

LLP user guide

PsN 4.6.0

Revised 2015-04-15

1 Introduction

The Log Likelihood Profiling (LLP) tool is used to calculate confidence intervals of parameter values. Without the LLP the confidence intervals can be calculated with the standard errors of the parameters under the assumption that the parameter values are normally distributed. The LLP, however, makes no assumption of the shape of the distribution. The LLP tool will calculate the confidence intervals for any number of parameters in the model, working with one parameter at a time. By first fitting the original model and then fixing the parameter at values close to the NONMEM estimate, the LLP obtains the difference in likelihood between the original model and new, reduced model. The logarithm of the difference in likelihood is chi2 distributed and when that value is 3.84, the parameter value is at the 95% confidence limit. The search for the limit is done on both sides of the original parameter value, and thus the LLP makes no assumption of symmetry or the parameter value distribution. Examples

```
llp run89.mod -thetas=1,2 -rse_thetas=20,30
```

2 Input and options

2.1 Required input

A model file is required on the command-line. Then, at least one of the options -thetas, -omegas or -sigmas must be specified, see below. If an lst-

file with standard error estimates already exists, no more input is needed. Otherwise, for each specified \langle parameter \rangle (theta/omega/sigma) there must be a corresponding rse-value given by option `-rse_` \langle parameter \rangle , see below.

-thetas = *theta number list*

A comma-separated list, specifying the thetas for which the llp should try to assess confidence intervals.

-omegas = *omega number list*

A comma-separated list, specifying the omegas for which the llp should try to assess confidence intervals. The numbers refer to the order number of the initial values in the model file. For example, if first there is a block record with size 2 (3 initial values) and then there is a diagonal record with size 3, then numbers 4,5 and 6 refer to the diagonal elements.

-sigmas = *sigma number list*

A comma-separated list, specifying the sigmas for which the llp should try to assess confidence intervals. The numbers refer to the order number of the initial values in the model file. For example, if first there is a block record with size 2 (3 initial values) and then there is a diagonal record with size 3, then numbers 4,5 and 6 refer to the diagonal elements.

2.2 Optional input

-rse_thetas = *list*

A comma-separated list of the relative standard error, specified in percent (%), for each theta listed by option `-thetas`.

-rse_omegas = *list*

A comma-separated list of the relative standard error, specified in percent (%), for each omega listed by option `-omegas`.

-rse_sigmas = *list*

A comma-separated list of the relative standard error, specified in percent (%), for each sigma listed by option `-sigmas`.

-max_iterations = *N*

Default value is 10. This number limits the number of search

iterations for each interval limit. If the llp has not found the upper limit for a parameter after `max_iteration` number of guesses it terminates.

-normq = X

Default value 1.96. The value is used for calculating the first guess of the confidence interval limits. If the standard errors (SE) exist, the guess will be maximum-likelihood estimate \pm normq * SE, otherwise it will be MLE \pm normq * rse_parameter/100 * MLE, where rse_parameter is rse_thetas, rse_omegas or rse_sigmas (optional input parameters). The default value of normq is 1.96 which translates to a 95% confidence interval assuming normal distribution of the parameter estimates.

-outputfile = *filename*

The name of the NONMEM output file. The default name is the name of the model file with '.mod' substituted with '.lst'. Example: if the model file is run89.mod, LLP will by default look for the output file run89.lst. If the name of the output file does not follow this standard, the name must be specified with this option.

-ofv_increase = X

Default value 3.84. The increase in objective function value associated with the desired confidence interval.

-significant_digits = N

Default 3. Specifies the number of significant digits that is required for the test of the increase in objective function value. The default is 3, which means that the method will stop once the difference in objective function value is between 3.835 and 3.845 if -ofv_increase is set to 3.84 (default).

2.3 Some important common PsN options

For a complete list see `common_options.pdf`, or `psn_options -h` on the commandline.

-h or -?

Print the list of available options and exit.

-help

With `-help` all programs will print a longer help message. If an option name is given as argument, help will be printed for this option. If no option is specified, help text for all options will be printed.

-directory = *'string'*

Default `llp_dirN`, where `N` will start at 1 and be increased by one each time you run the script. The `directory` option sets the directory in which `PsN` will run `NONMEM` and where `PsN`-generated output files will be stored. You do not have to create the directory, it will be done for you. If you set `-directory` to a the name of a directory that already exists, `PsN` will run in the existing directory.

-seed = *'string'*

You can set your own random seed to make `PsN` runs reproducible. The random seed is a string, so both `-seed=12345` and `-seed=JustinBieber` are valid. It is important to know that because of the way the Perl pseudo-random number generator works, for two similar string seeds the random sequences may be identical. This is the case e.g. with the two different seeds 123 and 122. Setting the same seed guarantees the same sequence, but setting two slightly different seeds does not guarantee two different random sequences, that must be verified.

-clean = *'integer'*

Default 1. The `clean` option can take four different values:

0 Nothing is removed

1 `NONMEM` binary and intermediate files except `INTER` are removed, and files specified with option `-extra_files`.

2 model and output files generated by `PsN` restarts are removed, and data files in the `NM_run` directory, and (if option `-nmqual` is used) the xml-formatted `NONMEM` output.

- 3 All `NM_run` directories are completely removed. If the PsN tool has created `modelfit_dir:s` inside the main run directory, these will also be removed.

-nm_version = *'string'*

Default is 'default'. If you have more than one NONMEM version installed you can use option `-nm_version` to choose which one to use, as long as it is defined in the `[nm_versions]` section in `psn.conf`, see `psn_configuration.pdf` for details. You can check which versions are defined, without opening `psn.conf`, using the command

```
psn -nm_versions
```

-threads = *'integer'*

Default 5 (if default PsN config file is used). Use the `threads` option to enable parallel execution of multiple models. This option decides how many models PsN will run at the same time, and it is completely independent of whether the individual models are run with serial NONMEM or parallel NONMEM. If you want to run a single model in parallel you must use options `-parafile` and `-nodes`. On a desktop computer it is recommended to not set `-threads` higher the number of CPUs in the system plus one. You can specify more threads, but it will probably not increase the performance. If you are running on a computer cluster, you should consult your system administrator to find out how many threads you can specify.

-version

Prints the PsN version number of the tool, and then exit.

2.4 Auto-generated R-plots from PsN

PsN can automatically generate R plots to visualize results for `llp`, using a default template found in the `R-scripts` subdirectory of the installation directory. The user can also create a custom template, see more details in section “Auto-generated R-plots from PsN” in `common_options.pdf`.

The default llp template requires the R libraries ggplot2, reshape and plyr. If the packages are not installed then no pdf will be generated, see the .Rout file in the main run directory for error messages.

-rplots = *level*

-rplots<0 means R script is not generated

-rplots=0 (default) means R script is generated but not run

-rplots=1 means basic plots are generated

-rplots=2 means basic and extended plots are generated

Troubleshooting

If no pdf was generated even if a template file is available and the appropriate options were set, check the .Rout-file in the main run directory for error messages. If no .Rout-file exists then check that R is properly installed, and that either command 'R' is available or that R is configured in psn.conf.

Basic plots

A basic llp rplot will be generated in file PsN_llp_plots.pdf if option -rplots is set >0, and the general rplots conditions fulfilled, see above. The plot has one panel per evaluated parameter. The parameter value is on the x-axis, and the delta-ofv relative the input model on the y-axis. The blue line shows the confidence interval computed from the standard error under the assumption of normality, and the red line shows the llp-computed confidence interval. Each black dot represents an evaluated parameter value. The IR-value at the top is the ratio between the center-to-upper-limit-distance and the center-to-lower-limit-distance. If the IR is equal to 1 it means the llp-computed confidence interval is symmetric.

3 Output

The file llp_results.csv contains statistics and summaries specific for the llp. The raw_results.csv file is a standard PsN file containing raw result data for termination status, parameter estimates, uncertainty estimates etc. for all model estimations. If option -rplots>0 the file PsN_llp_plots.pdf is also created.