The values of $\text{cond} \mathbf{A}\mathbf{D}^*$, s_1 and s_r can be found in Table 1. Note that in all cases, the condition number of $\mathbf{A}\mathbf{D}^*$ is substantially less than the condition number of \mathbf{A} , which is 1.2×10^{17} . Thus inequality (20) holds in these cases.

Plots of $\mathbf{x}^* = (\mathbf{A}\mathbf{D}^*)^\dagger \mathbf{b}$, which is the unique solution of (5), are given in Figure 2. These plots should not be viewed as reconstructions. Rather, they should be compared to the exact solutions of $\mathbf{A}\mathbf{x} = \mathbf{b}$ given in Figure 3.

When this is done, it is evident that noise amplification is much less pro-

ρ	$\text{cond } \mathbf{AD}^*$	s_1	s_r
0.01	2.448×10^2	3.821×10^{-5}	1.561×10^{-7}
0.04	6.589×10^4	5.784×10^{-5}	8.780×10^{-10}
0.07	2.330×10^4	5.392×10^{-5}	2.313×10^{-10}
0.1	1.678×10^4	5.115×10^{-5}	3.048×10^{-10}

Table 1: Condition numbers of \mathbf{AD}^* for various ρ

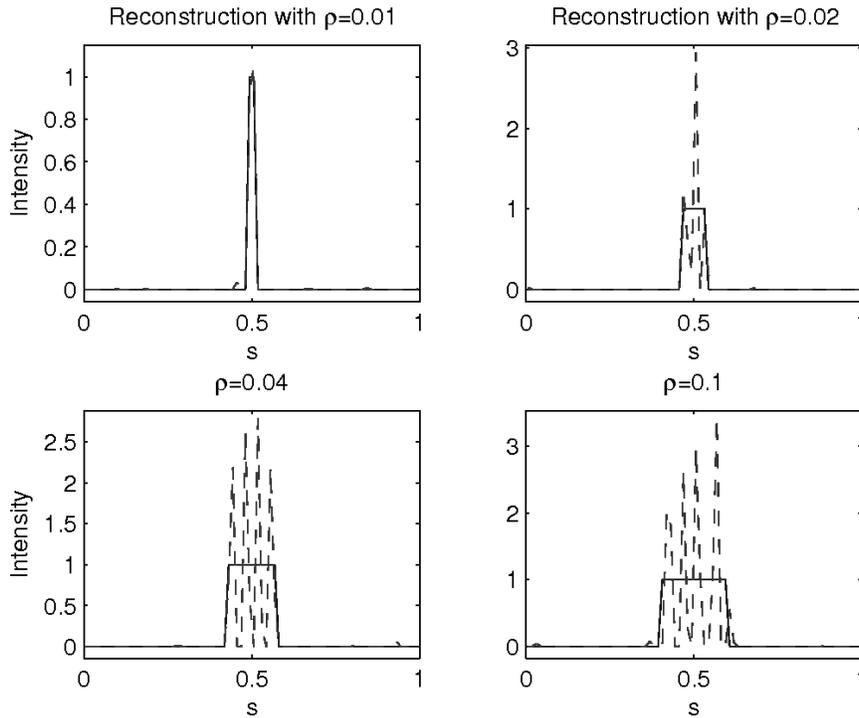


Figure 2: Plots of nonnegatively constrained reconstructions without regularization. The plot of the discrete object \mathbf{x}_{true} has the solid line. The plot of $\mathbf{x}^* = (\mathbf{AD}^*)^\dagger \mathbf{b}$ has the dashed line. The upper left-hand corner corresponds to $\rho = 0.01$, the upper right to $\rho = 0.04$, the lower left to $\rho = 0.07$, and the lower right to $\rho = 0.1$.

nounced for the solutions of $\mathbf{AD}^* \mathbf{x} = \mathbf{b}$. Nonetheless some noise amplification does occur in three of the four reconstructions, which indicates that \mathbf{AD}^* is

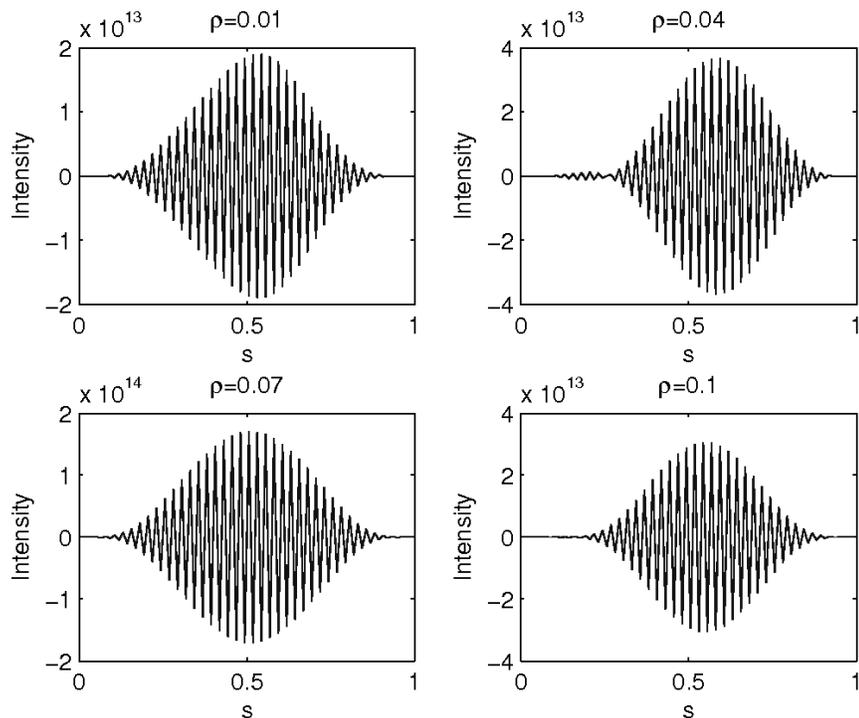


Figure 3: Plots of $\mathbf{A}^{-1}\mathbf{b}$. The upper left-hand corner corresponds to $\rho = 0.01$, the upper right to $\rho = 0.04$, the lower left to $\rho = 0.07$, and the lower right to $\rho = 0.1$.

mildly ill-conditioned in those cases. Hence, as expected, (5) requires some form regularization.

Figures 2 and 3, together with the condition number results in Table 1, provide strong evidence that nonnegativity constraints have a stabilizing effect.

We next consider an example in which the object \mathbf{x}_{true} consists of several sources. A plot of both \mathbf{x}_{true} and \mathbf{b} is given on the left in Figure 4. In this case, the object is composed of three sources, each of the type (24): one with $\rho = 0$, one with $\rho = 0.01$ and one with $\rho = 0.1$. We note that this object is indicative of what we will call a stellar-like object. Such objects are characterized by collections of point, or near-point, sources. In 2D astronomical imaging a star-field is an example of such an object. A plot of $(\mathbf{AD}^*)^\dagger\mathbf{b}$, with \mathbf{D}^* estimated as in the previous examples, is given on the right in Figure 4. Again, notice

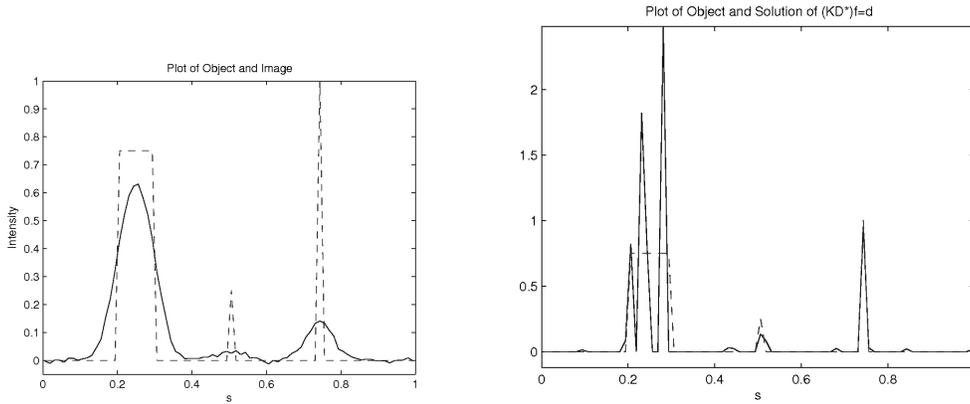


Figure 4: Plots of object, image and reconstruction. On both the left and right-hand side, the plot of the discrete object \mathbf{x}_{true} has the dashed line. On the left-hand side, the plot of the noisy, blurred image \mathbf{b} has the solid line. On the right-hand side, the plot of $\mathbf{x}^* = (\mathbf{AD}^*)^{\dagger} \mathbf{b}$ has the solid line.

that mild noise amplification occurs, suggesting the need for regularization in practice. The plot of $\mathbf{A}^{-1} \mathbf{b}$ is similar to those found in Figure 3.

The condition number of \mathbf{AD}^* for this example is given by $\text{cond } \mathbf{AD}^* \approx 1.069 \times 10^5$, which is 12 orders of magnitude smaller than $\text{cond } \mathbf{A} \approx 1.2 \times 10^{17}$. Thus, once again, (20) is satisfied.

Before continuing, we note that in the above examples \mathbf{x}_{true} has a zero background. This has the effect that many of the components of \mathbf{x}^* are also zero, and hence, \mathbf{A} and \mathbf{AD}^* are quite different. These results will not be significantly effected if the background of \mathbf{x}_{true} is effectively zero, i.e. very small when compared to the bright regions. But if \mathbf{x}_{true} has a background that is sufficiently bright, the addition of nonnegativity constraints will have very little, if any, effect.

4.2. The Two-Dimensional Case

We now consider the 2D analogue of the one-dimensional example above. The mathematical model we consider has the form

$$b(s, t) = \int_0^1 \int_0^1 A(s - s', t - t') x(s', t') ds' dt'. \tag{25}$$

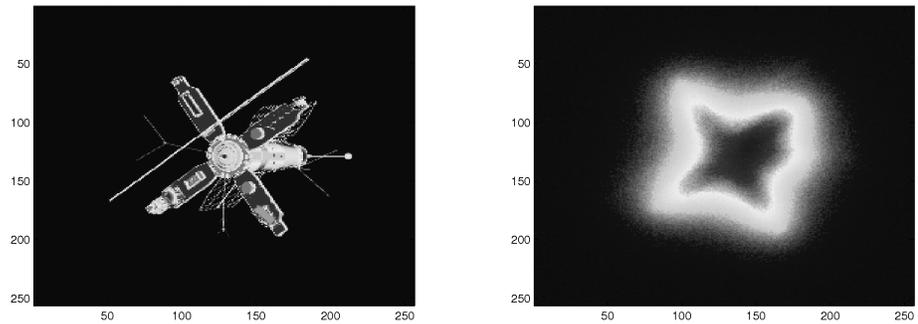


Figure 5: Satellite object on the left, and blurred, noisy data on the right.

As above, x represents the light intensity of the object as a function of spatial position, b represents the light intensity of the image, and A is the PSF, which characterizes the blurring effects of the imaging system.

In our example, we use data that was developed at the US Air Force Phillips Laboratory, Lasers and Imaging Directorate, Kirtland Air Force Base, New Mexico. The image is a computer simulation of a field experiment showing a satellite as taken from a ground based telescope. The true and blurred images have 256×256 pixels, and are shown in Figure 5. We remark that the $65,536 \times 65,536$ blurring matrix \mathbf{A} is not constructed explicitly, but is defined implicitly by the PSF. In this example, the full-width-at-half-maximum of the PSF is 13 pixels. It is not a discretized Gaussian. The resulting matrix \mathbf{A} has full column rank, and hence, the results above that required this assumption are valid for this example. The true image and PSF are contained in the *RestoreTools* image restoration package, see [7]. The data was generated with an SNR of 30.

In the 2D case, the computation of $\text{cond } \mathbf{A}\mathbf{D}^*$ is infeasible. Nonetheless, we can numerically compare the stability of the solutions of (4) and (5) by comparing the plots of the *relative solution error*, $\|\mathbf{x}_k - \mathbf{x}_{\text{true}}\|/\|\mathbf{x}_{\text{true}}\|$, for approximate solutions \mathbf{x}_k generated by comparable iterative algorithms. We know that the sensitivity of the solutions of (4) and (5) to noise in the data is closely tied to the condition numbers of \mathbf{A} and $\mathbf{A}\mathbf{D}^*$ respectively. Thus if an increase in the relative solution error occurs substantially earlier in the iterations of the method applied to (4) than for the method applied to (5), and if this increase is significantly more pronounced, i.e. if the curve is much steeper, we will conclude that the nonnegativity constraints stabilize the reconstruction

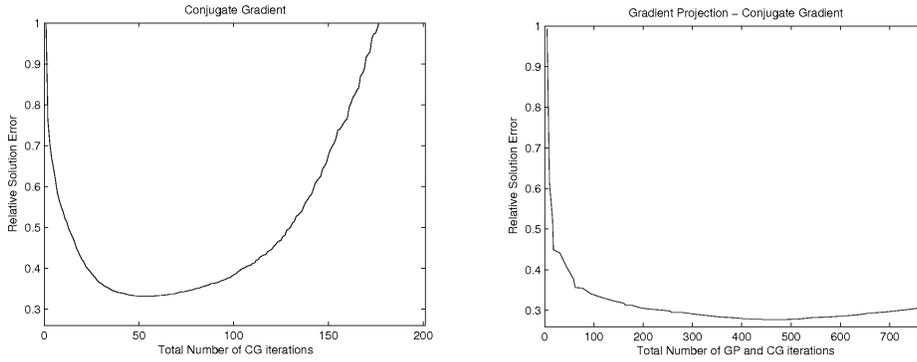


Figure 6: Plot of $\|\mathbf{x}_k - \mathbf{x}_{\text{true}}\|/\|\mathbf{x}_{\text{true}}\|$. The plot on the left is for conjugated gradient iterations. The plot on the right is for the gradient projection - conjugate gradient iterations.

problem.

For our comparison we use the conjugate gradient method (CG) (see [8]) for solving (4) and the gradient projection-conjugate gradient method (GPCG) [6] for solving (5). (Our implementation of the projected Newton method for solving (5) was not able to force an increase in the relative solution error. Thus a more robust method was needed.) We acknowledge that this method of comparison is problematic due to the fact that (4) and (5) are different problems, and CG and GPCG are different algorithms. Nonetheless, it is shown in [6] that GPCG applied to (5) will eventually reduce to CG applied to (12). Thus a comparison between these two algorithms seems natural. In any event, provided that the relative solution error curves are *significantly* different, we will assume that conclusions about problem stability can be made.

Finally, we emphasize that we are not comparing the convergence properties of CG and GPCG. Our focus, rather, is on the sensitivity of these algorithms to the noise in the data, as measured by the plots of the relative solution error curves.

The plots of the relative solution error for CG applied to (4) and for GPCG applied to (5) are given in Figure 6. For CG we plot the relative solution error versus the total number of CG iterations. For GPCG we plot the relative solution error versus the total number of gradient projection and CG iterations. An increase in the relative solution error for CG applied to (4) begins to occur between 50 and 60 iterations. For GPCG applied to (5), the increase begins after roughly 475 total GP and CG iterations. In addition, we see that the

curve is much steeper for CG than it is for GPCG. It is evident, therefore, that the CG iterations are *much* more sensitive to the noise in \mathbf{b} than is GPCG. Thus we conclude that the nonnegatively constrained problem is more stable.

5. Conclusions

Linear systems of the form (1) that arise in image reconstruction are inherently ill-conditioned. Exact solutions of such systems are often unusable due to the fact that they are unstable with respect to errors in the collected image \mathbf{b} .

As a naive solution method, we first considered (4), and then we proceeded to compare its stability with that of its nonnegatively constrained counterpart (5). The solutions of these two problems, given that \mathbf{A} has full column rank, are given by $\mathbf{A}^\dagger \mathbf{b}$ and $(\mathbf{A}\mathbf{D}^*)^\dagger \mathbf{b}$ respectively.

We then proved that $\text{cond } \mathbf{A}\mathbf{D}^* \leq \text{cond } \mathbf{A}$, but we argued heuristically that one can expect the much stronger inequality $\text{cond } \mathbf{A}\mathbf{D}^* \ll \text{cond } \mathbf{A}$ to hold for astronomical imaging examples. This was supported by our numerical experiments. For several one-dimensional examples, we were able to show explicitly that $\text{cond } \mathbf{A}\mathbf{D}^* \ll \text{cond } \mathbf{A}$.

Although the connection between the condition number of \mathbf{A} and $\mathbf{A}\mathbf{D}^*$ and the stability of the corresponding solutions is not direct, we argued that in astronomical imaging, where the singular values tend to be smoothly distributed, a large value for the condition number will likely indicate an instability in the solution with respect to the noise, which is assumed to be random in nature.

In our two-dimensional example, we could not compute $\mathbf{A}\mathbf{D}^*$ explicitly, but we provide a convincing argument, via the use of iterative methods, that for the example considered, the incorporation of nonnegativity constraints stabilizes the reconstruction problem.

The benefits of using a nonnegativity constraint suggest the importance of incorporating *a priori* knowledge about solutions when possible. In fact, if this the prior information is significantly strong, sophisticated likelihood functions and computational techniques may be unnecessary.

Finally, due to the fact that weighted least squares problems often arise in image restoration, it is natural to ask whether or not the results of this paper will extend to such problems. The theoretical arguments and discussion certainly will, since the weighted least squares problem

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_{\mathbf{W}}^2, \quad (26)$$

where \mathbf{W} is a symmetric, positive definite, weight matrix, can be written in the standard least squares form

$$\min_{\mathbf{x}} \|\tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}}\|^2,$$

where $\tilde{\mathbf{A}} = \mathbf{W}^{1/2}\mathbf{A}$ and $\tilde{\mathbf{b}} = \mathbf{W}^{1/2}\mathbf{b}$. A thorough numerical investigation would need to be done, though, in order to verify that in practice the inclusion of nonnegativity constraints in (26) has a stabilizing effect. Preliminary numerical experiments suggest that this is the case.

References

- [1] A. Ben-Israel, T.N.E. Greville, *Generalized Inverses: Theory and Applications*, Wiley (1974).
- [2] P.C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*, SIAM, Philadelphia (1998).
- [3] R.A. Horn, C.R. Johnson, *Matrix Analysis*, Cambridge Univ. (1985).
- [4] R.A. Horn, C.R. Johnson, *Topics in Matrix Analysis*, Cambridge Univ. (1991).
- [5] C.T. Kelley, *Iterative Methods for Optimization*, SIAM, Philadelphia (1999).
- [6] J.J. Moré, G. Toraldo, On the solution of large quadratic programming problems with bound constraints, *SIAM Journal on Optimization*, **1** (1991), 93-113.
- [7] J.G. Nagy, K. Palmer, L. Perrone, Iterative methods for image restoration: A Matlab object oriented approach, *Numerical Algorithms*, **36** (2003), 73-93.
- [8] J. Nocedal, S.J. Wright, *Numerical Optimization*, Springer (1999).
- [9] C.R. Vogel, *Computational Methods for Inverse Problems*, SIAM, Philadelphia (2002).

