

# THE FACTORIZE OBJECT for solving linear systems

Copyright 2011-2012, Timothy A. Davis,

<http://www.suitesparse.com>

[DrTimothyAldenDavis@gmail.com](mailto:DrTimothyAldenDavis@gmail.com)

This is a demonstration of the FACTORIZE object for solving linear systems and least-squares problems, and for computations with the matrix inverse and pseudo-inverse.

## Contents

- Rule Number One: never multiply by the inverse,  $\text{inv}(A)$
- Rule Number Two: never break Rule Number One
- How to use BACKSLASH solve  $A*x=b$
- BACKSLASH versus INV ... let the battle begin
- LU and LINSOLVE are fast and accurate but complicated to use
- INV is easy to use, but slow and inaccurate
- So the winner is ... nobody
- The FACTORIZE object to the rescue
- Least-squares problems
- Underdetermined systems
- Computing selected entries in the inverse or pseudo-inverse
- Computing the entire inverse or pseudo-inverse
- Update/downdate of a dense Cholesky factorization
- Caveat Executor
- Summary

## Rule Number One: never multiply by the inverse, $\text{inv}(A)$

Use backslash or a matrix factorization instead (LU, CHOL, or QR).

## Rule Number Two: never break Rule Number One

However, the problem with Rule Number One is that it can be hard to figure out which matrix factorization to use and how to use it. Using LU, CHOL, or QR is complicated, particularly if you want the best performance. BACKSLASH (MLDIVIDE) is great, but it can't be reused when solving multiple systems ( $x=A\backslash b$  and  $y=A\backslash c$ ). Its syntax doesn't match the use of the inverse in mathematical expressions, either.

The goal of the FACTORIZE object is to solve this problem ...

”Don’t let that INV go past your eyes; to solve that system, FACTORIZE!”

## How to use BACKSLASH solve $A*x=b$

First, let’s create a square matrix A and a right-hand-side b for a linear system  $A*x=b$ . There are many ways to solve this system. The best way is to use  $x=A\backslash b$ . The residual r is a vector of what’s left over in each equation, and its norm tells you how accurately the system was solved.

```
format compact ;
A = rand (3)
b = rand (3,1)
x = A\b
r = b-A*x ;
norm (r)

A =
    0.2748    0.2719    0.3696
    0.2217    0.8329    0.5006
    0.9315    0.5791    0.0753
b =
    0.1293
    0.9138
    0.4589
x =
   -0.3827
    1.4657
   -0.4437
ans =
    1.2413e-16
```

## BACKSLASH versus INV ... let the battle begin

The backslash operation  $x=A\backslash b$  is mathematically the same as  $x=\text{inv}(A)*b$ . However, backslash is faster and more accurate since it uses a matrix factorization instead of multiplying by the inverse. Even though your linear algebra textbook might write  $x=A^{(-1)}*b$  as the solution to the system  $A*x=b$ , your textbook author never means for you to compute the inverse.

These next statements give the same answer, so what’s the big deal?

```
S = inv(A) ;
x = S*b
x = A\b
```

```

x =
    -0.3827
     1.4657
    -0.4437
x =
    -0.3827
     1.4657
    -0.4437

```

The big deal is that you should care about speed and you should care even more about accuracy. BACKSLASH relies on matrix factorization (LU, CHOL, QR, or other specialized methods). It's faster and more reliable than multiplying by the inverse, particularly for large matrices and sparse matrices. Here's an illustration of how pathetic  $\text{inv}(A)*b$  can be.

```
A = gallery ('frank',16) ; xtrue = ones (16,1) ; b = A*xtrue ;
```

```

x = inv(A)*b ; norm (b-A*x)
x = A\b      ; norm (b-A*x)

```

```

ans =
    0.0606
ans =
    1.7764e-15

```

The performance difference between BACKSLASH and INV for even small sparse matrices is striking.

```

load west0479 ;
A = west0479 ;
n = size (A,1)
b = rand (n,1) ;
tic ; x = A\b ; toc
norm (b-A*x)
tic ; x = inv(A)*b ; toc
norm (b-A*x)

```

```

n =
    479
Elapsed time is 0.002425 seconds.
ans =
    1.9174e-10
Elapsed time is 0.041780 seconds.
ans =
    3.4438e-09

```

What if you want to solve multiple systems? Use a matrix factorization. But which one? And how do you use it? Here are some alternatives using LU for the sparse west0479 matrix, but some are faster than others.

```
tic ; [L,U]      = lu(A) ; x1 = U \ (L \ b)          ; t1=toc ; nz1=nnz(L+U);
tic ; [L,U,P]    = lu(A) ; x2 = U \ (L \ P*b)        ; t2=toc ; nz2=nnz(L+U);
tic ; [L,U,P,Q]  = lu(A) ; x3 = Q * (U \ (L \ P*b)) ; t3=toc ; nz3=nnz(L+U);

fprintf ('1: nnz(L+U): %5d time: %8.4f resid: %e\n', nz1,t1, norm(b-A*x1));
fprintf ('2: nnz(L+U): %5d time: %8.4f resid: %e\n', nz2,t2, norm(b-A*x2));
fprintf ('3: nnz(L+U): %5d time: %8.4f resid: %e\n', nz3,t3, norm(b-A*x3));

1: nnz(L+U): 16151 time:    0.0031 resid: 1.346712e-10
2: nnz(L+U): 15826 time:    0.0081 resid: 1.757002e-10
3: nnz(L+U):  3704 time:    0.0017 resid: 1.691750e-10
```

## LU and LINSOLVE are fast and accurate but complicated to use

A quick look at “help lu” will scroll off your screen. For full matrices, [L,U,p] = lu (A,'vector') is fastest. Then for the forward/back solves, use LINSOLVE instead of BACKSLASH for even faster performance. But for sparse matrices, use the optional 'Q' output of LU so you get a good fill-reducing ordering. But you can't use 'Q' if the matrix is full. But LINSOLVE doesn't work on sparse matrices.

But ... Ack! That's getting complicated ...

Here's the best way to solve  $A*x=b$  and  $A*y=c$  when A is full and unsymmetric:

```
n = 1000 ;
A = rand (n) ;
b = rand (n,1) ;
c = rand (n,1) ;
tic ; [L,U,p] = lu (A, 'vector') ; LUtime = toc

tic ; x = U \ (L \ b (p,:)) ;
      y = U \ (L \ c (p,:)) ; toc

tic ; opL = struct ('LT', true) ;
      opU = struct ('UT', true) ;
      x = linsolve (U, linsolve (L, b(p,:), opL), opU) ;
      y = linsolve (U, linsolve (L, c(p,:), opL), opU) ; toc
```

```

LUtime =
    0.0155
Elapsed time is 0.004917 seconds.
Elapsed time is 0.002092 seconds.

```

## INV is easy to use, but slow and inaccurate

Oh bother! Using LU and LINSOLVE is too complicated. You just want to solve your system. Let's just compute  $\text{inv}(A)$  and use it twice. Easy to write, but slower and less accurate ...

```

S = inv (A) ;
x = S*b ; norm (b-A*x)
y = S*c ; norm (c-A*y)

```

```

ans =
    2.3292e-11
ans =
    1.7292e-11

```

Sometimes using the inverse seems inevitable. For example, your textbook might show the Schur complement formula as  $S = A - B \cdot \text{inv}(D) \cdot C$ . This can be done without  $\text{inv}(D)$  in one of two ways: SLASH or BACKSLASH (MRDIVIDE or MLDIVIDE to be precise).

$\text{inv}(A) \cdot B$  and  $A \backslash B$  are mathematically equivalent, as are  $B \cdot \text{inv}(A)$  and  $B/A$ , so these three methods give the same results (ignoring computational errors, which are worse for  $\text{inv}(D)$ ). Only the first equation looks like the equation in your textbook, however.

```

A = rand (200) ; B = rand (200) ; C = rand (200) ; D = rand (200) ;

tic ; S1 = A - B*inv(D)*C ; toc ;
tic ; S2 = A - B*(D\C) ;      toc ;
tic ; S3 = A - (B/D)*C ;      toc ;

```

```

Elapsed time is 0.002398 seconds.
Elapsed time is 0.001632 seconds.
Elapsed time is 0.001474 seconds.

```

## So the winner is ... nobody

BACKSLASH: mostly simple to use (except remember that Schur complement formula?). Fast and accurate ... but slow if you want to solve two linear systems with the same matrix A.

LU, QR, CHOL: fast and accurate. Awful syntax to use. Drag out your linear algebra textbook if you want to use these in MATLAB. Whenever I use them I have to derive them from scratch, even though I **wrote** most of the sparse factorizations used in MATLAB!

INV: slow and inaccurate. Wins big on ease-of-use, though, since it's a direct plug-in for all your nice mathematical formulas.

No method is best on all three criterion: speed, accuracy, and ease of use.

Is there a solution? Yes ... keeping reading ...

## The FACTORIZE object to the rescue

The FACTORIZE method is just as easy to use as INV, but just as fast and accurate as BACKSLASH, LU, QR, CHOL, and LINSOLVE.

`F = factorize(A)` computes the factorization of A and returns it as an object that you can reuse to solve a linear system with `x=F\b`. It picks LU, QR, or Cholesky for you, just like BACKSLASH.

`S = inverse(A)` is simpler yet. It does NOT compute `inv(A)`, but factorizes A. When multiplying `S*b`, it doesn't multiply by the inverse, but uses the correct forward/backsolve equations to solve the linear system.

```
n = 1000 ;
A = rand (n) ;
b = rand (n,1) ;
c = rand (n,1) ;

tic ;                               x = A\b ; y = A\c ; toc
tic ; S = inv(A) ;                  x = S*b ; y = S*c ; toc
tic ; F = factorize(A) ;            x = F\b ; y = F\c ; toc
tic ; S = inverse(A) ;              x = S*b ; y = S*c ; toc
```

Elapsed time is 0.051483 seconds.

Elapsed time is 0.053813 seconds.

Elapsed time is 0.025811 seconds.

Elapsed time is 0.029840 seconds.

## Least-squares problems

Here are some different methods for solving a least-squares problem when your system is over-determined. The last two methods are the same.

```
A = rand (1000,200) ;
b = rand (1000,1) ;

tic ; x = A\b          ; toc, norm (A'*A*x-A'*b)
tic ; x = pinv(A)*b    ; toc, norm (A'*A*x-A'*b)
tic ; x = inverse(A)*b ; toc, norm (A'*A*x-A'*b)
tic ; x = factorize(A)\b ; toc, norm (A'*A*x-A'*b)
```

```
Elapsed time is 0.010705 seconds.
ans =
    2.7575e-12
Elapsed time is 0.025339 seconds.
ans =
    2.3477e-12
Elapsed time is 0.008903 seconds.
ans =
    2.5140e-12
Elapsed time is 0.008408 seconds.
ans =
    2.5140e-12
```

FACTORIZE is better than BACKSLASH because you can reuse the factorization for different right-hand-sides. For full-rank matrices, it's better than PINV because it's faster (and PINV fails for sparse matrices).

```
A = rand (1000,200) ;
b = rand (1000,1) ;
c = rand (1000,1) ;

tic ;          ; x = A\b ; y = A\c ; toc
tic ; S = pinv(A) ; x = S*b ; y = S*c ; toc
tic ; S = inverse(A) ; x = S*b ; y = S*c ; toc
tic ; F = factorize(A) ; x = F\b ; y = F\c ; toc
```

```
Elapsed time is 0.021934 seconds.
Elapsed time is 0.027198 seconds.
Elapsed time is 0.009692 seconds.
Elapsed time is 0.010197 seconds.
```

## Underdetermined systems

The under-determined system  $Ax=b$  where  $A$  has more columns than rows has many solutions.  $x=A\backslash b$  finds a basic solution (some of the entries in  $x$  are zero).  $\text{pinv}(A)*b$  finds a minimum 2-norm solution, but it's slow. QR factorization will do the same if  $A$  has full rank. That's what the `factorize(A)` and `inverse(A)` methods do.

```
A = rand (200,1000) ;
b = rand (200,1) ;

tic ; x = A\b          ; toc, norm (x)
tic ; x = pinv(A)*b     ; toc, norm (x)
tic ; x = inverse(A)*b  ; toc, norm (x)
tic ; x = factorize(A)\b ; toc, norm (x)
```

```
Elapsed time is 0.016247 seconds.
ans =
    2.5580
Elapsed time is 0.034924 seconds.
ans =
    0.5154
Elapsed time is 0.012460 seconds.
ans =
    0.5154
Elapsed time is 0.009924 seconds.
ans =
    0.5154
```

## Computing selected entries in the inverse or pseudo-inverse

If you want just a few entries from the inverse, it's still better to formulate the problem as a system of linear equations and use a matrix factorization instead of computing `inv(A)`. The `FACTORIZE` object does this for you, by overloading the subsref operator.

```
A = rand (1000) ;

tic ; S = inv (A)      ; S (2:3,4), toc
tic ; S = inverse (A) ; S (2:3,4), toc

ans =
   -0.2554
    0.0993
```



```
Elapsed time is 0.047095 seconds.
ans =
    -0.2554
     0.0993
Elapsed time is 0.030233 seconds.
```

## Computing the entire inverse or pseudo-inverse

Rarely, and I mean RARELY, you really do need the inverse. More frequently what you want is the pseudo-inverse. You can force a factorization to become a plain matrix by converting it to double. Note that `inverse(A)` only handles full-rank matrices (either dense or sparse), whereas `pinv(A)` works for all dense matrices (not sparse).

The explicit need for `inv(A)` (or `S=A\eye(n)`, which is the same thing) is RARE. If you ever find yourself multiplying by the inverse, then you know one thing for sure. You know with certainty that you don't know what you're doing.

```
A = rand (500) ;
tic ; S1 = inv (A) ;           ; toc
tic ; S2 = double (inverse (A)) ; toc
norm (S1-S2)
```

```
A = rand (500,400) ;
tic ; S1 = pinv (A)           ; toc
tic ; S2 = double (inverse (A)) ; toc
norm (S1-S2)
```

```
Elapsed time is 0.009174 seconds.
Elapsed time is 0.013046 seconds.
ans =
    1.6446e-12
Elapsed time is 0.073158 seconds.
Elapsed time is 0.016360 seconds.
ans =
    1.7565e-14
```

## Update/downdate of a dense Cholesky factorization

Wilkinson considered the update/downdate of a matrix factorization to be a key problem in computational linear algebra. The idea is that you first factorize a matrix. Next, make a low-rank change to  $A$ , and patch up (or down...) the factorization so that it becomes the factorization of the new matrix. In MATLAB,

this only works for dense symmetric positive definite matrices, via cholupdate. This is much faster than computing the new factorization from scratch.

```
n = 1000 ;
A = rand (n) ;
A = A*A' + n*eye (n) ;
w = rand (n,1) ; t = rand (n,1) ; b = rand (n,1) ;
F = factorize (A) ;
```

```
tic ; F = cholupdate (F,w,'+') ; x = F\b ; toc
tic ; y = (A+w*w')\b ;          toc
norm (x-y)
```

```
tic ; F = cholupdate (F,t,'-') ; x = F\b ; toc
tic ; y = (A+w*w'-t*t')\b ; toc
norm (x-y)
```

```
Elapsed time is 0.010284 seconds.
Elapsed time is 0.015608 seconds.
ans =
    3.6395e-17
Elapsed time is 0.011662 seconds.
Elapsed time is 0.019173 seconds.
ans =
    3.5119e-17
```

## Caveat Executor

One caveat: If you have a large number of very small systems to solve, the object-oriented overhead of creating and using an object can dominate the run time, at least in MATLAB R2011a. For this case, if you want the best performance, stick with BACKSLASH, or LU and LINSOLVE (just extract the appropriate formulas from the M-files in the FACTORIZE package).

Hopefully the object-oriented overhead will drop in future versions of MATLAB, and you can ignore this caveat.

```
A = rand (10) ; b = rand (10,1) ; F = factorize (A) ;

tic ; for k = 1:10000, x = F\b ; end ; toc

tic ; for k = 1:10000, x = A\b ; end ; toc

[L,U,p] = lu (A, 'vector') ;
```

```

opL = struct ('LT', true) ;
opU = struct ('UT', true) ;
tic ;
for k = 1:10000
    x = linsolve (U, linsolve (L, b(p,:), opL), opU) ;
end
toc

```

```

Elapsed time is 1.156836 seconds.
Elapsed time is 0.123091 seconds.
Elapsed time is 0.057961 seconds.

```

## Summary

So ... don't use INV, and don't worry about how to use LU, CHOL, or QR factorization. Just install the FACTORIZE package, and you're on your way. Assuming you are now in the Factorize/ directory, cut-and-paste these commands into your command window:

```

addpath (pwd)
savepath

```

And remember ...

*Don't let that INV go past your eyes; to solve that system, FACTORIZE!*