Subject: chisq value
Posted by gunvicsin11 on Mon, 16 Sep 2013 05:11:11 GMT
View Forum Message <> Reply to Message

Hello everyone,
I am trying to understand what is the chisq keyword means in the routine,
poly_fit,
svdfit
linfit

So I took a simple example, x=[1,2,3,4],y=[1,2,3,4]
when I give
p=poly_fit(x,y,1,chisq=c)
c= 1.14631e-30
p=svdfit(x,y,2,chisq=c)
c= 2.86139e-13
p=linfit(x,y,chisqr=c)
c=0.00000

for poly_fit and linfit the definition of chisq is the same
"Set this keyword to a named variable that will contain the value of the unreduced chi-square goodness-of-fit statistic"

But the chisq values are different in both these cases even though the inputvalues given are same.

Could anyone please let me know what this chisq fit actually means.

Thanking you in advance,
sid

Subject: Re: chisq value
Posted by on Mon, 16 Sep 2013 06:16:01 GMT
View Forum Message <> Reply to Message

On 2013-09-16 07:11, sid wrote:
> Hello everyone,
> I am trying to understand what is the chisq keyword means in the routine,
> poly_fit,
> svdfit
> linfit
>
> So I took a simple example, x=[1,2,3,4],y=[1,2,3,4]
> when I give
> p=poly_fit(x,y,1,chisq=c)
> c= 1.14631e-30
> p=svdfit(x,y,2,chisq=c)

> c= 2.86139e-13
> p=linfit(x,y,chisqr=c)
> c=0.00000
>
> for poly_fit and linfit the definition of chisq is the same
> "Set this keyword to a named variable that will contain the value of the unreduced chi-square goodness-of-fit statistic"
>
> But the chisq values are different in both these cases even though the inputvalues given are same.

You example data makes a "prefect" fit, so the chisq values are mostly
numerical precision errors. If you choose data that do not fit perfectly
to a straight line, the results make more sense:

```
IDL> x=[1,2,3,4]
IDL> y=[1.1,2.0,3.3,3.9]
IDL> p=poly_fit(x,y,1,chisq=c)
IDL> print,c
    0.0830000
IDL> p=svdfit(x,y,2,chisq=c)
IDL> print,c
    0.0830001
IDL> p=linfit(x,y,chisqr=c)
IDL> print,c
    0.0829998
```

> Could anyone please let me know what this chisq fit actually means.

Try this:

```
IDL> p=linfit(x,y,chisqr=c,yfit=yfit)
IDL> print,yfit
      1.12000      2.09000      3.06000      4.03000
IDL> print,total((y-yfit)^2)
    0.0829998
```

Subject: Re: chisq value
Posted by on Mon, 16 Sep 2013 06:23:33 GMT
View Forum Message <> Reply to Message

On 2013-09-16 08:16, Mats Löfdahl wrote:
> On 2013-09-16 07:11, sid wrote:
>> Hello everyone,
>> I am trying to understand what is the chisq keyword
>> means in the routine,

>> poly_fit,
>> svdfit
>> linfit
>>
>> So I took a simple example, x=[1,2,3,4],y=[1,2,3,4]
>> when I give
>> p=poly_fit(x,y,1,chisq=c)
>> c= 1.14631e-30
>> p=svdfit(x,y,2,chisq=c)
>> c= 2.86139e-13
>> p=linfit(x,y,chisqr=c)
>> c=0.00000
>>
>> for poly_fit and linfit the definition of chisq is the same
>> "Set this keyword to a named variable that will contain the value of
>> the unreduced chi-square goodness-of-fit statistic"
>>
>> But the chisq values are different in both these cases even though the
>> inputvalues given are same.
>
> You example data makes a "prefect" fit, so the chisq values are mostly
> numerical precision errors. If you choose data that do not fit perfectly
> to a straight line, the results make more sense:
>
> IDL> x=[1,2,3,4]
> IDL> y=[1.1,2.0,3.3,3.9]
> IDL> p=poly_fit(x,y,1,chisq=c)
> IDL> print,c
> 0.0830000
> IDL> p=svdfit(x,y,2,chisq=c)
> IDL> print,c
> 0.0830001
> IDL> p=linfit(x,y,chisqr=c)
> IDL> print,c
> 0.0829998
>
>
>> Could anyone please let me know what this chisq fit actually means.
>
> Try this:
>
> IDL> p=linfit(x,y,chisqr=c,yfit=yfit)
> IDL> print,yfit
> 1.12000 2.09000 3.06000 4.03000
> IDL> print,total((y-yfit)^2)
> 0.0829998

And by "prefect", I mean "perfect". :o)

Subject: Re: chisq value
Posted by Craig Markwardt on Tue, 17 Sep 2013 16:31:09 GMT
View Forum Message <> Reply to Message

On Monday, September 16, 2013 1:11:11 AM UTC-4, sid wrote:
> Hello everyone,
>
> I am trying to understand what is the chisq keyword means in the routine,
>
> poly_fit,
>
> svdfit
>
> linfit
>
>
>
> So I took a simple example, x=[1,2,3,4],y=[1,2,3,4]
>
> when I give
>
> p=poly_fit(x,y,1,chisq=c)
>
> c= 1.14631e-30
>
> p=svdfit(x,y,2,chisq=c)
>
> c= 2.86139e-13
>
> p=linfit(x,y,chisqr=c)
>
> c=0.00000
>
>
>
> for poly_fit and linfit the definition of chisq is the same
>
> "Set this keyword to a named variable that will contain the value of the unreduced chi-square goodness-of-fit statistic"
>
>
>
> But the chisq values are different in both these cases even though the inputvalues given are same.
>
>
>
> Could anyone please let me know what this chisq fit actually means.

I agree with what Mats said. This is numerical round-off error, which is computer-dependent.

An additional point is that these routines are using different levels of numerical precision. Round-off error for floating point (which is usually the default precision), is typically ~1e-7, and for double precision is ~1d-16. After you square the residuals, you get squared round-off errors of ~1e-14 and ~1d-32, which is close to some of the numbers you get.

If you force your X and Y arrays to be double precision, then the values of chi-square are much smaller.

Craig
