Constellation School on X-rays from Star Forming Regions (May 18-22, 2009)

SPECTRAL FITTING WITH XSPEC

Edited by Costanza Argiroffi and Antonio Maggio

In this example it is shown how to fit a spectrum with XSPEC. The files needed are: the spectrum (already rebinned), the effective area (or auxiliary response), the response matrix, and the background spectrum.

Basic commands

To start an XSPEC session type `xspec` on the terminal. Once entered in XSPEC to obtain the help on any command, type:

`XSPEC> help` *command*

where *command* is the name of the relevant command. Omitting the command name, the `help` command displays the entire manual. To logout from XSPEC type `exit`.

How to load a dataset

To load the spectrum contained in the file *spectrum.pha* (for XMM, or *spectrum.pi* for Chandra) type:

`XSPEC> data` *spectrum.pha*

The `data` command tells the program to read the data. Background, effective area, and response matrix are automatically loaded if they are named in the header of the data file. You can check what files were loaded with the command

`XSPEC> show file`

If you need to re-load any of them, use the following commands:

`XSPEC> backgrnd` *bgspectrum.pha*

`XSPEC> arf` *effarea.arf*

`XSPEC> response` *respmatrix.rmf*

where *bgspectrum.pha* (or *bgspectrum.pi* for Chandra), *effarea.arf*, and *respmatrix.rmf* are supposed to be the file names. If the filename input of the `data` command is `none` the current spectrum is removed from XSPEC.

Finally, the command:

`XSPEC> show rate`

displays information on total counts, exposure time and background-subtracted count rates for the current dataset.

Plot

One of the first things most users will want to do at this stage, even before fitting models, is to look at their data. The plotting device should be set first, with the command `cpd` (change plotting device). Here, we use the pgplot X-Window server, /xw.

```
XSPEC> cpd /xw
```

There are more than 50 different things that can be plotted, all related in some way to the data, the model, the fit and the instrument. To see them, type:

```
XSPEC> plot ?
```

The most fundamental is the data plotted against instrument channel (data); next most fundamental, and more informative, is the data plotted against channel energy. To do this plot, use the XSPEC command:

```
XSPEC> setplot energy
```

while use `setplot wave` if you prefer wavelength. The `setplot` command can be used to modify other values used by the plot routines. To check the current plot setting type:

```
XSPEC> show plot
```

The command:

```
XSPEC> plot data      or      XSPEC> plot ldata
```

produces a linear or a log plot. The command:

```
XSPEC> plot data delchi
```

produces a plot with the residuals with respect to the current model (see below for instructions about models). To produce a color postscript file containing the plot type:

```
XSPEC> cpd plotfile.ps/cps
```

```
XSPEC> plot
```

In this case the plot will be saved in the *plotfile.ps* file. Type again `cpd /xw` to return to the pgplot X-Window. To improve the plot appearance it could be useful to discard some data channels. To do that see the next section.

How to discard data channels

To discard some data channels, XSPEC provides the `ignore` command. To use it type:

`XSPEC> ignore` $n_1 - n_2$

where n_1 and n_2, if integers, indicate the channel range to be discarded, while, if n_1 and n_2 are reals, they indicate the energies (or wavelengths if `setplot wave` has been specified). To ignore the first or last channels you can type:

`XSPEC> ignore` $** - n_2$

`XSPEC> ignore` $n_1 - **$

While the command:

```
XSPEC> ignore bad
```

discards channels listed as bad channels in the headers of the spectral files because of known issues. To include again some data channels previously discarded use the `notice` command, which has the same syntax.

How to define a model and its parameters

The `model` command allows you to define the model to be used for the data fitting. The model is selected by combining (adding or multiplying) the model components contained in XSPEC. The XSPEC model components are divided into two categories: *additive model components*, which represent X-ray sources of different kinds (i.e. power laws, thermal models, emission lines, etc.); *multiplicative model components*, which represent phenomena that modify the observed X-ray spectrum (by an energy-dependent multiplicative parameter) produced the additive components (i.e. photoelectric absorption models, edges, absorption lines, etc.).
A list of all the currently installed models is given in response to the command:

```
XSPEC> model ?
```

Model components are combined in the obvious algebraic way, with + separating additive models, * separating multiplicative models, and parentheses to show which additive models the multiplicative models act on. The * need not be included next to parentheses, where it is redundant. As an example, typing:

```
XSPEC> model wabs ( apec + apec )
```

you define a model composed by two emission components of collisionally-ionized plasma (`apec`, whose parameters are the normalization, the temperature, and the plasma metallicity), subjected to the same photo-electric absorption (`wabs`, whose parameter is the hydrogen column). While the command:

```
XSPEC> model wabs * vapec + wabs * vapec
```

defines a model composed by two emission components of collisionally-ionized plasma (`vapec` is a `apec` variant that allows the user to set the abundance of each chemical element), subjected to different photo-electric absorption. The user is then prompted for the initial values of the parameters.
Type:

```
XSPEC> show model
```

to obtain information on the current model. The commands `addcomp` or `delcomp` allows you to add a component to the model or to remove it.

How to fit the spectrum

To decide which parameters have to vary or have to be frozen during the fit use the `thaw` and `freeze` commands. The commands:

`XSPEC> thaw` n_1 n_2 `XSPEC> freeze` n_3 n_4

let the n_1 and n_2 parameters vary, and instead freezes n_3 and n_4. The number associated to each parameter is listed in the first column of the output produced by the `show model` command. To check which are the free parameters type:

```
XSPEC> show free
```

To change the current value of a parameter type:

```
XSPEC> newpar n x
```

where n is the parameter number and x is its new value. The command:

```
XSPEC> newpar n1 = n2
```

link the n_1 parameter to the value of the n_2 one. If you type:

```
XSPEC> newpar n1
```

XSPEC then prompt for the initial value of the n_1 parameter (and other settings on the parameter variation). The command:

```
XSPEC> fit k Δχ²
```

finds the best fit model parameters for the current data by minimizing the current statistic. k (integer) is the maximum number of iterations (then query the user whether he or she wants to continue[1]), and $\Delta\chi^2$ (real) is the change in the fit statistic between iterations less than which the fit is deemed to have converged. Both k and $\Delta\chi^2$ can be omitted in the `fit` command.

To calculate the flux of the current model use the command:

```
XSPEC> flux E1 E2
```

where E_1 and E_2 define the energy range over which the flux is calculated (initial default values are 2 and 10 keV). The flux is given in units of ph $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and erg $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The energy range must be contained within the range covered by the current spectrum (which determines the range over which the model is evaluated). Note that if you need the unabsorbed flux (i.e. to determine the source X-ray luminosity), you should set to zero the model absorption before running the `flux` command.

How to save and reload data and models

To save the model after having performed the fit procedure use the command:

```
XSPEC> save model modelfile
```

This command creates the file *modelfile.xcm* which contains the current model. A saved model can be reloaded with the command:

```
XSPEC> @modelfile
```

You can also save the commands required to load the data (source spectrum, response matrix and effective area files) with

```
XSPEC> save all commandfile
```

The file *commandfile.xcm* can be reloaded and executed as above.

How to determine statistical uncertainties

To determine the confidence region for the model parameters type:

[1]The `query yes` command will suppress all questions and assume that their answer is yes.

`XSPEC> error` $\Delta\chi^2$ n_1 n_2 n_3

where $\Delta\chi^2$ (a real value that can also be omitted) indicates the delta fit statistic and n_1 n_2 n_3 (integer values) are the parameters to be examined. Each indicated parameter is varied until the value of the fit statistic, minimized by allowing all the other non-frozen parameters to vary, is equal to the last value of fit statistic determined by the fit command plus the indicated delta fit statistic. Note that before the error command is executed, the data must be fitted. The initial default values for the delta fit statistic is 2.706 (equivalent to the 90% confidence region for a single interesting parameter).
The procedure above explores the parameter space in the neighborhood of the best-fit values and may find better minima (lower χ^2 values).
You can also compute a grid of χ^2 values in a given region (2-D section) of the parameter space with the command

`XSPEC> steppar` n_1 $lowval_1$ $hival_1$ $nsteps_1$ n_2 $lowval_2$ $hival_2$ $nsteps_2$

where n_1 and n_2 (integer values) are two selected parameters, *lowval* and *hival* define a range of values for each parameter, and *nsteps* is the number of grid points in each range. This command computes the lowest χ^2 by keeping fixed the two selected parameters at each point in the grid and allowing all the other non-frozen parameters to vary until the lowest fit statistic is found. The χ^2 map created in this way can be displayed as a contour plot with the command

```
XSPEC> plot contour
```

By default you will see contour levels drawn at pre-defined values of $\Delta\chi^2$ (2.3, 4.61, and 9.21) with respect to the last best-fit model, corresponding to the 68%, 90%, and 99% statistical confidence regions for the two interesting parameters you have chosen. This map is useful to spot possible correlations among the two parameters, and to find possible local minima of the fit statistics. However a proper choice of the boundary values of the grid and of its resolution (given by the number of steps) is required to obtain a clear χ^2 map.