ctools - Bug #1074

Optimizer does not converge for parameters close to truth

01/09/2014 11:10 PM - Knödlseder Jürgen

Status:	Closed		**Start date:**	01/09/2014
Priority:	Urgent		**Due date:**	
Assigned To:	Knödlseder Jürgen		**% Done:**	100%
Category:			**Estimated time:**	0.00 hour
Target version:	1.0.0			

Description

Using the "Getting started example" of the ctools page, the unbinned ctlike fit did not properly converge for the following run:

$ ctlike
Event list, counts map or observation definition file [selected_events.fits]
Calibration database [$GAMMALIB/share/caldb/cta]
Instrument response function [cta_dummy_irf]
Source model [crab.xml] $GAMMALIB/share/models/crab.xml
Source model output file [crab_results.xml]

The resulting log file is:

```
2014-01-09T21:49:39:  Initial iteration: func=44391.3, Lambda=0.001
2014-01-09T21:49:39:  Iteration 1: func=44391.3, Lambda=0.01, delta=0, max(grad)=9.59982 [6] (stalled)
2014-01-09T21:49:39:  Iteration 2: func=44391.3, Lambda=0.1, delta=0, max(grad)=8.93495 [6] (stalled)
2014-01-09T21:49:39:  Iteration 3: func=44391.3, Lambda=1, delta=0, max(grad)=3.15298 [6] (stalled)
2014-01-09T21:49:39:  Iteration 4: func=44388.3, Lambda=0.1, delta=2.9745, max(grad)=0 [-1]
2014-01-09T21:49:39:  Iteration 5: func=44388.3, Lambda=1, delta=0, max(grad)=8.9891 [6] (stalled)
2014-01-09T21:49:39:  Iteration 6: func=44387.4, Lambda=0.1, delta=0.897666, max(grad)=2.69155 [7]
2014-01-09T21:49:39:  Iteration 7: func=44387.4, Lambda=1, delta=0, max(grad)=11.4893 [6] (stalled)
2014-01-09T21:49:39:  Iteration 8: func=44387.4, Lambda=10, delta=0, max(grad)=12.8144 [7] (stalled)
2014-01-09T21:49:39:  Iteration 9: func=44387.4, Lambda=1, delta=0.0603885, max(grad)=4.33702 [7]
2014-01-09T21:49:39:  Iteration 10: func=44387.4, Lambda=10, delta=0, max(grad)=13.2321 [7] (stalled)
2014-01-09T21:49:39:  Iteration 11: func=44387.3, Lambda=1, delta=0.0520817, max(grad)=5.68515 [7]
2014-01-09T21:49:39:  Iteration 12: func=44387.3, Lambda=10, delta=0, max(grad)=13.5394 [7] (stalled)
2014-01-09T21:49:39:  Iteration 13: func=44387.3, Lambda=1, delta=0.0451309, max(grad)=6.77871 [7]
2014-01-09T21:49:39:  Iteration 14: func=44387.3, Lambda=10, delta=0, max(grad)=13.7533 [7] (stalled)
2014-01-09T21:49:39:  Iteration 15: func=44387.3, Lambda=100, delta=0, max(grad)=12.7579 [7] (stalled)
2014-01-09T21:49:39:  Iteration 16: func=44387.3, Lambda=10, delta=0.00431746, max(grad)=12741.4 [6]
2014-01-09T21:49:39:  Iteration 17: func=44387.3, Lambda=100, delta=0, max(grad)=0 [-1] (stalled)
2014-01-09T21:49:39:  Iteration 18: func=44387.3, Lambda=1000, delta=0, max(grad)=0 [-1] (stalled)
2014-01-09T21:49:39:  Iteration 19: func=44387.3, Lambda=10000, delta=0, max(grad)=0 [-1] (stalled)
2014-01-09T21:49:39:  Iteration 20: func=44387.3, Lambda=100000, delta=0, max(grad)=5.46593 [7] (stalled)
2014-01-09T21:49:39:  Iteration 21: func=44387.3, Lambda=1e+06, delta=0, max(grad)=6.74412 [7] (stalled)
2014-01-09T21:49:39:  Iteration 22: func=44387.3, Lambda=1e+07, delta=0, max(grad)=12741.1 [6] (stalled)
2014-01-09T21:49:40:  Iteration 23: func=44387.3, Lambda=1e+08, delta=0, max(grad)=12741.4 [6] (stalled)
2014-01-09T21:49:40:  Iteration 24: func=44387.3, Lambda=1e+09, delta=0, max(grad)=12741.4 [6] (stalled)
2014-01-09T21:49:40:  Iteration 25: func=44387.3, Lambda=1e+10, delta=0, max(grad)=12741.4 [6] (stalled)
2014-01-09T21:49:40:  Iteration 26: func=44387.3, Lambda=1e+11, delta=0, max(grad)=12741.4 [6] (stalled)
```

Related issues:

Related to GammaLib - Feature # 311: Tune Levenberg-Marquard optimizer	**Closed**	**07/13/2012**

History

#1 - 10/30/2014 12:22 PM - Knödlseder Jürgen

- Priority changed from Normal to High

- Target version set to 1.0.0

This definitely should be investigate before we get the release 1.0.0 out.

#2 - 10/31/2014 12:20 AM - Knödlseder Jürgen

If you follow the quickstart example from the documentation, doing an unbinned ctlike analysis produces exactly the problem:

```
2014-10-30T22:55:20: +=================================+
2014-10-30T22:55:20: | Maximum likelihood optimisation |
2014-10-30T22:55:20: +=================================+
2014-10-30T22:55:20: >Iteration   0: -logL=44395.244, Lambda=1.0e-03
2014-10-30T22:55:20:  Iteration   1: -logL=44395.244, Lambda=1.0e-03, delta=-0.979, max(|grad|)=9.644743 [Sigma:6] (stalled)
2014-10-30T22:55:20:  Iteration   2: -logL=44395.244, Lambda=1.0e-02, delta=-0.976, max(|grad|)=8.980893 [Sigma:6] (stalled)
2014-10-30T22:55:20:  Iteration   3: -logL=44395.244, Lambda=1.0e-01, delta=-0.997, max(|grad|)=-4.897460 [Index:3] (stalled)
2014-10-30T22:55:20: >Iteration   4: -logL=44392.328, Lambda=1.0e+00, delta=2.916, max(|grad|)=-37.029192 [Sigma:6]
2014-10-30T22:55:20:  Iteration   5: -logL=44392.328, Lambda=1.0e-01, delta=-3.884, max(|grad|)=8.963805 [Sigma:6] (stalled)
2014-10-30T22:55:20: >Iteration   6: -logL=44391.437, Lambda=1.0e+00, delta=0.891, max(|grad|)=-19.625468 [Sigma:6]
2014-10-30T22:55:20:  Iteration   7: -logL=44391.437, Lambda=1.0e-01, delta=-4.781, max(|grad|)=11.452963 [Sigma:6] (stalled)
2014-10-30T22:55:20:  Iteration   8: -logL=44391.437, Lambda=1.0e+00, delta=-4.840, max(|grad|)=13.273568 [Normalization:7] (stalled)
2014-10-30T22:55:20: >Iteration   9: -logL=44391.377, Lambda=1.0e+01, delta=0.060, max(|grad|)=-18.011138 [Sigma:6]
2014-10-30T22:55:20:  Iteration  10: -logL=44391.377, Lambda=1.0e+00, delta=-4.888, max(|grad|)=13.655269 [Normalization:7] (stalled)
2014-10-30T22:55:20: >Iteration  11: -logL=44391.325, Lambda=1.0e+01, delta=0.052, max(|grad|)=-16.578132 [Sigma:6]
2014-10-30T22:55:20:  Iteration  12: -logL=44391.325, Lambda=1.0e+00, delta=-4.930, max(|grad|)=13.930485 [Normalization:7] (stalled)
2014-10-30T22:55:20: >Iteration  13: -logL=44391.279, Lambda=1.0e+01, delta=0.045, max(|grad|)=-15.303160 [Sigma:6]
2014-10-30T22:55:20:  Iteration  14: -logL=44391.279, Lambda=1.0e+00, delta=-4.967, max(|grad|)=14.115798 [Normalization:7] (stalled)
2014-10-30T22:55:20:  Iteration  15: -logL=44391.279, Lambda=1.0e+01, delta=-5.070, max(|grad|)=13.352603 [Normalization:7] (stalled)
2014-10-30T22:55:20: >Iteration  16: -logL=44391.275, Lambda=1.0e+02, delta=0.004, max(|grad|)=-15.176887 [Sigma:6]
2014-10-30T22:55:20:
2014-10-30T22:55:20: +======================================+
2014-10-30T22:55:20: | Maximum likelihood optimization results |
2014-10-30T22:55:20: +======================================+
2014-10-30T22:55:20: === GOptimizerLM ===
2014-10-30T22:55:20:  Optimized function value ..: 44391.271
2014-10-30T22:55:20:  Absolute precision ........: 0.005
2014-10-30T22:55:20:  Optimization status .......: converged
2014-10-30T22:55:20:  Number of parameters ......: 9
2014-10-30T22:55:20:  Number of free parameters .: 4
2014-10-30T22:55:20:  Number of iterations ......: 6
2014-10-30T22:55:20:  Lambda ....................: 10
2014-10-30T22:55:20:  Maximum log likelihood ....: -44391.271
2014-10-30T22:55:20:  Observed events  (Nobs) ...: 6127.000
2014-10-30T22:55:20:  Predicted events (Npred) ..: 6120.693 (Nobs - Npred = 6.3066)
```

When I then changed the spectral slope from the initial value of -2.48 to -2.41 I got:

```
2014-10-30T22:55:43: +=================================+
2014-10-30T22:55:43: | Maximum likelihood optimisation |
2014-10-30T22:55:43: +=================================+
2014-10-30T22:55:43: >Iteration   0: -logL=44400.561, Lambda=1.0e-03
2014-10-30T22:55:43: >Iteration   1: -logL=44396.253, Lambda=1.0e-03, delta=4.308, max(|grad|)=9.645475 [Sigma:6]
2014-10-30T22:55:43: >Iteration   2: -logL=44396.162, Lambda=1.0e-04, delta=0.091, max(|grad|)=-0.336139 [Sigma:6]
2014-10-30T22:55:43: >Iteration   3: -logL=44396.162, Lambda=1.0e-05, delta=0.000, max(|grad|)=0.005603 [Sigma:6]
2014-10-30T22:55:43:
2014-10-30T22:55:43: +======================================+
2014-10-30T22:55:43: | Maximum likelihood optimization results |
2014-10-30T22:55:43: +======================================+
2014-10-30T22:55:43: === GOptimizerLM ===
2014-10-30T22:55:43:  Optimized function value ..: 44396.162
2014-10-30T22:55:43:  Absolute precision ........: 0.005
2014-10-30T22:55:43:  Optimization status .......: converged
2014-10-30T22:55:43:  Number of parameters ......: 9
2014-10-30T22:55:43:  Number of free parameters .: 4
2014-10-30T22:55:43:  Number of iterations ......: 3
2014-10-30T22:55:43:  Lambda ....................: 1e-06
2014-10-30T22:55:43:  Maximum log likelihood ....: -44396.162
2014-10-30T22:55:43:  Observed events  (Nobs) ...: 6127.000
2014-10-30T22:55:43:  Predicted events (Npred) ..: 6127.000 (Nobs - Npred = 0.000141861)
```

#3 - 10/31/2014 12:21 AM - Knödlseder Jürgen

- Priority changed from High to Urgent

#4 - 10/31/2014 03:45 PM - Knödlseder Jürgen

- Status changed from New to Feedback

- Assigned To set to Knödlseder Jürgen

- % Done changed from 0 to 90

Allowing to continue iterations in case of a slight log-likelihood decrease solves the issue. I did this allow already before, but the threshold was set to 1e-6, hence a very tiny decrease. Now I allow for a likelihood decrease one 1.0 which may help to get out of a local minimum that hampers convergence. I added for this purpose an accept_dec() method to GOptimizerLM that allows for a user adjustment of the threshold (by default set to 1.0). Below the ctlike results with this new approach:


```
2014-10-31T14:45:29: +=================================+
2014-10-31T14:45:29: | Maximum likelihood optimisation |
2014-10-31T14:45:29: +=================================+
2014-10-31T14:45:29: >Iteration   0: -logL=44395.244, Lambda=1.0e-03
2014-10-31T14:45:29:  Iteration   1: -logL=44396.223, Lambda=1.0e-03, delta=-0.979, max(|grad|)=9.644743 [Sigma:6] (stalled)
2014-10-31T14:45:29: >Iteration   2: -logL=44396.162, Lambda=1.0e-02, delta=0.061, max(|grad|)=-0.309683 [Normalization:7]
2014-10-31T14:45:29: >Iteration   3: -logL=44396.162, Lambda=1.0e-03, delta=0.000, max(|grad|)=-0.008931 [Index:3]
2014-10-31T14:45:29:
2014-10-31T14:45:29: +========================================+
2014-10-31T14:45:29: | Maximum likelihood optimization results |
2014-10-31T14:45:29: +========================================+
2014-10-31T14:45:29: === GOptimizerLM ===
2014-10-31T14:45:29:  Optimized function value ..: 44396.162
2014-10-31T14:45:29:  Absolute precision ........: 0.005
2014-10-31T14:45:29:  Acceptable function decrease : 1
2014-10-31T14:45:29:  Optimization status .......: converged
2014-10-31T14:45:29:  Number of parameters ......: 9
2014-10-31T14:45:29:  Number of free parameters .: 4
2014-10-31T14:45:29:  Number of iterations ......: 3
2014-10-31T14:45:29:  Lambda ....................: 0.0001
2014-10-31T14:45:29:  Maximum log likelihood ....: -44396.162
2014-10-31T14:45:29:  Observed events  (Nobs) ...: 6127.000
2014-10-31T14:45:29:  Predicted events (Npred) ..: 6127.000 (Nobs - Npred = 0.000303534)
2014-10-31T14:45:29: === GModels ===
2014-10-31T14:45:29:  Number of models ..........: 2
2014-10-31T14:45:29:  Number of parameters ......: 9
2014-10-31T14:45:29: === GModelSky ===
2014-10-31T14:45:29:  Name ......................: Crab
2014-10-31T14:45:29:  Instruments ...............: all
2014-10-31T14:45:29:  Instrument scale factors ..: unity
2014-10-31T14:45:29:  Observation identifiers ...: all
2014-10-31T14:45:29:  Model type ................: PointSource
2014-10-31T14:45:29:  Model components ..........: "SkyDirFunction" * "PowerLaw" * "Constant"
2014-10-31T14:45:29:  Number of parameters ......: 6
2014-10-31T14:45:29:  Number of spatial par's ...: 2
2014-10-31T14:45:29:   RA .......................: 83.6331 [-360,360] deg (fixed,scale=1)
2014-10-31T14:45:29:   DEC ......................: 22.0145 [-90,90] deg (fixed,scale=1)
2014-10-31T14:45:29:  Number of spectral par's ..: 3
2014-10-31T14:45:29:   Prefactor ...............: 6.13325e-16 +/- 2.05738e-17 [1e-23,1e-13] ph/cm2/s/MeV (free,scale=1e-16,gradient)
2014-10-31T14:45:29:   Index ...................: -2.5057 +/- 0.0250818 [-0,-5]  (free,scale=-1,gradient)
2014-10-31T14:45:29:   PivotEnergy .............: 300000 [10000,1e+09] MeV (fixed,scale=1e+06,gradient)
2014-10-31T14:45:29:  Number of temporal par's ..: 1
2014-10-31T14:45:29:   Constant ................: 1 (relative value) (fixed,scale=1,gradient)
2014-10-31T14:45:29: === GCTAModelRadialAcceptance ===
2014-10-31T14:45:29:  Name ......................: Background
2014-10-31T14:45:29:  Instruments ...............: CTA
2014-10-31T14:45:29:  Instrument scale factors ..: unity
2014-10-31T14:45:29:  Observation identifiers ...: all
2014-10-31T14:45:29:  Model type ................: "Gaussian" * "FileFunction" * "Constant"
2014-10-31T14:45:29:  Number of parameters ......: 3
2014-10-31T14:45:29:  Number of radial par's ....: 1
2014-10-31T14:45:29:   Sigma ...................: 3.04429 +/- 0.0330227 [0.01,10] deg2 (free,scale=1,gradient)
2014-10-31T14:45:29:  Number of spectral par's ..: 1
2014-10-31T14:45:29:   Normalization ...........: 0.996767 +/- 0.0175227 [0,1000]  (free,scale=1,gradient)
2014-10-31T14:45:29:  Number of temporal par's ..: 1
2014-10-31T14:45:29:   Constant ................: 1 (relative value) (fixed,scale=1,gradient)
2014-10-31T14:45:29:
```

2014-10-31T14:45:29: +==============+
2014-10-31T14:45:29: | Save results |
2014-10-31T14:45:29: +==============+
2014-10-31T14:45:29:
2014-10-31T14:45:29:
2014-10-31T14:45:29: Application "ctlike" terminated after 2 wall clock seconds, consuming 0.255421 seconds of CPU time.

#5 - 10/31/2014 03:50 PM - Knödlseder Jürgen

Keep this under feedback to see how this new feature works in real live.

#6 - 11/03/2014 09:41 AM - Knödlseder Jürgen

- % Done changed from 90 to 50

The new behavior leads to a problem when fitting the diffuse model in unbinned mode. Here the output:

```
+================================================+
| Unbinned maximum likelihood fitting of CTA data |
+================================================+
>Iteration   0: -logL=32632.072, Lambda=1.0e-03
>Iteration   1: -logL=32632.022, Lambda=1.0e-03, delta=0.050, max(|grad|)=-70.083938 [Index:7]
 Iteration   2: -logL=32632.540, Lambda=1.0e-04, delta=-0.518, max(|grad|)=-30.223946 [Index:2] (stalled)
 Iteration   3: -logL=32633.289, Lambda=1.0e-03, delta=-0.749, max(|grad|)=-80.038846 [Index:7] (stalled)
>Iteration   4: -logL=32631.386, Lambda=1.0e-02, delta=1.903, max(|grad|)=-21.672392 [Index:2]
 Iteration   5: -logL=32632.256, Lambda=1.0e-03, delta=-0.870, max(|grad|)=-63.668797 [Index:7] (stalled)
>Iteration   6: -logL=32631.334, Lambda=1.0e-02, delta=0.922, max(|grad|)=-21.387280 [Index:2]
 Iteration   7: -logL=32632.109, Lambda=1.0e-03, delta=-0.775, max(|grad|)=-60.208218 [Index:7] (stalled)
>Iteration   8: -logL=32631.328, Lambda=1.0e-02, delta=0.781, max(|grad|)=-21.375283 [Index:2]
 Iteration   9: -logL=32632.089, Lambda=1.0e-03, delta=-0.761, max(|grad|)=-59.682426 [Index:7] (stalled)
>Iteration  10: -logL=32631.328, Lambda=1.0e-02, delta=0.761, max(|grad|)=-21.378861 [Index:2]
 Iteration  11: -logL=32632.086, Lambda=1.0e-03, delta=-0.759, max(|grad|)=-59.619641 [Index:7] (stalled)
>Iteration  12: -logL=32631.328, Lambda=1.0e-02, delta=0.759, max(|grad|)=-21.379756 [Index:2]
 Iteration  13: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.613251 [Index:7] (stalled)
>Iteration  14: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379921 [Index:2]
 Iteration  15: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612789 [Index:7] (stalled)
>Iteration  16: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379945 [Index:2]
 Iteration  17: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612786 [Index:7] (stalled)
>Iteration  18: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  19: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  20: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379948 [Index:2]
 Iteration  21: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612794 [Index:7] (stalled)
>Iteration  22: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  23: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  24: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  25: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  26: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  27: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  28: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  29: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  30: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  31: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  32: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  33: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  34: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  35: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  36: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  37: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  38: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  39: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612798 [Index:7] (stalled)
>Iteration  40: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  41: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  42: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
```

```
 Iteration  43: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  44: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  45: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  46: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  47: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  48: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  49: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  50: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  51: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  52: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  53: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  54: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  55: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  56: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  57: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612798 [Index:7] (stalled)
>Iteration  58: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  59: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  60: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  61: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  62: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379948 [Index:2]
 Iteration  63: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612794 [Index:7] (stalled)
>Iteration  64: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  65: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  66: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  67: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  68: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  69: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  70: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  71: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  72: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  73: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  74: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  75: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  76: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  77: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  78: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  79: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  80: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  81: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  82: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  83: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  84: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379948 [Index:2]
 Iteration  85: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  86: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  87: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  88: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379948 [Index:2]
 Iteration  89: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  90: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  91: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration  92: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  93: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612797 [Index:7] (stalled)
>Iteration  94: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379948 [Index:2]
 Iteration  95: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612794 [Index:7] (stalled)
>Iteration  96: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  97: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612795 [Index:7] (stalled)
>Iteration  98: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
 Iteration  99: -logL=32632.086, Lambda=1.0e-03, delta=-0.758, max(|grad|)=-59.612796 [Index:7] (stalled)
>Iteration 100: -logL=32631.328, Lambda=1.0e-02, delta=0.758, max(|grad|)=-21.379949 [Index:2]
=== GOptimizerLM ===
 Optimized function value ..: 32631.328
 Absolute precision ........: 0.005
 Acceptable value decrease .: 1
 Optimization status .......: converged
 Number of parameters ......: 10
 Number of free parameters .: 5
 Number of iterations ......: 101
 Lambda ....................: 0.001
=== GModels ===
 Number of models ..........: 2
 Number of parameters ......: 10
=== GModelSky ===
 Name ......................: Crab radio template
 Instruments ...............: all
 Instrument scale factors ..: unity
 Observation identifiers ...: all
```

Model type ....: DiffuseSource
Model components ....: "SpatialMap" * "PowerLaw" * "Constant"
Number of parameters ......: 5
Number of spatial par's ...: 1
 Prefactor ....: 1 [0.001,1000] (fixed,scale=1,gradient)
Number of spectral par's ..: 3
 Prefactor ....: 5.62221e-16 +/- 1.83243e-16 [1e-23,1e-13] ph/cm2/s/MeV (free,scale=1e-16,gradient)
 Index ....: -2.42874 +/- 0.197857 [-0,-5] (free,scale=-1,gradient)
 PivotEnergy ....: 300000 [10000,1e+09] MeV (fixed,scale=1e+06,gradient)
Number of temporal par's ..: 1
 Constant ....: 1 (relative value) (fixed,scale=1,gradient)
=== GCTAModelRadialAcceptance ===
 Name ....: Background
 Instruments ....: CTA
 Instrument scale factors ..: unity
 Observation identifiers ...: all
 Model type ....: "Gaussian" * "PowerLaw" * "Constant"
 Number of parameters ......: 5
 Number of radial par's ....: 1
 Sigma ....: 3.03361 +/- 0.0409843 [0.01,10] deg2 (free,scale=1,gradient)
 Number of spectral par's ..: 3
 Prefactor ....: 6.45749e-05 +/- 3.81354e-06 [0,0.001] ph/cm2/s/MeV (free,scale=1e-06,gradient)
 Index ....: -1.83205 +/- 0.0330376 [-0,-5] (free,scale=-1,gradient)
 PivotEnergy ....: 1e+06 [10000,1e+09] MeV (fixed,scale=1e+06,gradient)
 Number of temporal par's ..: 1
 Constant ....: 1 (relative value) (fixed,scale=1,gradient)
 Elapsed time ....: 5.560 sec

#7 - 11/03/2014 09:43 AM - Knödlseder Jürgen

From iteration 13 on the algorithm enters obviously an oscillation between improvement and worsening of the logL without any net improvement in the long run.

#8 - 11/03/2014 09:14 PM - Knödlseder Jürgen

- % Done changed from 50 to 100

To prevent the problem, I limited the maximum number of allowed function decreases to 3. Hope that this is well adapted to real life situations.

#9 - 06/29/2015 11:16 PM - Knödlseder Jürgen

- Status changed from Feedback to Closed

So far no problems encountered anymore, close now.

#9 - 06/29/2015 11:16 PM - Knödlseder Jürgen

- Status changed from Feedback to Closed

So far no problems encountered anymore, close now.