

ctools - Bug #3712

very high negative values for test statistic when combining observations

05/28/2021 12:53 AM - Stuani Pereira Luiz Augusto

<b>Status:</b>	Closed	<b>Start date:</b>	05/27/2021
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>	Knödseder Jürgen	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>			
<b>Description</b> I am running ctlike to calculate the test statistic (TS) for an unbinned combined observation (two different observations of the same source) of a source with weak flux and all the time I run ctlike I get a different TS value and it is always lower than the TS calculated for the observations individually. Moreover, sometimes the optimization status converges and other times it stalls (this also happens for sources with high detection significance). Furthermore, I often get very high negative TS values such as -204298. I am using ctools version 1.7.4 and calibration database prod3b-v2 for the simulation. I appended a table with the TS values for a hypothetical source after running ctlike 20 times.			

History

#1 - 05/28/2021 02:42 AM - Knödseder Jürgen

- Status changed from New to In Progress
- Assigned To set to Knödseder Jürgen

Could you please provide a bit more information so that I can reproduce your problem? I would need the observation and model definition XML files that you use for the simulations and model fitting, as well as a typical ctlike output, once without problems and once with problems.

When you say that you get different results with ctlike, are you changing the seed value for the simulations between the runs? Computations should always be fully reproducible, i.e. with the same input parameters you always get the same results.

#2 - 05/28/2021 12:36 PM - Stuani Pereira Luiz Augusto

- File spec\_src2330.txt added
- File obs.xml added
- File model\_TS.xml added
- File model.xml added
- File logfile\_stalled\_low\_negative\_TS.log added
- File logfile\_stalled\_high\_negative\_TS.log added
- File logfile\_converged\_example2.log added
- File logfile\_converged\_example1.log added
- File ctobssim.py added
- File ctlike.py added

I am using two different python files for the simulation: ctobssim.py and ctlike.py  
First I run ctobssim.py to generate event1.fits and event2.fits and after I create an observation definition file (obs.xml) and use it as an input in ctlike.py file, which I run 20 times. I did not change the seed in ctobssim.py, I . When I do the same analysis procedure but for a single observation I always get the same TS value.  
I am sending the files I used for simulation. The model file I use in ctobssim is called model.xml (the file input is a customized flux called spec\_src2330.txt) and the one I use in ctlike is model\_TS.xml.  
I am also sending three logfiles: two for converged fitting, one for a stalled fitting with low negative TS value and another for a stalled fitting with very high negative TS value.  
If you need any additional information just let me know.  
Thank you very much for your help.

### #3 - 05/28/2021 01:53 PM - Knödseder Jürgen

- % Done changed from 0 to 10

Thanks for posting the files. I could download and use the files, yet I could not reproduce the problem. Every time I run `ctlike.py` I get exactly the same result, as I would have expected. Here the first few lines of the fitting for a first run:

```
2021-05-28T10:56:06: +=====+
2021-05-28T10:56:06: | Maximum likelihood optimisation |
2021-05-28T10:56:06: +=====+
2021-05-28T10:56:07: >Iteration 0: -logL=1482911.076, Lambda=1.0e-03
2021-05-28T10:56:07: Parameter "Prefactor" drives optimization step (step=0.00167941)
2021-05-28T10:56:07: >Iteration 1: -logL=1282289.201, Lambda=1.0e-03, delta=200621.876, step=1.7e-03, max(|grad|)=-2965.509270
[Prefactor:2]
2021-05-28T10:56:07: Parameter "Prefactor" does not drive optimization step anymore.
2021-05-28T10:56:07: Parameter "Index" drives optimization step (step=0.000757938)
```

and here for a second run

```
2021-05-28T10:57:19: +=====+
2021-05-28T10:57:19: | Maximum likelihood optimisation |
2021-05-28T10:57:19: +=====+
2021-05-28T10:57:20: >Iteration 0: -logL=1482911.076, Lambda=1.0e-03
2021-05-28T10:57:20: Parameter "Prefactor" drives optimization step (step=0.00167941)
2021-05-28T10:57:20: >Iteration 1: -logL=1282289.201, Lambda=1.0e-03, delta=200621.876, step=1.7e-03, max(|grad|)=-2965.509270
[Prefactor:2]
2021-05-28T10:57:20: Parameter "Prefactor" does not drive optimization step anymore.
2021-05-28T10:57:20: Parameter "Index" drives optimization step (step=0.000757938)
```

I noticed in your files that the results for Iteration 0 differs in one of the log files, which it should not since Iteration 0 computes only the log-likelihood for the input model.

Are you sure that you always feed the same model on input? You can check this with putting `chatter=4`.

I did the test using the current development version of the tools, I still need to make some tests on `ctools 1.7.3`.

#### #4 - 05/28/2021 02:09 PM - Knödseder Jürgen

Just did the test also for 1.7.3, again results are identical. Just for reference, here are the fit results that I found in every run:

```
2021-05-28T12:01:24: +=====+
2021-05-28T12:01:24: | Maximum likelihood optimisation results |
2021-05-28T12:01:24: +=====+
2021-05-28T12:01:24: === GOptimizerLM ===
2021-05-28T12:01:24: Optimized function value ...: 1282287.561
2021-05-28T12:01:24: Absolute precision .....: 0.005
2021-05-28T12:01:24: Acceptable value decrease ..: 2
2021-05-28T12:01:24: Optimization status .....: converged
2021-05-28T12:01:24: Number of parameters .....: 10
2021-05-28T12:01:24: Number of free parameters ..: 4
2021-05-28T12:01:24: Number of iterations .....: 14
2021-05-28T12:01:24: Lambda .....: 1e-05
2021-05-28T12:01:24: Maximum log likelihood ....: -1282287.561
2021-05-28T12:01:24: Observed events (Nobs) ...: 253794.000
2021-05-28T12:01:24: Predicted events (Npred) ..: 253793.634 (Nobs - Npred = 0.365719267167151)
2021-05-28T12:01:24: === GModels ===
2021-05-28T12:01:24: Number of models .....: 2
2021-05-28T12:01:24: Number of parameters .....: 10
2021-05-28T12:01:24: === GModelSky ===
2021-05-28T12:01:24: Name .....: src2330
2021-05-28T12:01:24: Instruments .....: all
2021-05-28T12:01:24: Test Statistic .....: 1.52692038984969
2021-05-28T12:01:24: Observation identifiers ...: all
2021-05-28T12:01:24: Model type .....: PointSource
2021-05-28T12:01:24: Model components .....: "PointSource" * "PowerLaw" * "Constant"
2021-05-28T12:01:24: Number of parameters .....: 6
2021-05-28T12:01:24: Number of spatial par's ...: 2
2021-05-28T12:01:24: RA .....: 184.5492 [-360,360] deg (fixed,scale=1)
2021-05-28T12:01:24: DEC .....: -5.8137 [-90,90] deg (fixed,scale=1)
2021-05-28T12:01:24: Number of spectral par's ...: 3
2021-05-28T12:01:24: Prefactor .....: 1.02729682946507e-20 +/- 3.5428155307913e-20 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2021-05-28T12:01:24: Index .....: -3.8508026662252 +/- 1.89455693104743 [-5,-0] (free,scale=-1,gradient)
2021-05-28T12:01:24: PivotEnergy .....: 1000000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2021-05-28T12:01:24: Number of temporal par's ...: 1
2021-05-28T12:01:24: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
2021-05-28T12:01:24: Number of scale par's .....: 0
2021-05-28T12:01:24: === GCTAModelIrfBackground ===
2021-05-28T12:01:24: Name .....: Background
2021-05-28T12:01:24: Instruments .....: CTA
2021-05-28T12:01:24: Observation identifiers ...: all
2021-05-28T12:01:24: Model type .....: "PowerLaw" * "Constant"
2021-05-28T12:01:24: Number of parameters .....: 4
2021-05-28T12:01:24: Number of spectral par's ...: 3
2021-05-28T12:01:24: Prefactor .....: 0.996979979887099 +/- 0.00384627869042233 [0.001,1000] ph/cm2/s/MeV (free,scale=1,gradient)
2021-05-28T12:01:24: Index .....: -0.00262329024370695 +/- 0.00220344906702262 [-5,5] (free,scale=1,gradient)
2021-05-28T12:01:24: PivotEnergy .....: 1000000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2021-05-28T12:01:24: Number of temporal par's ...: 1
2021-05-28T12:01:24: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)
```

As a side note, results were obtained on a MacBook running Mac OS 10.14.6.

**#5 - 05/28/2021 03:11 PM - Stuari Pereira Luiz Augusto**

Thank you very much for running my codes. I think I found the reason why I am getting different results. I am running ctools 1.7.4 via conda in Linux Ubuntu and now I used the same version installed without conda and I get the same results as you did. So I think the problem may be related to anaconda.

**#6 - 05/28/2021 03:22 PM - Knödlseider Jürgen**

This is weird, why should anaconda impact the results? Which Ubuntu version are you using? And what is the hardware that you are using?

By the way: from your log files I see that you ran 1.7.3, but again, I had no problems with that version.

**#7 - 05/28/2021 06:37 PM - Stuari Pereira Luiz Augusto**

I am sorry I was not clear on my last answer. For the problems I got I was running ctools 1.7.3 via conda in Linux Mint 17.2 and the last test I did (which worked well) was without conda in a CentOS Linux release 7.6.1810. I am going to install ctools without conda in my Linux Mint to check if it works properly.

**#8 - 05/28/2021 10:05 PM - Stuari Pereira Luiz Augusto**

I installed ctools 1.7.3 without conda in Linux Mint 17.2 and I get the same TS problem as before. So conda is not the problem and it may be related to Linux Mint. But this fitting problem only occurs when combining observations. Thank you very much for all your help.

**#9 - 05/29/2021 10:34 AM - Knödlseider Jürgen**

I don't test Linux Mint so far, but from their web page I see it should be based on Debian and Ubuntu. Having non-reproducible results can point towards a memory leak, I will try to setup a Linux Mint 17.2 system on my side to see whether I can reproduce the problem.

Could you please indicate which version from <https://www.linuxmint.com/release.php?id=25> you have installed?

**#10 - 05/29/2021 12:31 PM - Stuari Pereira Luiz Augusto**

```
RELEASE=17.2
CODENAME=rafaela
EDITION="Cinnamon 64-bit"
DESCRIPTION="Linux Mint 17.2 Rafaela"
DESKTOP=Gnome
TOOLKIT=GTK
NEW_FEATURES_URL=http://www.linuxmint.com/rel_rafaela_cinnamon_whatsnew.php
RELEASE_NOTES_URL=http://www.linuxmint.com/rel_rafaela_cinnamon.php
USER_GUIDE_URL=help:linuxmint
GRUB_TITLE=Linux Mint 17.2 Cinnamon 64-bit
```

Thank you very much for your help

**#11 - 05/29/2021 09:06 PM - Knödlseider Jürgen**

- % Done changed from 10 to 50

I installed Linux Mint 17.2 on a VM and installed ctools version 1.7.3. I repeated your test and I always get the same result. So I really don't

understand what happens on your system.

```
2021-05-29T15:36:12: +=====+
2021-05-29T15:36:12: | Maximum likelihood optimisation |
2021-05-29T15:36:12: +=====+
2021-05-29T15:36:14: >Iteration 0: -logL=1482911.076, Lambda=1.0e-03
2021-05-29T15:36:14: Parameter "Prefactor" drives optimization step (step=0.00167941)
2021-05-29T15:36:16: >Iteration 1: -logL=1282289.201, Lambda=1.0e-03, delta=200621.876, step=1.7e-03, max(|grad|)=-2965.509270
[Prefactor:2]
2021-05-29T15:36:16: Parameter "Prefactor" does not drive optimization step anymore.
2021-05-29T15:36:16: Parameter "Index" drives optimization step (step=0.000757938)
2021-05-29T15:36:17: >Iteration 2: -logL=1282289.186, Lambda=1.0e-04, delta=0.015, step=7.6e-04, max(|grad|)=-126562.237505 [Prefactor:2]
2021-05-29T15:36:17: Parameter "Index" does not drive optimization step anymore.
2021-05-29T15:36:17: Parameter "Index" hits minimum: -178.838 < 0 (1)
2021-05-29T15:36:19: Iteration 3: -logL=1282289.186, Lambda=1.0e-05, delta=-86.835, step=1.0e+00, max(|grad|)=1746031.166814 [Prefactor:2]
(stalled)
2021-05-29T15:36:19: Parameter "Index" drives optimization step (step=0.0275381)
2021-05-29T15:36:20: Iteration 4: -logL=1282289.186, Lambda=1.0e-04, delta=-2.518, step=2.8e-02, max(|grad|)=1717878.323001 [Prefactor:2]
(stalled)
2021-05-29T15:36:20: Parameter "Index" does not drive optimization step anymore.
2021-05-29T15:36:20: Parameter "Index" hits minimum: -156.352 < 0 (1)
2021-05-29T15:36:22: Iteration 5: -logL=1282289.186, Lambda=1.0e-03, delta=-78.008, step=1.0e+00, max(|grad|)=1744401.130097 [Prefactor:2]
(stalled)
2021-05-29T15:36:22: Parameter "Index" drives optimization step (step=0.0709794)
2021-05-29T15:36:23: Iteration 6: -logL=1282289.186, Lambda=1.0e-02, delta=-3.130, step=7.1e-02, max(|grad|)=1718319.629436 [Prefactor:2]
(stalled)
2021-05-29T15:36:23: Parameter "Index" does not drive optimization step anymore.
2021-05-29T15:36:23: Parameter "Index" hits maximum: 9.51785 > 5 (1)
2021-05-29T15:36:24: >Iteration 7: -logL=1282287.846, Lambda=1.0e-01, delta=1.339, step=1.0e+00, max(|grad|)=-27623.590493 [Prefactor:2]
2021-05-29T15:36:26: Iteration 8: -logL=1282287.979, Lambda=1.0e-02, delta=-0.133, step=1.0e+00, max(|grad|)=-10123.301165 [Prefactor:2]
(stalled)
2021-05-29T15:36:27: Iteration 9: -logL=1282288.747, Lambda=1.0e-01, delta=-0.768, step=1.0e+00, max(|grad|)=55395.297199 [Prefactor:2]
(stalled)
2021-05-29T15:36:29: >Iteration 10: -logL=1282287.761, Lambda=1.0e+00, delta=0.986, step=1.0e+00, max(|grad|)=17157.430662 [Prefactor:2]
2021-05-29T15:36:30: >Iteration 11: -logL=1282287.583, Lambda=1.0e-01, delta=0.178, step=1.0e+00, max(|grad|)=1308.947277 [Prefactor:2]
2021-05-29T15:36:32: >Iteration 12: -logL=1282287.577, Lambda=1.0e-02, delta=0.007, step=1.0e+00, max(|grad|)=-1987.788142 [Prefactor:2]
2021-05-29T15:36:33: >Iteration 13: -logL=1282287.561, Lambda=1.0e-03, delta=0.015, step=1.0e+00, max(|grad|)=219.433700 [Prefactor:2]
2021-05-29T15:36:35: >Iteration 14: -logL=1282287.561, Lambda=1.0e-04, delta=0.000, step=1.0e+00, max(|grad|)=-160.321683 [Prefactor:2]
2021-05-29T15:36:38: +=====+
2021-05-29T15:36:38: | Maximum likelihood optimisation results |
2021-05-29T15:36:38: +=====+
2021-05-29T15:36:38: === GOptimizerLM ===
2021-05-29T15:36:38: Optimized function value ..: 1282287.561
2021-05-29T15:36:38: Absolute precision ..: 0.005
2021-05-29T15:36:38: Acceptable value decrease ..: 2
2021-05-29T15:36:38: Optimization status ..: converged
2021-05-29T15:36:38: Number of parameters ..: 10
2021-05-29T15:36:38: Number of free parameters ..: 4
2021-05-29T15:36:38: Number of iterations ..: 14
2021-05-29T15:36:38: Lambda ..: 1e-05
2021-05-29T15:36:38: Maximum log likelihood ..: -1282287.561
2021-05-29T15:36:38: Observed events (Nobs) ..: 253794.000
2021-05-29T15:36:38: Predicted events (Npred) ..: 253793.634 (Nobs - Npred = 0.365718831191771)
2021-05-29T15:36:38: === GModels ===
2021-05-29T15:36:38: Number of models ..: 2
2021-05-29T15:36:38: Number of parameters ..: 10
2021-05-29T15:36:38: === GModelSky ===
2021-05-29T15:36:38: Name ..: src2330
2021-05-29T15:36:38: Instruments ..: all
2021-05-29T15:36:38: Test Statistic ..: 1.52692040009424
2021-05-29T15:36:38: Observation identifiers ..: all
2021-05-29T15:36:38: Model type ..: PointSource
2021-05-29T15:36:38: Model components ..: "PointSource" * "PowerLaw" * "Constant"
2021-05-29T15:36:38: Number of parameters ..: 6
2021-05-29T15:36:38: Number of spatial par's ..: 2
2021-05-29T15:36:38: RA ..: 184.5492 [-360,360] deg (fixed,scale=1)
2021-05-29T15:36:38: DEC ..: -5.8137 [-90,90] deg (fixed,scale=1)
2021-05-29T15:36:38: Number of spectral par's ..: 3
2021-05-29T15:36:38: Prefactor ..: 1.02729682942727e-20 +/- 3.54281553081389e-20 [1e-23,1e-13] ph/cm2/s/MeV
(free,scale=1e-16,gradient)
2021-05-29T15:36:38: Index ..: -3.85080266621993 +/- 1.89455693112968 [-5,-0] (free,scale=-1,gradient)
2021-05-29T15:36:38: PivotEnergy ..: 1000000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)
2021-05-29T15:36:38: Number of temporal par's ..: 1
```

2021-05-29T15:36:38: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)  
 2021-05-29T15:36:38: Number of scale par's .....: 0  
 2021-05-29T15:36:38: === GCTAModellrfBackground ===  
 2021-05-29T15:36:38: Name .....: Background  
 2021-05-29T15:36:38: Instruments .....: CTA  
 2021-05-29T15:36:38: Observation identifiers ...: all  
 2021-05-29T15:36:38: Model type .....: "PowerLaw" \* "Constant"  
 2021-05-29T15:36:38: Number of parameters .....: 4  
 2021-05-29T15:36:38: Number of spectral par's ...: 3  
 2021-05-29T15:36:38: Prefactor .....: 0.996979979896491 +/- 0.00384627869046336 [0.001,1000] ph/cm2/s/MeV (free,scale=1,gradient)  
 2021-05-29T15:36:38: Index .....: -0.00262329023858107 +/- 0.00220344906702464 [-5,5] (free,scale=1,gradient)  
 2021-05-29T15:36:38: PivotEnergy .....: 1000000 [10000,1000000000] MeV (fixed,scale=1000000,gradient)  
 2021-05-29T15:36:38: Number of temporal par's ...: 1  
 2021-05-29T15:36:38: Normalization .....: 1 (relative value) (fixed,scale=1,gradient)

#### #12 - 05/30/2021 05:35 PM - Stuari Pereira Luiz Augusto

Thank you very much for doing this test. There might be then a problem with my operational system.

#### #13 - 05/31/2021 09:14 AM - Knödlseider Jürgen

- Status changed from In Progress to Closed

- % Done changed from 50 to 100

I close the issue now since It was not able to reproduce it. In case that you do encounter new problems, or if you have some additional information, please do not hesitate to signal it. I can then reopen the issue.

#### Files

test_statistic_table.png	44.2 KB	05/27/2021	Stuari Pereira Luiz Augusto
spec_src2330.txt	2.15 KB	05/28/2021	Stuari Pereira Luiz Augusto
obs.xml	376 Bytes	05/28/2021	Stuari Pereira Luiz Augusto
model_TS.xml	1.1 KB	05/28/2021	Stuari Pereira Luiz Augusto
model.xml	1.07 KB	05/28/2021	Stuari Pereira Luiz Augusto
logfile_stalled_low_negative_TS.log	16 KB	05/28/2021	Stuari Pereira Luiz Augusto
logfile_stalled_high_negative_TS.log	14.9 KB	05/28/2021	Stuari Pereira Luiz Augusto
logfile_converged_example2.log	15 KB	05/28/2021	Stuari Pereira Luiz Augusto
logfile_converged_example1.log	14.7 KB	05/28/2021	Stuari Pereira Luiz Augusto
ctobssim.py	830 Bytes	05/28/2021	Stuari Pereira Luiz Augusto
ctlake.py	350 Bytes	05/28/2021	Stuari Pereira Luiz Augusto