



## UPDATING AND DOWNDATING MATRIX FACTORIZATIONS: A CHANGE IN PLANS

By Dianne P. O'Leary

**W**E SELDOM GET IT RIGHT THE FIRST TIME. WHETHER WE'RE COMPOSING AN EMAIL, SEASONING A STEW, PAINTING A PIC-

TURE, OR PLANNING AN EXPERIMENT, WE

almost always make improvements on our original thought. The same is true of engineering design; we draft a plan, but changes are almost always made. Perhaps the customer changes the performance specifications, or perhaps a substitution of building materials leads to a redesign.

In this homework assignment, we consider numerical methods that make it easier to reanalyze the design after small changes are made. For definiteness, we focus on truss design, but the same principles apply to any linear model.

### Truss Design

Consider the first truss of Figure 1. Beginning engineering students learn to compute the force acting on each truss member by considering two equations for each node in the truss, ensuring that the sum of the horizontal forces is zero, the sum of the vertical forces is zero, and the moment is zero. They can “march” through the equations, solving for the horizontal and vertical forces by a clever ordering. If the design is changed—by moving a node, for example, as in the second truss in the figure—then the resulting forces are just as easy to determine.

For more complex models—for example, a finite element model of the forces on a wing surface—marching no longer works, and the system of equations must be solved using a method such as Gauss elimination. We would like to have an algorithm that enabled us to easily recompute the forces if the wing's shape changed slightly.

To introduce methods for solving modified models, let's forget the marching trick and return to the familiar truss example. We start by writing a system of equations  $Af = \ell$  for the unknown force on each member; the matrix of the system has one column for each unknown force and two rows

per node (for horizontal and vertical forces). The load on node  $E$  is put in the right-hand side. For Figure 1's very simple first truss, for example, we have  $n = 10$  forces to compute, the solution to the linear system

$$\begin{bmatrix} -1 & 0 & c & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & s & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -c & c & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & s & s & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -c & c & 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & s & s & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -c & c & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & s & s & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -c & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & s & 0 & 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} f_{Ab} \\ f_{Av} \\ f_{BC} \\ f_{CD} \\ f_{DE} \\ f_{BD} \\ f_{AC} \\ f_{CE} \\ F_{Ev} \end{bmatrix} = \begin{bmatrix} \ell_{Ab} \\ \ell_{Av} \\ \ell_{Bb} \\ \ell_{Bv} \\ \ell_{Cb} \\ \ell_{Cv} \\ \ell_{Db} \\ \ell_{Dv} \\ \ell_{Eb} \\ \ell_{Ev} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ -50 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (1)$$

(The subscript  $Ab$ , for example, denotes a horizontal force at node  $A$ , and the subscript  $BC$  refers to the member connecting nodes  $B$  and  $C$ .) The truss members are all equal length, so  $c = \cos(\pi/3)$  and  $s = \sin(\pi/3)$ .

Several simple kinds of changes to the truss design produce “small” changes in the system of equations. For example,

- If we change the loading on the truss, then we keep the

## HANDLING SMALL CHANGES

The problem of efficiently handling small changes in the model matrix arises in many situations other than engineering design. For example,

- Suppose we're solving a system of linear inequalities  $Ax \geq b$ , with  $A$  of dimension  $m \times n$  ( $m \geq n$ ), and we think that the first  $n$  of them should be active:  $a_i^T x = b_i$ ,  $i = 1, \dots, n$ . Suppose then that the solution to these equations violates the  $k$ th inequality ( $k > n$ ), and we want to add it to our current system of equations and delete the  $j$ th equation. Can we solve our new linear system easily? This problem routinely arises in minimization problems when we have

linear inequality constraints, and it's the basic computation in the simplex algorithm for solving linear programming problems.

- Suppose that we're solving a linear least-squares problem,  $Ax \approx b$ , and we get some new measurements. This adds rows to  $A$  and  $b$ . Can we solve our new least-squares problem easily?
- Suppose we've computed the eigenvalues and eigenvectors of  $A$ , and then  $A$  is changed by the addition of a rank-1 matrix  $\hat{A} = A + cr^T$ . What are the eigenvalues of  $\hat{A}$ ?

Each of these problems (and similar ones) can be solved cheaply by using the techniques discussed in this homework assignment.

same matrix but change the right-hand side  $\ell$ .

- If we add a new node along with two new truss members, then our new matrix has two additional rows and columns and contains the old matrix as a submatrix.
- If we remove a set of nodes and their truss members, then we delete the columns of the matrix corresponding to the forces exerted by the members, and we delete the pair of rows corresponding to each removed node.
- If we move a node, then we change the two rows corresponding to the horizontal and vertical forces on that node, and we change the columns corresponding to its members.

For a problem with  $n = 10$  unknowns, we could easily recompute the answer after any of these changes, but if  $n = 1,000,000$ , we might want to take advantage of our solution to the original problem to more quickly obtain a solution to the modified problem. In this homework assignment, we develop the tools to do this.

### Changes to the Right-Hand Side

If we need to analyze the truss for several different loadings, then it's a good idea to compute a factorization of the matrix  $A$  once and save it for multiple uses. For example, if we compute the  $LU$  decomposition with partial pivoting


$$PA = LU, \quad (2)$$

where  $P$  is a permutation matrix that interchanges rows of  $A$ ,  $L$  is a lower triangular matrix, and  $U$  is an upper triangular matrix, then each of the loads can be handled by solving

$$Ly = P\ell$$

by forward-substitution, and then computing

$$Uf = y$$



**Figure 1. Truss design.** (a) A truss with five nodes and seven equal-length members, loaded with a force of 50 Newtons. Node A is fixed (supported horizontally and vertically) and node E is rolling (supported only vertically). (b) A change in the truss. Member CE is now two times the length of the other members.

by back-substitution. If  $A$  were a dense matrix, with very few nonzeros, then the initial  $LU$  decomposition would cost  $O(n^3)$  operations, whereas forward- and back-substitution would cost only  $O(n^2)$ . Taking advantage of the sparsity of  $A$  can reduce the cost of both the  $LU$  decomposition and the substitution (see "Solving Sparse Linear Systems: Taking the Direct Approach," *Computing in Science & Eng.*, vol. 7, no. 5, 2005, pp. 62–67), but substitution will still be significantly less costly than factorization when  $n$  is large.

## Changes to the Matrix

If we change the shape of the truss, then the matrix changes. We can handle this in two ways: by using the *Sherman-Morrison-Woodbury* formula or by updating a matrix decomposition.

### Sherman-Morrison-Woodbury Formula

Sometimes our matrix changes in a rather simple way, and we want to reconsider our problem, making use of the original decomposition without explicitly forming the update.

Suppose we have the factorization from Equation 2 and now want to solve the linear system  $(A - \mathbf{z}\mathbf{v}^T)\mathbf{f} = \ell$ , where  $\mathbf{z}$  and  $\mathbf{v}$  are column vectors of length  $n$ , so that  $\mathbf{z}\mathbf{v}^T$  is an  $n \times n$  matrix. For example, if we decide to move node  $E$  as in the truss of Figure 1, then we need to change column 6 in our matrix. To do this, we set  $\mathbf{v}$  to be the 6th unit vector (1 in position 6 and zeroes elsewhere) and  $\mathbf{z}$  to be the difference between the new column and the old one.

In the next problem, we see how to apply this principle to more than one set of changes in the matrix.

#### PROBLEM 1.

- Suppose we want to change columns 6 and 7 in our matrix  $A$ . Express the new matrix as  $A - \mathbf{Z}\mathbf{V}^T$ , where  $\mathbf{Z}$  and  $\mathbf{V}$  have dimension  $n \times 2$ .
- Suppose we want to change both column 6 and row 4 of  $A$ . Find  $\mathbf{Z}$  and  $\mathbf{V}$  so that our new matrix is  $A - \mathbf{Z}\mathbf{V}^T$ .

In Problem 2, we see how this formulation of our new matrix as a small-rank change in our old matrix leads to an efficient computational algorithm.

#### PROBLEM 2.

- Assume that  $A$  and  $A - \mathbf{Z}\mathbf{V}^T$  are both nonsingular. Show that

$$(A - \mathbf{Z}\mathbf{V}^T)^{-1} = A^{-1} + A^{-1}\mathbf{Z}(\mathbf{I} - \mathbf{V}^T A^{-1}\mathbf{Z})^{-1}\mathbf{V}^T A^{-1}$$

by verifying that the product of this matrix with  $A - \mathbf{Z}\mathbf{V}^T$  is the identity matrix  $\mathbf{I}$ . This is called the Sherman-Morrison-Woodbury formula.

- Suppose we have an  $LU$  decomposition of  $A$  as in Equation 2. Assume that  $\mathbf{Z}$  and  $\mathbf{V}$  are  $n \times k$  and  $k \ll n$ . Show that we can use this decomposition and the Sherman-Morrison-Woodbury formula to solve the linear system  $(A - \mathbf{Z}\mathbf{V}^T)\mathbf{f} = \ell$

without forming any matrix inverses. (If  $A$  is dense, then we perform  $O(n^2)$  operations using Sherman-Morrison-Woodbury, rather than the  $O(n^3)$  operations needed to solve the linear system from scratch.) Hint: Remember that we can compute  $A^{-1}\mathbf{y}$  by solving a linear system using forward- and back-substitution with the factors  $L$  and  $U$ .

In the following two problems, let's experiment with the Sherman-Morrison-Woodbury algorithm to see when it can be useful.

#### PROBLEM 3.

Implement the Sherman-Morrison-Woodbury algorithm from Problem 2b. Debug it by factoring the matrix in Equation 1, modeling the first truss in Figure 1, and then changing the model to the second truss.

#### PROBLEM 4.

For  $n$  taken to be various numbers between 10 and 1,000, generate a random  $n \times n$  matrix  $A$ . Find the number of updates  $k_0$  that makes the cost of the Sherman-Morrison-Woodbury method comparable to computing  $A \setminus \ell$ . Plot  $k_0$  as a function of  $n$ .

### Updating a Matrix Decomposition

The Sherman-Morrison-Woodbury formula lets us solve modified linear systems without explicitly modifying our matrix factorization. It's very efficient when we need to make only a few changes.

In some problems, though, we must do a long series of updates to the matrix, and it's better to explicitly update the factorization. We could consider updating an  $LU$  factorization, but this can be complicated because pivoting is necessary to preserve stability. Instead, let's use a factorization that's stable without pivoting. This also enables us to consider matrices that have more rows than columns, such as those that arise in least-squares problems.

Suppose we've factored

$$A = QR,$$

where  $A$  is  $m \times n$  with  $m \geq n$ ,  $Q$  is  $m \times m$  and orthogonal

( $Q^T Q = I$ ), and  $R$  is  $m \times n$  and has zeros below its main diagonal. For definiteness, we'll let  $A$  have dimensions  $5 \times 3$ .

As examples, let's consider two kinds of common changes:

- **Adding a row.** In least squares, this happens when new data comes in; in our truss problem it means that we have a new node.
- **Deleting a column.** In least squares, this happens when we decide to reduce the number of parameters in the model; in our truss problem, it could result from removing a member.

Let's look at both in more detail.

### Adding a Row

Denote the new matrix as  $\hat{A}$ . Our decomposition can be written as

$$\hat{A} = \begin{bmatrix} Q & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ 0 & r_{22} & r_{23} \\ 0 & 0 & r_{33} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} a_{61} & a_{62} & a_{63} \end{bmatrix} \quad (3)$$

To complete the decomposition, we need to reduce the  $a$ 's to zeros. We can do this by using  $n$  Givens rotations—much cheaper than recomputing the entire decomposition.

We'll write the Givens matrix as

$$G = \begin{bmatrix} c & s \\ -s & c \end{bmatrix},$$

where  $c^2 + s^2 = 1$ . Thus,  $c$  and  $s$  have the geometric interpretation of the cosine and sine of angle  $\theta$ ; multiplying a vector by this matrix rotates the vector by an angle  $\theta$ .

Let's see how we can use Givens rotation matrices.

### PROBLEM 5.

- Given a vector  $z \neq 0$  of dimension  $2 \times 1$ , find  $G$  so that  $Gz = x\mathbf{e}_1$ , where  $x = \|z\|$  and  $\mathbf{e}_1$  is the vector with a 1 in the first position and zeros elsewhere.
- We will use the notation  $G_{ij}$  to denote an  $m \times m$  identity matrix with its  $i$ th and  $j$ th rows modified to include the

Givens rotation: for example, if  $m = 6$ , then

$$G_{25} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & c & 0 & 0 & s & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & s & 0 & 0 & -c & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Multiplication of a vector by this matrix leaves all but rows 2 and 5 of the vector unchanged. Show that we can finish our  $QR$  decomposition by (left) multiplying the  $R$  matrix in Equation 3 first by  $G_{16}$ , then by  $G_{26}$ , and finally by  $G_{36}$ , where the angle defining each of these matrices is suitably chosen. To preserve the equality, we multiply the  $Q$ -matrix by  $G_{16}^T G_{26}^T G_{36}^T$  on the right, and we have the updated factorization.

### Deleting a Column

If we delete column 1 from  $A$ , for example, we can write the decomposition as

$$\hat{A} = Q \begin{bmatrix} r_{12} & r_{13} \\ r_{22} & r_{23} \\ 0 & r_{33} \\ 0 & 0 \\ 0 & 0 \end{bmatrix}.$$

The resulting  $R$  is almost upper triangular; we just need rotations to reduce the elements labeled  $r_{22}$  and  $r_{33}$  to zero. In general, we need  $n - k$  rotations when column  $k$  is deleted.

Algorithms for deleting a row and adding a column are similar to those for adding a row and deleting a column. In the next problem, we construct algorithms for changing a column.

### PROBLEM 6.

Write a Matlab function that updates a  $QR$  decomposition of a matrix  $A$  when a single column is changed. Apply it to the truss examples in Problem 3.

### The Point of Updating

It might seem silly to worry so much about whether to update

## Tools

Standard textbooks on scientific computing, such as the one by Charles van Loan, discuss the **QR** factorization and the use of Givens rotations.<sup>1</sup>

Other stable methods for modifying matrix factorizations are considered by P.E. Gill, Gene Golub, Walter Murray, and Michael Saunders.<sup>2</sup> Gene Golub also discusses the solution of eigenvalue problems when the matrix is modified.<sup>3</sup>

### References

1. C.F. van Loan, *Introduction to Scientific Computing*, Prentice Hall, 2000, section 7.2.
2. P.E. Gill et al., "Methods for Modifying Matrix Factorizations," *Mathematics of Computation*, vol. 28, no. 126, 1974, pp. 505–535.
3. G.H. Golub, "Some Modified Matrix Eigenvalue Problems," *SIAM Rev.*, vol. 15, 1973, pp. 318–335.

or recompute; computers are fast, and if we make one change to the matrix, it really doesn't matter which we do.

But when we need to do the task over and over again, perhaps in a loop that solves a more complicated problem, it's essential to use appropriate updating techniques to reduce the cost.

Unfortunately, the literature contains many unstable updating algorithms, so it's important to use stable and trusted algorithms such as those given here and in the references in the "Tools" sidebar.

### Acknowledgments

I'm grateful to Brendan O'Leary for helpful discussions regarding trusses.

## Partial Solution to Last Issue's Homework Assignment

# COMPUTATIONAL SOFTWARE: WRITING YOUR LEGACY

By Dianne P. O'Leary

IN THE JANUARY/FEBRUARY ISSUE, WE DISCUSSED SITUATIONS IN WHICH WE'RE ASKED TO WORK WITH *LEGACY CODE*—CODE THAT HAS BEEN IN USE FOR A WHILE AND NOW needs maintenance by someone other than its author. We used a sample Matlab code for all our examples.

### PROBLEM 1.

Add documentation to `mgs`.

#### Answer:

See the program on the Web site ([www.computer.org/cise/homework](http://www.computer.org/cise/homework)).

### PROBLEM 2.

Judge `mgs` according to each of the first six design principles.

#### Answer:

- Data that a function needs should be specified in variables, not constants. This is fine; `C` is a variable.
- Code should be modular, so that a user can pull out one piece and substitute another when necessary. The program `mgs` factors a matrix into the product of two other matrices, and it would be easy to substitute a different factorization algorithm.
- On the other hand, considerable overhead is involved in function calls, so each module should involve a substantial computation to mask this overhead. This is also satisfied; `mgs` performs a significant computation ( $O(mn^2)$  operations).
- Input parameters should be tested for validity, and clear error messages should be generated for invalid input. The factorization can be performed for any matrix or scalar, so input should be tested to be sure it isn't a string, cell variable, and so on.
- "Spaghetti code" should be avoided. In other words, the sequence of instructions should be top-to-bottom (including loops), without a lot of jumps in control. The `mgs` program is fine in this regard, although there is a lot of nesting of loops.
- The names of variables should be chosen to remind the



reader of their purpose. The letter  $q$  is often used for an orthogonal matrix, and  $r$  is often used for an upper triangular one, but it would probably be better practice to use uppercase names for these matrices.

### PROBLEM 3.

- What does `mgs` do?
- Develop a testing code for `mgs`.
- Users have complained that `mgs` doesn't seem to behave well when the matrix  $C$  has more columns than rows. In particular, they've come to expect that  $q' * q$  is an identity matrix, which is no longer true. Investigate this bug complaint and see what can be done.

#### Answer:


- This program computes a  $QR$  factorization of the matrix  $C$  using the modified Gram-Schmidt algorithm.
- See the Web site ([www.computer.org/cise/homework/](http://www.computer.org/cise/homework/)).
- This is corrected in `mgsfact.m` on the Web site. The columns of  $q$  should be mutually orthogonal, but the number of columns in  $q$  should be the minimum of the row and column dimensions of  $C$ . Nonzero columns after that are just the result of rounding errors.

### PROBLEM 4.

The users have run a profiler on their software system to determine timing for each part of their code, and they've found that `mgs` takes 22 percent of the total time. Their typical input matrices  $C$  have roughly 200 rows and 100 columns. Change `mgs` to make it run faster. Test the original and modified versions, graphing the time required for problems with 200 rows and 50, 60, ..., 200 columns.

#### Answer:

The resulting program is on the Web site ([www.computer.org/cise/homework/](http://www.computer.org/cise/homework/)), and Figure A shows the timing results. The program `mgs` has been modified in `mgsfact` to use vector operations and internal functions like `norm` when possible, and preallocate storage for  $q$  and  $r$ . The `mgsfact` function runs 150 to 200 times faster than `mgs` on matrices with 200 rows, using a Sun UltraSPARC-III with clock speed 750 MHz running Matlab 6. It's an interesting exercise to determine the relative importance of the three changes.



**Figure A.** Time taken by the two algorithms for matrices with 200 rows. The modified algorithm runs 150 to 200 times faster than the original on a Sun UltraSPARC-III.

You also might think about how an efficient implementation in your favorite programming language might differ from this one.

**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. She has a BS in mathematics from Purdue University and a PhD in computer science from Stanford. O'Leary is also a member of SIAM, the ACM, and the Association for Women in Mathematics. Contact her at [oleary@cs.umd.edu](mailto:oleary@cs.umd.edu); [www.cs.umd.edu/users/oleary/](http://www.cs.umd.edu/users/oleary/).

## DON'T RUN THE RISK.

## BE SECURE.

IEEE  
**SECURITY & PRIVACY**

Ensure that your networks operate safely and provide critical services even in the face of attacks. Develop lasting security solutions, with this peer-reviewed publication.

Top security professionals in the field share information you can rely on:

- Wireless Security • Security Infrastructure Security • Privacy Issues • Legal Issues • Cybercrime • Digital Rights Management • Intellectual Property Protection and Piracy • The Security Profession • Education

Order your subscription today.  
[www.computer.org/security/](http://www.computer.org/security/)