Subject: Re: Speaking of curve fitting...
Posted by Paul Van Delst[1] on Thu, 31 Jan 2008 17:04:53 GMT
View Forum Message <> Reply to Message

Vince Hradil wrote:
> On Jan 31, 10:17 am, Paul van Delst <Paul.vanDe...@noaa.gov> wrote:
>> Lasse Clausen wrote:
>>> ... run the following code, spot the difference and explain, s'il vous
>>> plait.
>>> nn = 1000
>>> xx1 = dindgen(nn)
>>> xx2 = timegen(nn, start=julday(5,25,1980,11,23))
>>> yy1 = sin(2.*2.*!pi*xx1/(nn-1.))
>>> d = poly_fit(xx1, yy1, 6, yfit=yfit1, /double)
>>> d = poly_fit(xx2, yy1, 6, yfit=yfit2, /double)
>> Try
>> d = poly_fit(xx2-xx2[0], yy1, 6, yfit=yfit2, /double)
>>
>>> !p.multi = [0,1,2]
>>> plot, xx1, yy1, /xstyle
>>> oplot, xx1, yfit1, linestyle=1
>>> plot, xx2, yy1,/xstyle
>>> oplot, xx2, yfit2, linestyle=1
>>> end
>>> I had a quick look at POLY_FIT.PRO but I can spot nothing which could
>>> explain the above behaviour. I run 32bit IDL 6.4 on some Linux.
>>> Cheers
>>> Lasse Clausen
>
> Sure that works, but the underlying issue is still there - why should
> it matter?

It shouldn't. But it does because...

> My guess: propagation of (roundoff) errors when poly_fit.pro
> calculates the b-matrix.

I agree that it is likely a precision/roundoff issue... the exact mechanism may be different, but I reckon you're right. An easy test would be to see what happens using SVDFIT.

cheers,

paulv