Subject: Re: SVDFIT Problems
Posted by Paul Van Delst[1] on Fri, 30 Aug 2002 17:08:31 GMT
View Forum Message <> Reply to Message

Bill wrote:
>
> In my occaisional attempts to use SVDFIT I have had problems in that it
> reports singular solutions when other fitting routines, e.g., REGRESS,
> have had not problems.

Deja vu! I was using the SVDFIT routine just yesterday to fit a second order polynomial. I kept getting an error report for a singular value (even with /DOUBLE). Used POLY_FIT with no problems
(SVDFIT _was_ overkill for what I wanted to do.)

<snip>

> Unfortunately it appeas that the code currently relies on a simplified
> version of the Numerical Recipes routine SVDCMP, that does not return
> the singular value vector that Numerical Recipes uses to compute wmax. A
> rewrite to give SVDFIT the proper properties is more work than I am
> willing to put in at this time.

Thanks for taking the time to nut the problem out and report it here. It really pains me (just having finished the yearly license renewal tooth-pulling process) that these sorts of things are *still* found in incarnations of such well-worn linear algebra problems like SVD. I spent about half a day trying to figure out what I was doing wrong until I tossed the lot and surfed over to netlib.org for some fortran routines. I thought IDL was supposed to *save* time for this sort of thing?

paulv

--
Paul van Delst
CIMSS @ NOAA/NCEP/EMC
Ph: (301)763-8000 x7274
Fax:(301)763-8545

Subject: Re: SVDFIT Problems
Posted by William Clodius on Fri, 30 Aug 2002 19:52:55 GMT
View Forum Message <> Reply to Message

Paul van Delst wrote:

> Bill wrote:
>>
>> In my occaisional attempts to use SVDFIT I have had problems in that it

>> reports singular solutions when other fitting routines, e.g., REGRESS,
>> have had not problems.
>
> Deja vu! I was using the SVDFIT routine just yesterday to fit a second order polynomial. I kept
> getting an error report for a singular value (even with /DOUBLE). Used POLY_FIT with no problems
> (SVDFIT _was_ overkill for what I wanted to do.)

Unfortunately it is not overkill for me. Usually I use REGRESS in IDL, but I am trying to do an analysis of some calibrations that imposes mutual consistency requirements on disparate measurements.
None of these measurements has a natural offset, and REGRESS's solutions assume that such an offset
exists. I found a way around that for separate analyses, by transforming the equations so that what is
not normally thought of as an offfset is treated as one. However, the resulting "offsets" for the separate analyses have independent meanings which complicates combining the analyses to ensure mutual
consistency.

For the combined analyses SVDFIT is a reasonable approach, if the code works. Now I either have to
kluge things, or implement an interface to non-IDL code.

>
> <snip>
> Thanks for taking the time to nut the problem out and report it here. It really pains me (just
> having finished the yearly license renewal tooth-pulling process) that these sorts of things are
> *still* found in incarnations of such well-worn linear algebra problems like SVD. I spent about half
> a day trying to figure out what I was doing wrong until I tossed the lot and surfed over to
> netlib.org for some fortran routines. I thought IDL was supposed to *save* time for this sort of
> thing?
>

<snip>

There are many things that bother me:

1. RSI uses Numerical Recipes as the basis for its core numerical library, when there are more robust
libraries available.

2. RSI uses Numerical Recipes in C as the basis of its IDL implementations when the syntax of IDL is
more similar to Fortran, particularly Fortran 90/95, than to C, increasing the probability that the resulting libraries will be incorrect.

3. RSI probably used a non-numericist for SVDFIT. The error I found appears to be deliberate, as that
section of Numerical Recipes is correct, and the differences are not those that would result from the
dirrect translation of C to IDL. The changes appear to indicate a complete lack of understanding of the
algorithm. The changes also appear to be gratuitous. The only justification I can think of is to enhance performance, but they are not the performance critical part of the code.

4. The problem will show up in simple cases, e.g., polynomial fits with an independent variable range
that is much less or much larger than 1.0. The LEGENDRE keyword suggests that someone found problems
running polynomials, and assumed that the problem was simply due to the inherent problems with polynomials, and not to problems with the code itself. I don't think that they did appropriate testing of the code.

5. This problem appears to have existed ever since RSI incorporated the Numerical Recipes algorithms.
There have been complaints here about problems with SVDFIT about once or twice a year ever since. It
still hasn't been fixed.

6. The documentation has minor problems in that CHISQ returns the reduced chi-sq. and not the simple
chi-sq. and the description of COVAR, SIGMA, and VARIANCE should at least give their dimensionality

7. This is the routine that RSI recommends for general linear least-squares fitting.

Subject: Re: SVDFIT Problems
Posted by Robert Stockwell on Fri, 30 Aug 2002 21:25:43 GMT
View Forum Message <> Reply to Message

Paul van Delst wrote:
> Bill wrote:
>
>> In my occaisional attempts to use SVDFIT I have had problems in that it
>> reports singular solutions when other fitting routines, e.g., REGRESS,
>> have had not problems.
>
>
> Deja vu! I was using the SVDFIT routine just yesterday to fit a second order polynomial. I kept
> getting an error report for a singular value (even with /DOUBLE). Used POLY_FIT with no problems
> (SVDFIT _was_ overkill for what I wanted to do.)
>

> <snip>
>
>
>> Unfortunately it appeas that the code currently relies on a simplified
>> version of the Numerical Recipes routine SVDCMP, that does not return
>> the singular value vector that Numerical Recipes uses to compute wmax. A
>> rewrite to give SVDFIT the proper properties is more work than I am
>> willing to put in at this time.
>
>
> Thanks for taking the time to nut the problem out and report it here. It really pains me (just
> having finished the yearly license renewal tooth-pulling process) that these sorts of things are
> *still* found in incarnations of such well-worn linear algebra problems like SVD. I spent about
half
> a day trying to figure out what I was doing wrong until I tossed the lot and surfed over to
> netlib.org for some fortran routines. I thought IDL was supposed to *save* time for this sort of
> thing?
>
> paulv
>

I've noticed that using a cholesky decomp to solve a fit was
twice as fast as an SVD, in case any cares.

Cheers,
bob

Subject: Re: SVDFIT Problems
Posted by Robert Stockwell on Fri, 30 Aug 2002 22:47:41 GMT
View Forum Message <> Reply to Message

Do people find the posted problems with SVDFIT when one uses
the basic routines
SVDC, A, W, U, V,/double
result1 = SVSOL(U, W, V, zon,/double)
in IDL 5.5 ?

I use these routines in my solver, and the results seem to
be reasonable.

Cheers,
bob
