



IMAGE DEBLURRING: I CAN SEE CLEARLY NOW

By James G. Nagy and Dianne P. O'Leary

INVERSE PROBLEMS ARE AMONG THE MOST CHALLENGING COMPUTATIONS IN SCIENCE AND ENGINEERING BECAUSE THEY INVOLVE DETERMINING THE PARAMETERS OF A SYSTEM THAT

is only observed indirectly. For example, we might have a spectrum and want to determine the species that produced it as well as their relative proportions. Or we may have sonar measurements of a containment tank and want to know whether it has an internal crack.

Here is this issue's homework assignment: given a blurred image and a linear model for the blurring, reconstruct the original image. This linear inverse problem illustrates the impact of *ill-conditioning* on the choice of algorithms.

Ill-Conditioning

Consider a linear system of equations

$$Kf = g$$

Tools

The major tool used in this project is the singular value decomposition of a matrix.² Any real matrix A of dimension $m \times n$ (with $m \geq n$) has a representation as

$$A = U\Sigma V^T$$

where $U^T U = I$, $V^T V = I$, and Σ has nonnegative entries σ_i ($i = 1, \dots, n$) on its main diagonal and zeros elsewhere.

The matrix U is $m \times m$, V is $n \times n$, and Σ is $m \times n$. The singular values σ_i are the square roots of the eigenvalues of $A^T A$, and the columns of V are the eigenvectors of that matrix. The columns of U are the eigenvectors of AA^T . Computation of the singular value decomposition is more stable than forming $A^T A$ and computing the eigendecomposition.

where K is an $n \times n$ matrix, and f and g are vectors. Suppose K is scaled so that its largest singular value is $\sigma_1 = 1$. If the smallest singular value is $\sigma_n \approx 0$, then K is *ill-conditioned*. We distinguish two types of ill-conditioning:

- The matrix K is considered numerically rank deficient if there is a j such that $\sigma_j \gg \sigma_{j+1} \approx \dots \approx \sigma_n \approx 0$. That is, there is an obvious gap between large and small singular values.
- If the singular values decay to zero with no particular gap in the spectrum, we say the linear system $Kf = g$ is a *discrete ill-posed problem*.

Computing accurate approximate solutions of discrete ill-posed problems is extremely difficult, especially because in most real applications, g is not known exactly. Rather, the collected data typically has the form

$$g = Kf + \eta,$$

where η is a vector representing (unknown) noise or measurement errors. The goal, then, is given an ill-conditioned matrix K and a vector g , compute an approximation of the unknown vector f .

Naïvely solving $Kf = g$ usually does not work because the matrix K is so ill-conditioned. Instead, *regularization* is used to make the problem less sensitive to the noise.

Tikhonov Regularization

The best-known regularization procedure—*Tikhonov regularization*—computes a solution of the damped least-squares problem:

$$\min_f \left\{ \|g - Kf\|_2^2 + \alpha^2 \|f\|_2^2 \right\}. \quad (1)$$

The extra term $\alpha^2 \|f\|_2^2$ imposes a penalty for making the norm of the solution too big, which reduces the effect of small singular values. This regularized problem is also a least-squares problem.

Your Homework Assignment

How does a computational scientist learn the craft? Most of us either took traditional courses in numerical analysis or picked up the information “on the street.” We probably developed our expertise over a small set of application areas. As a result, we may lack knowledge of the computational tools we don’t use every day in our work; worse, we may have an overly narrow view of the role of computation across the science and engineering disciplines.

Despite 50 years of scientific computing research and practice, graduate education in computational science and engineering is still in its infancy. Although the number of programs is increasing rapidly, there is a dearth of educational resources.

Last semester, I taught the first course of a two-course graduate sequence in computational science at the University of Maryland, so I experienced this difficulty first hand. The primary obstacle is that no single text covers our course material (numerical linear algebra, optimization, solution of ordinary differential equations and differential-algebraic equations, and Monte Carlo methods) at a suitable level. (The situation for the second course, which focuses on sparse ma-

trix computation and partial differential equations, is a little better.) The second obstacle is that students really need to see a lot of case studies and problem sets to observe methods in action, and developing these resources is a lot of work!

So the purpose of this new column, Your Homework Assignment, is to provide resources for computational science education in the form of projects. In each issue, we will publish one project and the written part of the solution to the previous issue’s project, with code and data sets available on the column’s Web page (<http://computer.org/cise/homework/v5n3.htm>).

These computational science projects should be useful to students (for understanding and motivating the algorithms they study in class), to educators (who can use the projects as a basis for designing their own homework assignments), and I hope to practitioners (who might find them interesting to read as a means of continuing education or uncovering useful new techniques).

In conjunction with CISE’s Education and Computing Prescriptions departments, this new feature should confirm CISE as the premier source of information for CS&E teachers and students, beginners as well as seasoned veterans.

Problem 1. Show that Equation 1 is equivalent to the linear least-squares problem,

$$\min_{\mathbf{f}} \left\| \begin{bmatrix} \mathbf{g} \\ 0 \end{bmatrix} - \begin{bmatrix} \mathbf{K} \\ \alpha \mathbf{I} \end{bmatrix} \mathbf{f} \right\|_2^2. \quad (2)$$

The scalar α (called a *regularization parameter*) controls the solution’s degree of smoothness. Note that $\alpha = 0$ implies no regularization; the computed solution to Equation 2 with $\alpha = 0$ will likely be horribly corrupted with noise. On the other hand, if α is large, then the computed solution cannot be a good approximation of the exact $\mathbf{f}$. Choosing an appropriate value for α is not a trivial matter. Various algorithms appear elsewhere in the literature,¹ but we use a manual approach here.

Let’s turn to the problem of solving the least-squares problem encountered in Equation 2 in Problem 1.

Problem 2. Show that if $\mathbf{K}$ has a singular value decomposition $\mathbf{K} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, then Equation 2 can be transformed into the equivalent least-squares problem,

$$\min_{\mathbf{f}} \left\| \begin{bmatrix} \hat{\mathbf{g}} \\ 0 \end{bmatrix} - \begin{bmatrix} \mathbf{\Sigma} \\ \alpha \mathbf{I} \end{bmatrix} \hat{\mathbf{f}} \right\|_2^2, \quad (3)$$

where $\hat{\mathbf{f}} = \mathbf{V}^T \mathbf{f}$ and $\hat{\mathbf{g}} = \mathbf{U}^T \mathbf{g}$.

Problem 3. Determine a formula for the solution to Equation 3. Hint: you should set the derivative of the minimization function to zero and solve for $\hat{\mathbf{f}}$.

This gives us an algorithm to determine the Tikhonov solution to a discrete ill-posed problem.

Truncated Singular Value Decomposition

Another way of regularizing the problem is to truncate the singular value decomposition (SVD). Problem 4 demonstrates how to express the solution to the least-squares problem in terms of the SVD.

We can see that trouble occurs in $\mathbf{f}_{ls}$ if a small value of σ_i divides a term $\mathbf{u}_i^T \mathbf{g}$ that is dominated by error. In such cases, $\mathbf{f}_{ls}$ will be dominated by error.

Problem 4. Show that the solution to the problem

$$\min_{\mathbf{f}} \left\| \mathbf{g} - \mathbf{K}\mathbf{f} \right\|_2^2$$

is

$$\mathbf{f}_{ls} = \mathbf{V}\mathbf{\Sigma}^\dagger \mathbf{U}^T \mathbf{g} \equiv \sum_{i=1}^n \frac{\mathbf{u}_i^T \mathbf{g}}{\sigma_i} \mathbf{v}_i,$$

where $\mathbf{u}_i$ is the i th column of $\mathbf{U}$, and $\mathbf{v}_i$ is the i th column of $\mathbf{V}$.




Figure 1. Use two algorithms to read the text in this blurred message.

To overcome this, Richard Hanson (as well as James Varah) suggested truncating the previously mentioned expansion,^{3,4}

$$\mathbf{f}_t = \sum_{i=1}^p \frac{\mathbf{u}_i^T \mathbf{g}}{\sigma_i} \mathbf{v}_i,$$

for some value of $p < n$.

Now we have all the tools in place to solve a deblurring problem in image processing. Suppose we have a blurred, noisy image $\mathbf{G}$ (along with some knowledge of the blurring operator), and we want to reconstruct the true original image $\mathbf{F}$. This is an example of a discrete ill-posed problem, in which the vectors in the linear system $\mathbf{g} = \mathbf{K}\mathbf{f} + \boldsymbol{\eta}$ represent the image arrays stacked by columns to form vectors. In Matlab notation, it looks like this:

- $\mathbf{f} = \text{reshape}(\mathbf{F}, \mathbf{n}, 1)$,
- $\mathbf{g} = \text{reshape}(\mathbf{G}, \mathbf{n}, 1)$.

The goal in this problem is given $\mathbf{K}$ and $\mathbf{G}$, reconstruct an approximation of the unknown image $\mathbf{F}$.

If we assume $\mathbf{F}$ and $\mathbf{G}$ contain $\sqrt{n} \times \sqrt{n}$ pixels, then $\mathbf{f}$ and $\mathbf{g}$ are vectors of length n , and $\mathbf{K}$ is an $n \times n$ matrix representing the blurring operation. In general, this matrix is too large to use the SVD. However, in some cases, we can write $\mathbf{K}$ as a Kronecker product, $\mathbf{K} = \mathbf{A} \otimes \mathbf{B}$, and then we can use the SVD.

A Few Facts on Kronecker Products⁵

The Kronecker product $\mathbf{A} \otimes \mathbf{B}$, in which $\mathbf{A}$ is an $m \times m$ matrix, is defined as

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \dots & a_{1m}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \dots & a_{2m}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1}\mathbf{B} & a_{m2}\mathbf{B} & \dots & a_{mm}\mathbf{B} \end{bmatrix}.$$

Theorem 1. If $\mathbf{A} = \mathbf{U}_A \boldsymbol{\Sigma}_A \mathbf{V}_A^T$ and $\mathbf{B} = \mathbf{U}_B \boldsymbol{\Sigma}_B \mathbf{V}_B^T$, then $\mathbf{K} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T$, where $\mathbf{U} = \mathbf{U}_A \otimes \mathbf{U}_B$, $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_A \otimes \boldsymbol{\Sigma}_B$, and $\mathbf{V} = \mathbf{V}_A \otimes \mathbf{V}_B$.

Therefore, computing the SVD of a large matrix is possible if it is the Kronecker product of two smaller ones. On the Web page for this column (<http://computer.org/cise/homework/v5n3.htm>), there is a sample Matlab program, `projdemo.m`, illustrating this property.

To solve our image-deblurring problem, we must operate carefully with the small matrices; otherwise, storage quickly becomes an issue. Again, see the sample program for guidance. With the Kronecker product as a tool, we are ready to compute.

Problem 5. Write a program that takes matrices $\mathbf{A}$ and $\mathbf{B}$ and image $\mathbf{G}$ and computes approximations to image $\mathbf{F}$ using Tikhonov regularization and truncated SVD. For each of these two algorithms, experiment to find the value of the regularization parameter (α for Tikhonov or p for truncated SVD) that gives the clearest image.

The Web page contains sample data: a blurred image $\mathbf{G}$ (see Figure 1) and the matrices $\mathbf{A}$ and $\mathbf{B}$. Your job is to restore the image well enough so that you can read the text in it. Compare the effectiveness of the two algorithms. 

References

1. P.C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*, SIAM Press, 1998.
2. G.H. Golub and C.F. Van Loan, *Matrix Computations*, 2nd ed., Johns Hopkins Press, 1989.
3. R.J. Hanson, "A Numerical Method for Solving Fredholm Integral Equations of the First Kind Using Singular Values," *SIAM J. Numerical Analysis*, vol. 8, 1971, pp. 616–622.
4. J.M. Varah, "On the Numerical Solution of Ill-Conditioned Linear Systems with Applications to Ill-Posed Problems," *SIAM J. Numerical Analysis*, vol. 10, 1973, pp. 257–267.
5. R.A. Horn and C.R. Johnson, *Topics in Matrix Analysis*, Cambridge Univ. Press, 1991.

James G. Nagy is an associate professor in the Mathematics and Computer Science Department and the current Distinguished Teaching Professor in the Social and Natural Sciences at Emory University. His technical interests include numerical linear algebra, ill-posed problem, and image processing. Contact him at nagy@mathcs.emory.edu.

Dianne P. O'Leary is a professor of computer science and a faculty member in the Institute for Advanced Computer Studies and the Applied Mathematics Program at the University of Maryland. Her technical interests include numerical linear algebra, optimization, and scientific computing. Contact her at oleary@cs.umd.edu.