X-ray Astronomy School

Introduction to X-ray Data Analysis

David Huenemoerder (MIT)

with help from John Davis, John Houck, & Mike Nowak

Analysis:

a. A detailed examination or study of something so as to determine its nature, structure, or essential features.

b. The action or process of carrying out such a detailed examination; the methodical or systematic investigation of something complex in order to explain or understand it.

"*A more intimate analysis*..*matured my conjecture into full conviction*."
(1817 S. T. Coleridge *Biogr*. *Lit*. I. iv. 86)

One essential feature we wish to determine is the source flux:

$$S(E, \hat{p}, t)\,[\mathrm{photons\,cm^{-2}\,s^{-1}\,keV^{-1}\,sr^{-1}}]$$

Or some marginalized form - the spectrum, image, or light curve: $S(E)$, $S(\hat{p})$, or $S(t)$.

However, our source flux can be

- ⭐Confused with other sources
- ⭐Combined with a diffuse background source
- ⭐Absorbed by foreground material

The photons then encounter

- ⭐Optics and detectors with <100% efficiency
- ⭐Optics which redistribute (scatter) photon directions
- ⭐Detectors which distribute photons into *channels*
- ⭐Detectors with limiting count rates, macroscopic pixels
- ⭐Detectors with internal or local background

What we get are a list of counts with discretized times, positions, channels, ...

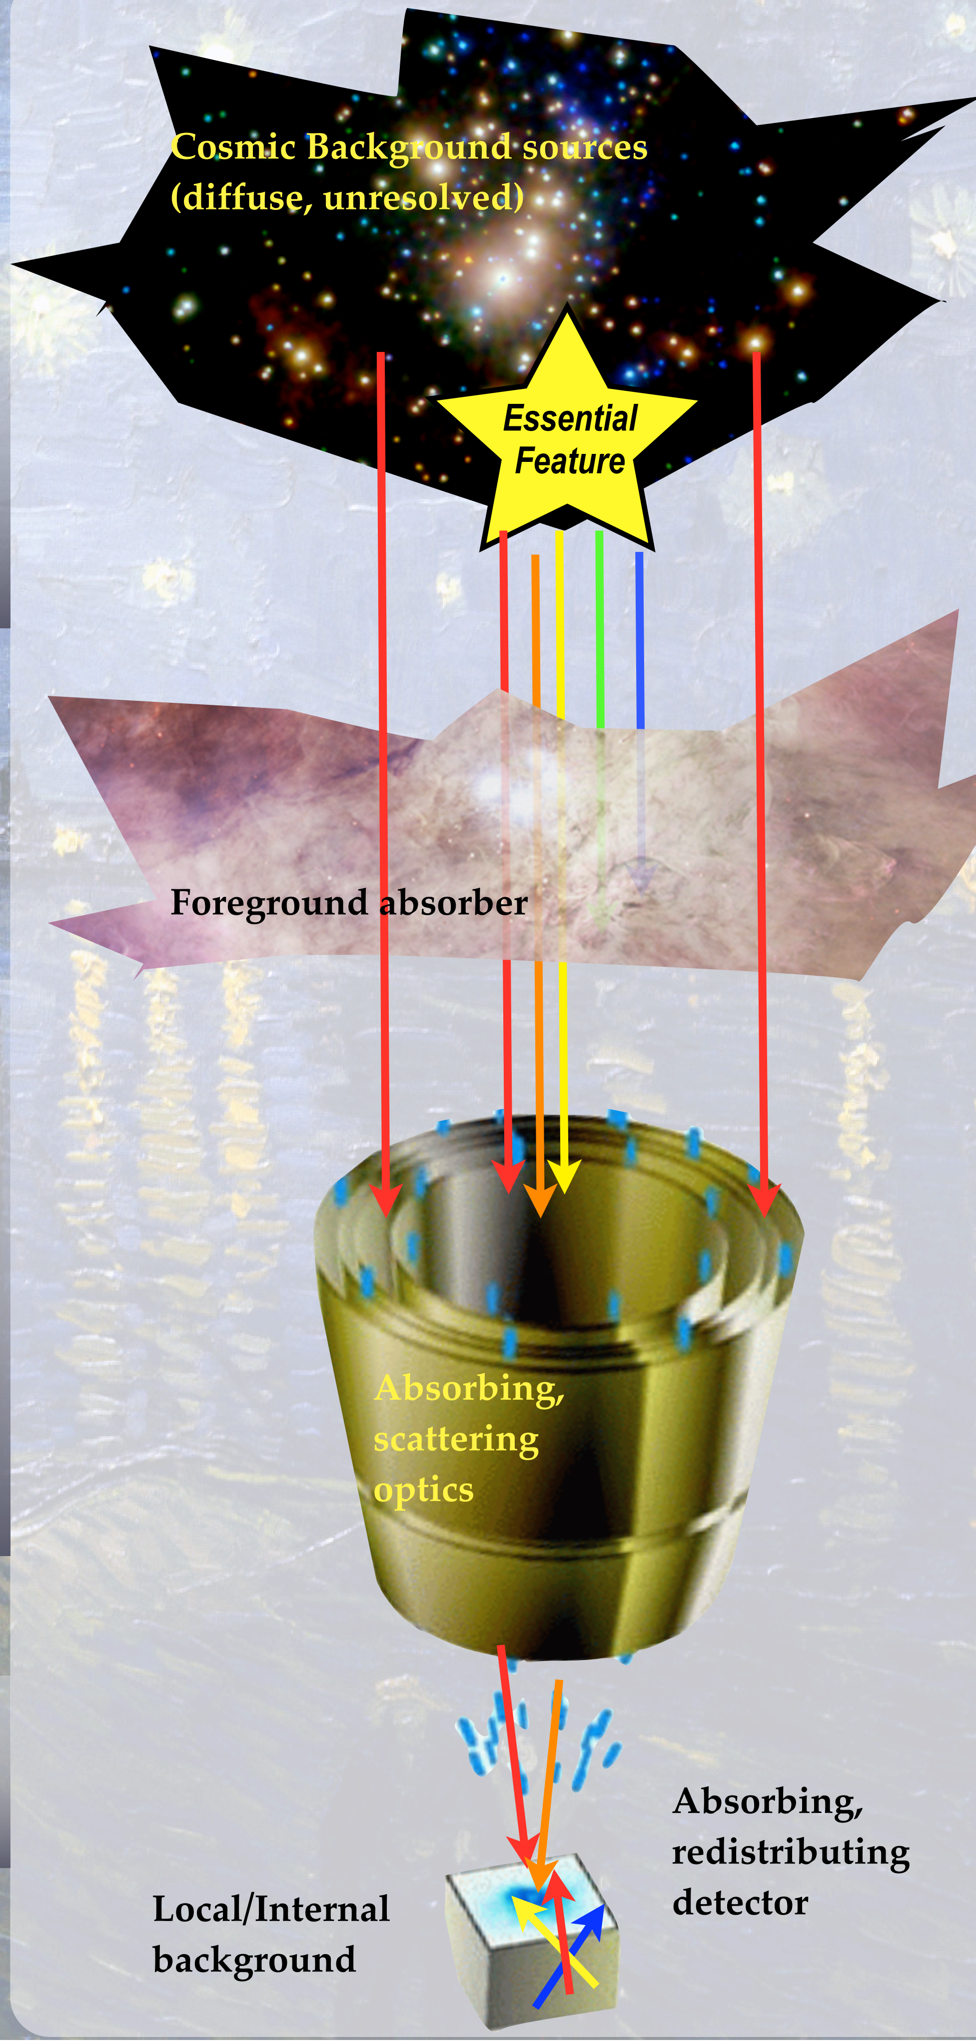

Formal description of the transformation from *photons* to *counts*:

(Note: not all photons are counted; not all counts are from photons.)

$$C(h) = T \int_0^\infty \sum_i R_i(E,h)\, A_i(E)\, S_i(E)\, dE + B(h)$$

$C(h)$	Counts in detector channel *h* (*h* is NOT energy)
T	Exposure time
$R_i(E,h)$	Detector redistribution from energy *E* to channel *h* ("response matrix", or RMF) for the source component *i*.
$A_i(E)$	Effective area (geometric area times efficiency) (Ancillary Response File, or ARF) at energy *E* for source component *i*.
$S_i(E)$	Source photon flux at energy *E* for component *i*.
$B(h)$	"Background" counts in channel *h*. (This could include both cosmic and internal sources, empirical or modeled, possibly from other observations.)

$$C(h) = T \int_0^\infty \sum_i R_i(E,h)\, A_i(E)\, S_i(E)\, dE + B(h)$$

Note: this is the *spectral forward-folding* equation. It does not explicitly include spatial redistribution by the Point Spread Function (*PSF*) or time dependent source flux or time dependent background. (For detailed treatment see J.E. Davis (2001) ApJ, 548, 1010.)

Counterparts* for imaging would be:

C(x,y): counts vs detector (or sky) position bin (an image, or 2D histogram) (over some channel range).
PSF(E, p, p') the angular redistribution from input angle *p* to output angle *p'* at energy *E*.

Or for timing,

C(h, t): a count-rate histogram for channel *h* vs time, *t* (over some angular region).

---

* Details deferred to the advanced analysis lecture, or as an exercise for the reader.

Observational Data Products

X-ray detectors are *photon counting*.
X-ray photons are *sparse.*

Hence, the basic data product from which all analysis follows is the

event list,

provided as a FITS binary table. The event list tabulates attributes of each detected *count,* such as:

- Time
- Detector element ID
- Detector X pixel
- Detector Y pixel
- Pulse height ("PHA")
- Sky x pixel
- Sky y pixel
- ... And much more

You must be able to visualize, browse, and filter event lists.

Working With Event Lists

Viewing header or data (as text):

- `dmlist` (ciao)
- `fdump` (ftools)
- `fitsdump` (marx)
- `fv` (ftools)
- `prism` (ciao)

Plotting/imaging the data:

- `chips` (ciao)
- `ds9` (sao)
- `IDL` (custom)
- `ISIS` (cfitsio, pgplot)
- `prism` (ciao)
- `vwhere` (ISIS module)

Working With Event Lists

Viewing header or data (as text):

- `dmlist` (ciao)

```
$ dmlist secondary/acisf07437_000N002_evt1.fits cols

--------------------------------------------------------------------------------
Columns for Table Block EVENTS
--------------------------------------------------------------------------------

ColNo  Name                 Unit        Type             Range
   1   time                 s            Real8          288836005.5346599817:288996206.8922700286 S/C TT corresponding to mid-exposure
   2   ccd_id                            Int2           0:9                     CCD reporting event
   3   node_id                           Int2           0:3                     CCD serial readout amplifier node
   4   expno                             Int4           0:2147483647            Exposure number of CCD frame containing event
   5   chip(chipx,chipy)    pixel        Int2           1:1024                  Chip coords
   6   tdet(tdetx,tdety)    pixel        Int2           1:8192                  ACIS tiled detector coordinates
   7   det(detx,dety)       pixel        Real4          0.50:     8192.50       ACIS detector coordinates
   8   sky(x,y)             pixel        Real4          0.50:     8192.50       sky coordinates
   9   phas[3,3]            adu          Int2(3x3)      -4096:4095              array of pixel pulse heights
  10   pha                  adu          Int4           0:36855                 total pulse height of event
  11   pha_ro               adu          Int4           0:36855                 total read-out pulse height of event
  12   energy               eV           Real4          0:  1000000.0           nominal energy of event (eV)
  13   pi                   chan         Int4           1:1024                  pulse invariant energy of event
  14   fltgrade                          Int2           0:255                   event grade, flight system
  15   grade                             Int2           0:7                     binned event grade
  16   status[4]                         Bit(4)                                 event status bits
```

- `prism` (ciao)
- `vwhere` (ISIS module)

Working With Event Lists

Event lists come in different *levels* of processing; generally, the higher the level, the more has been done to clean, or add information. E.g., for Chandra:

- Level 0 - telemetry - very raw (don't go there)
- Level 1 - physical coordinates and scalings - unfiltered.
- Level 1.5 - some source-dependent information (for grating spectra)
- Level 2 - filtered of bad pixels, bad time intervals

Observatory centers do a good job of data preparation*, to Level 2 (or equivalent). ***However***, they do not know the specific characteristics of your source of interest, or the details of your analysis; you might require more stringent filters, or could tolerate more relaxed ones. It's up to you to review the data.

---

* "Trust everyone, but cut the cards."

Working With Event Lists: Image Display

(Good, bad, or just ugly?)

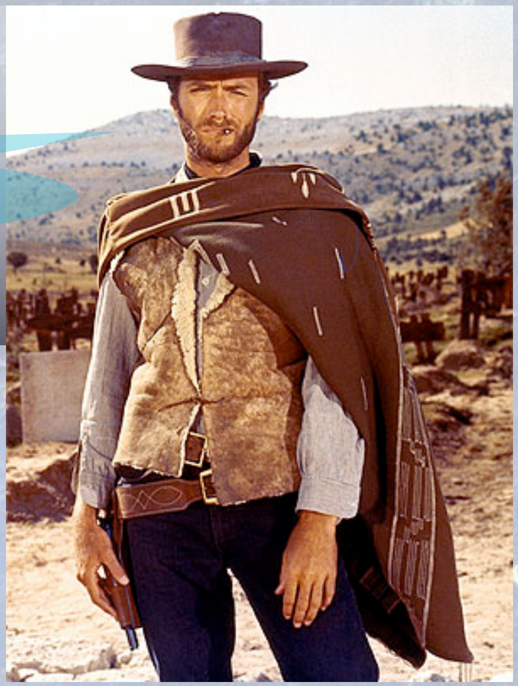

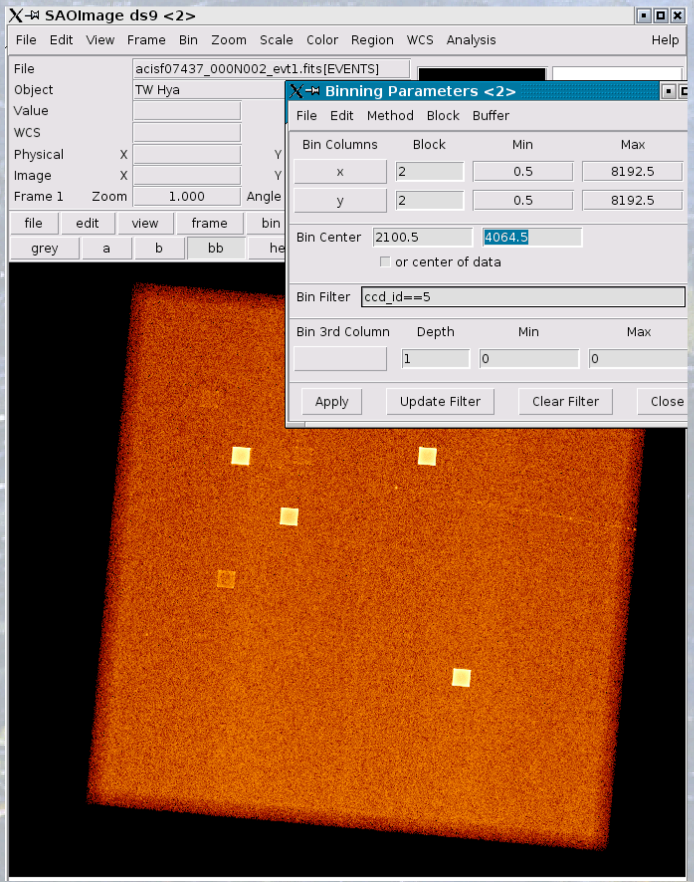

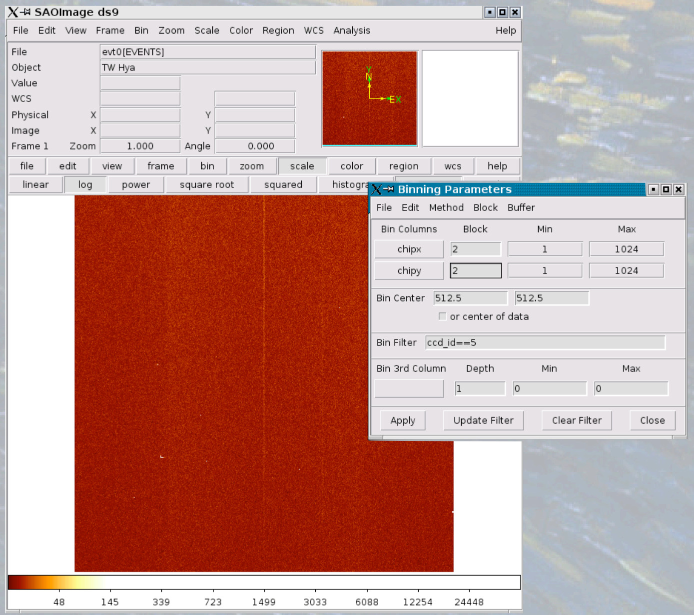

Coordinate systems, and dither.

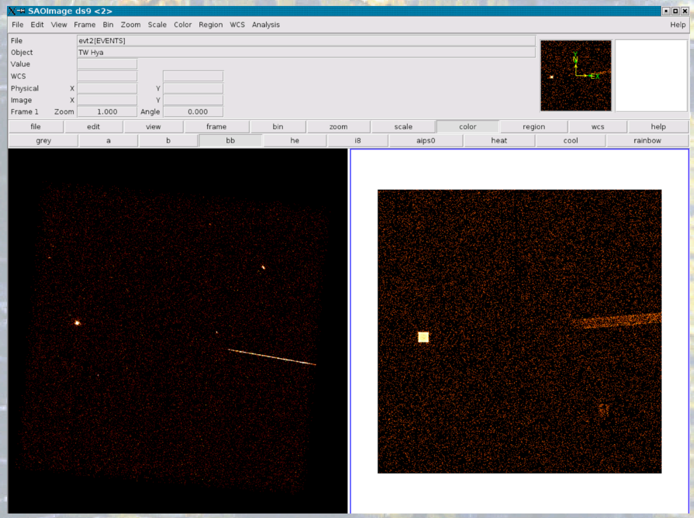

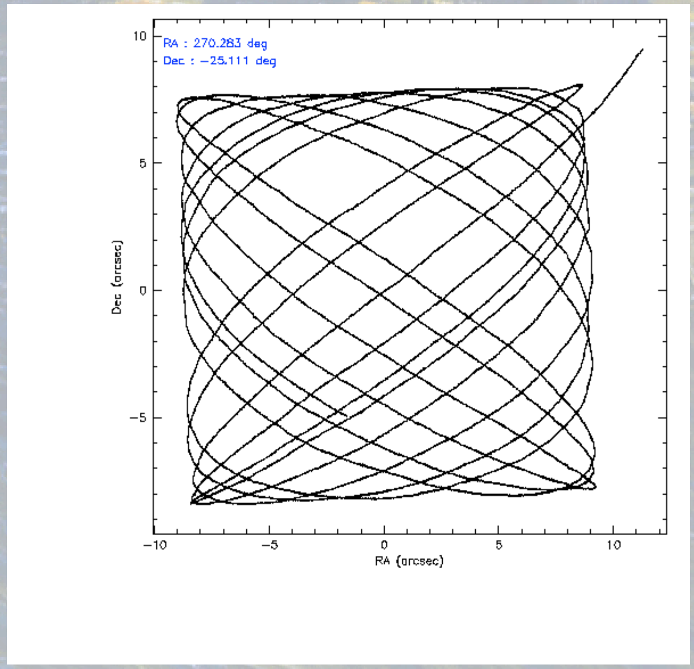

Unfiltered vs filtered

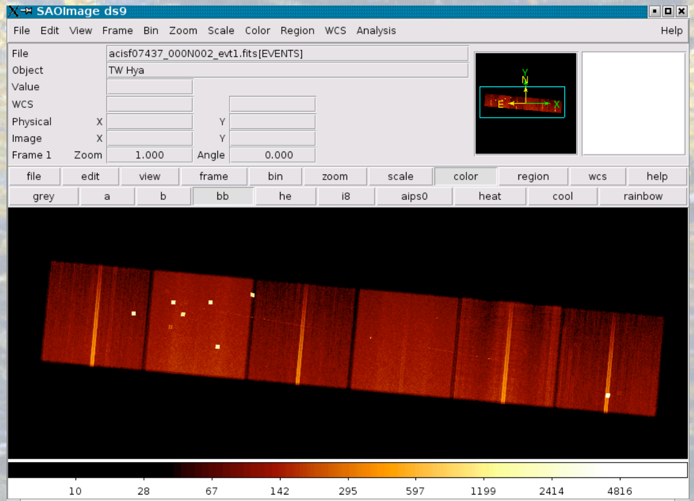

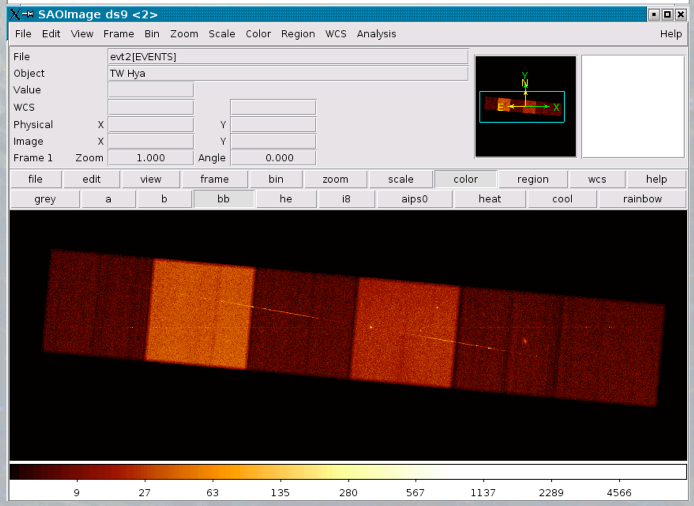

Working With Event Lists: Image Display w/ ds9

Sky coordinates, ACIS-5

Chip coordinates, ACIS-5

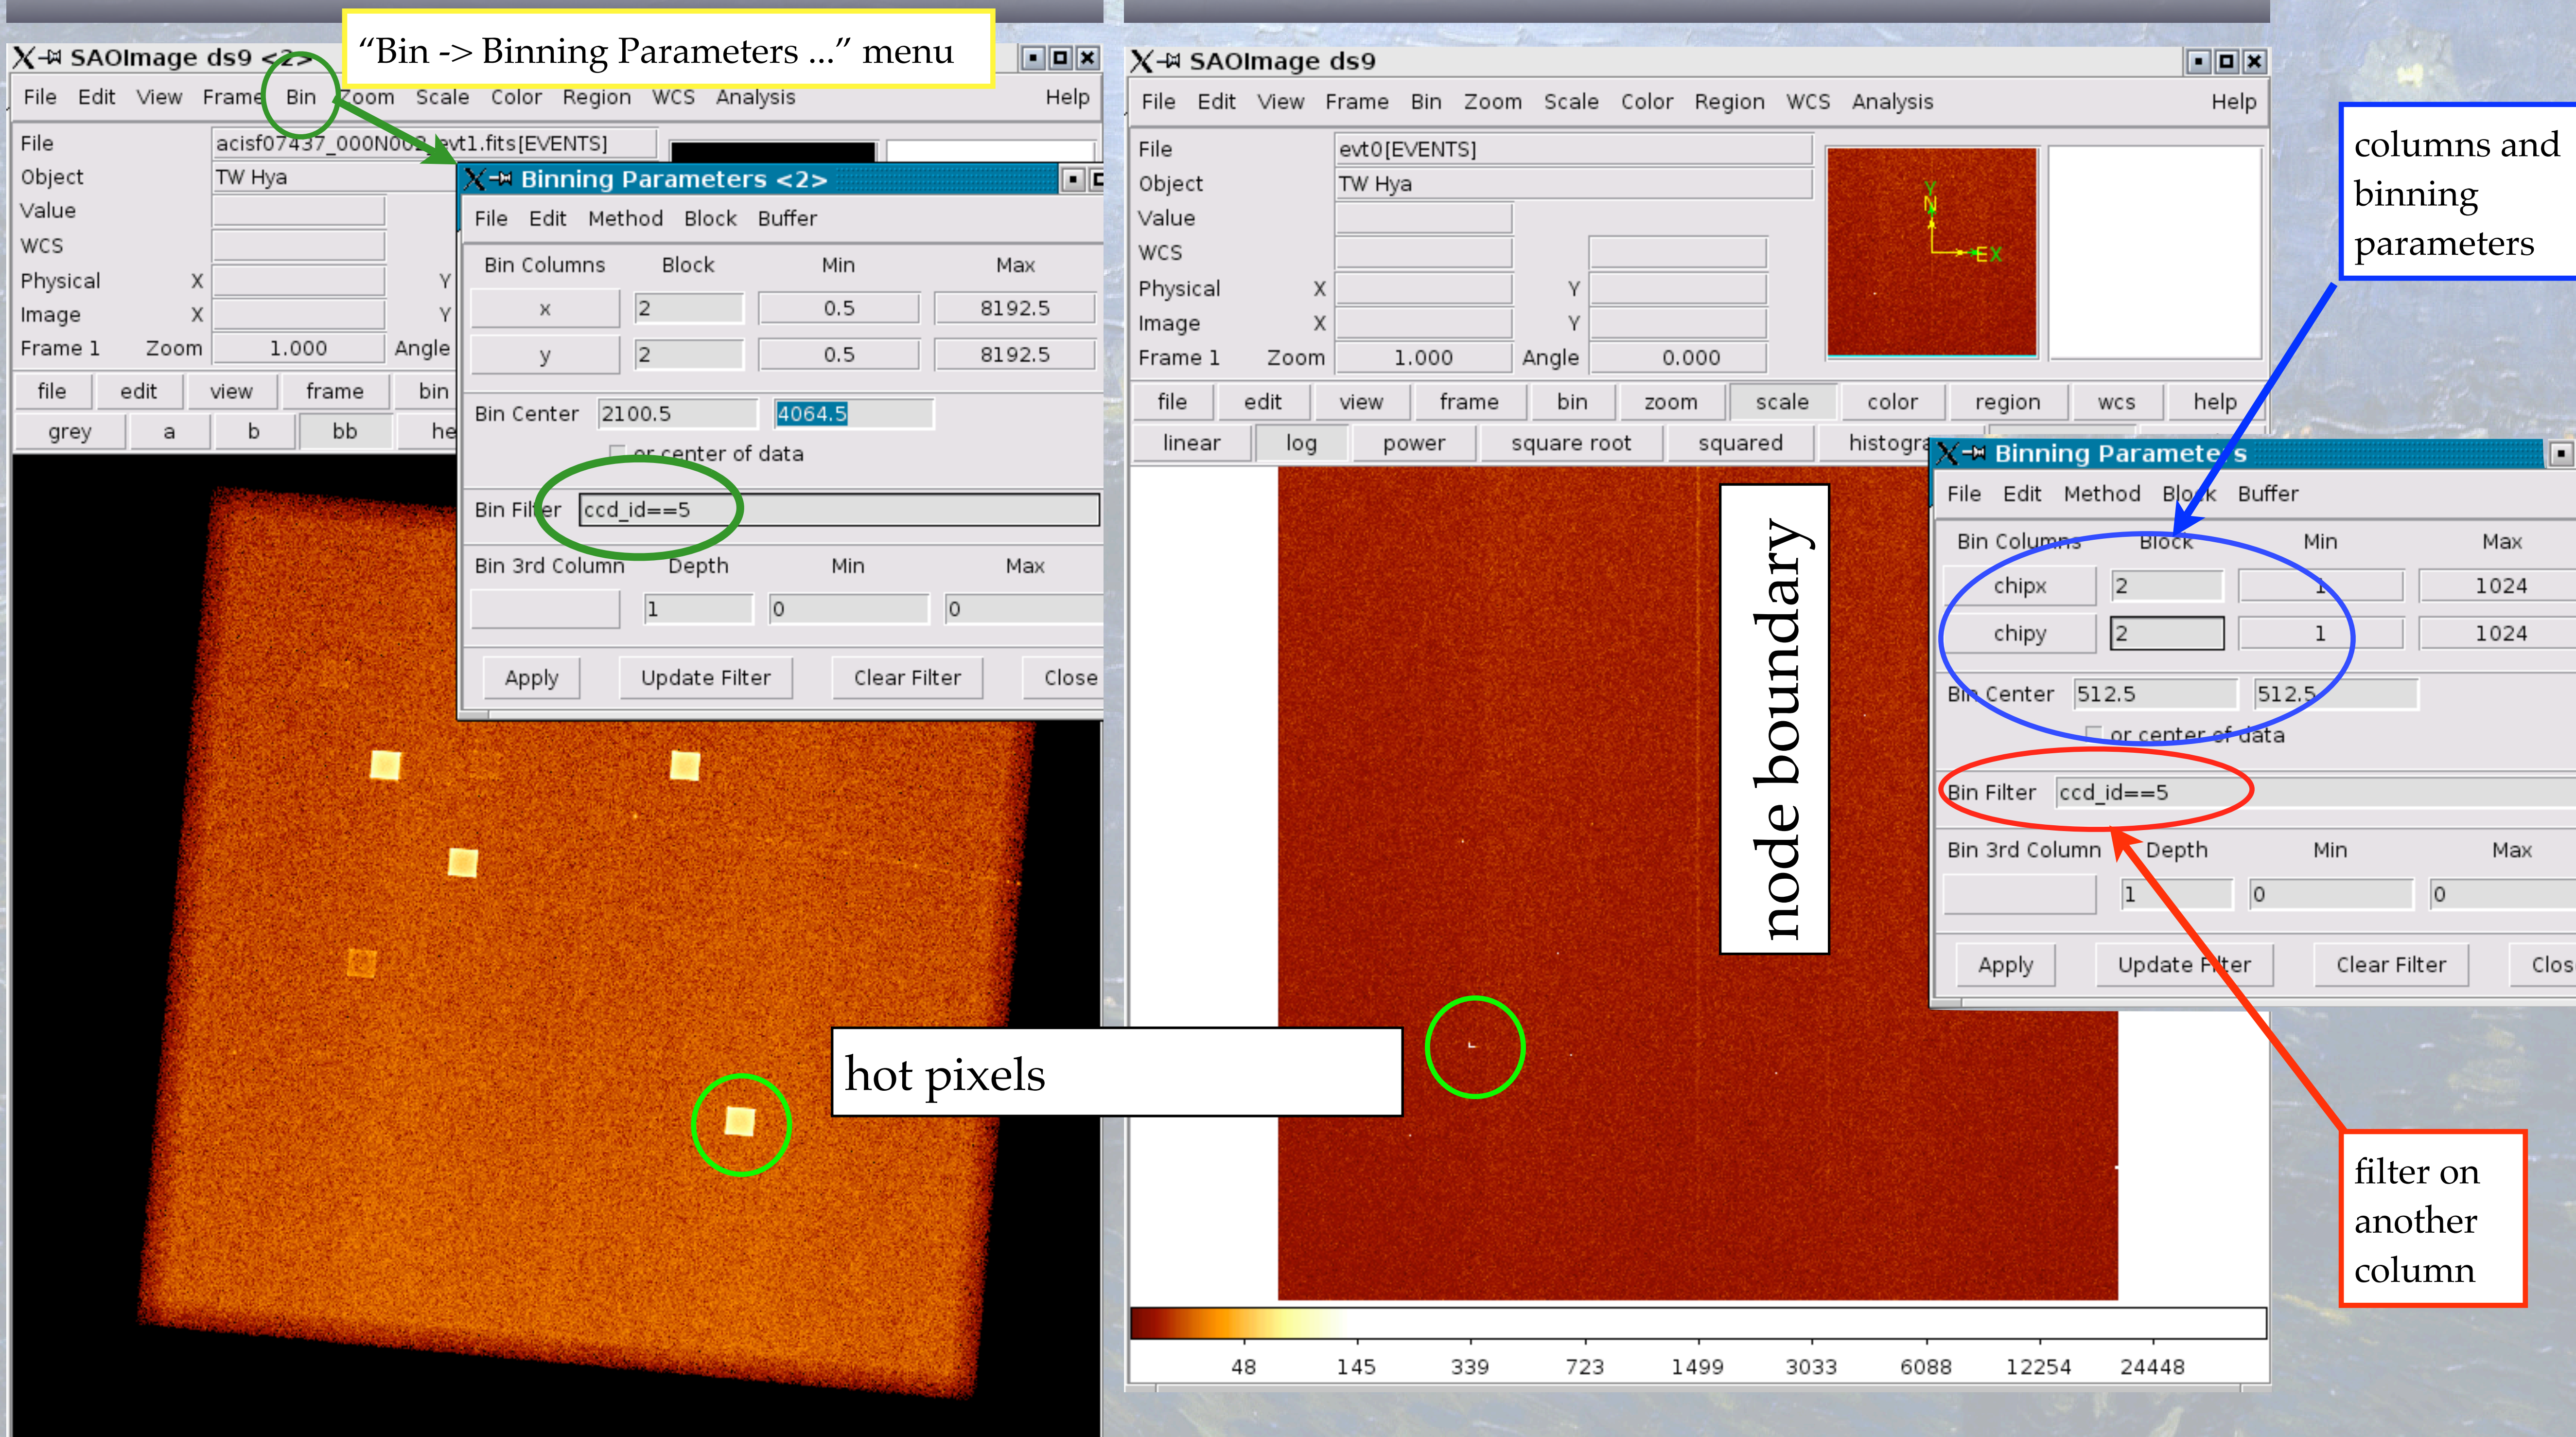

Working With Event Lists: *Chandra* Dithers...

Aspect Offsets

5 ks

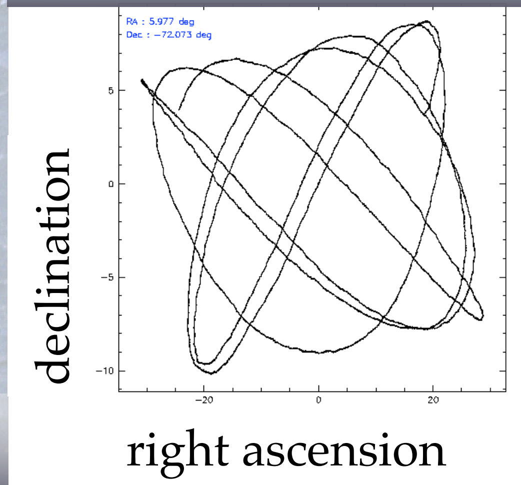

10 ks

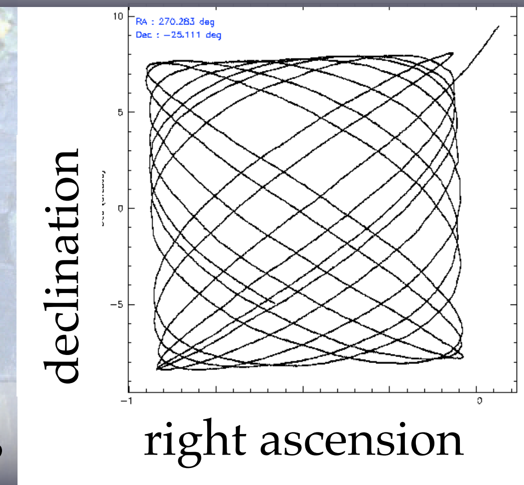

50 ks

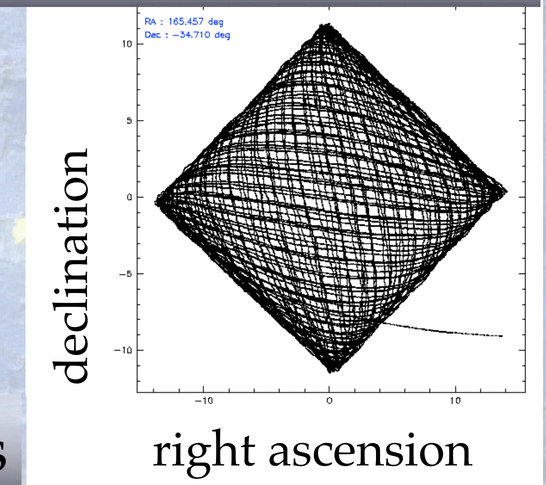

(Plots from your V&V report.)

Dither squares the circle (and vice versa)

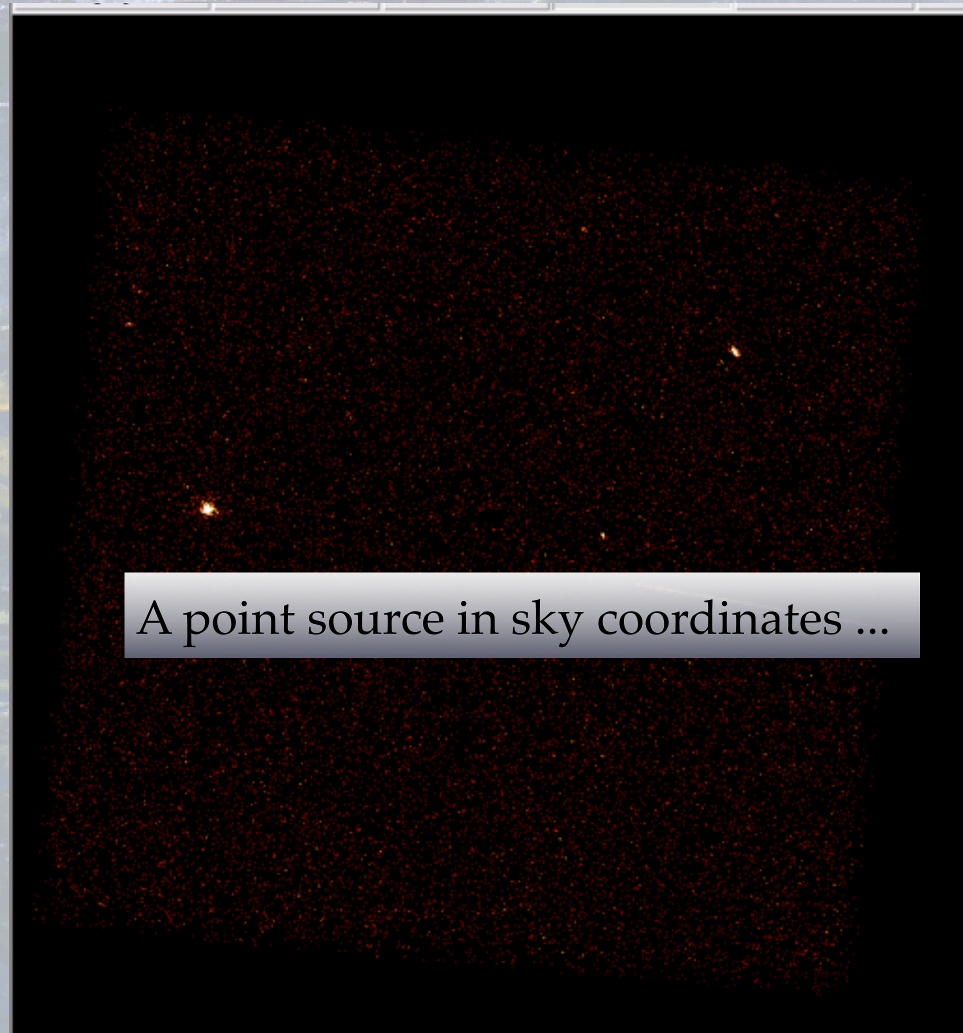

A point source in sky coordinates ...

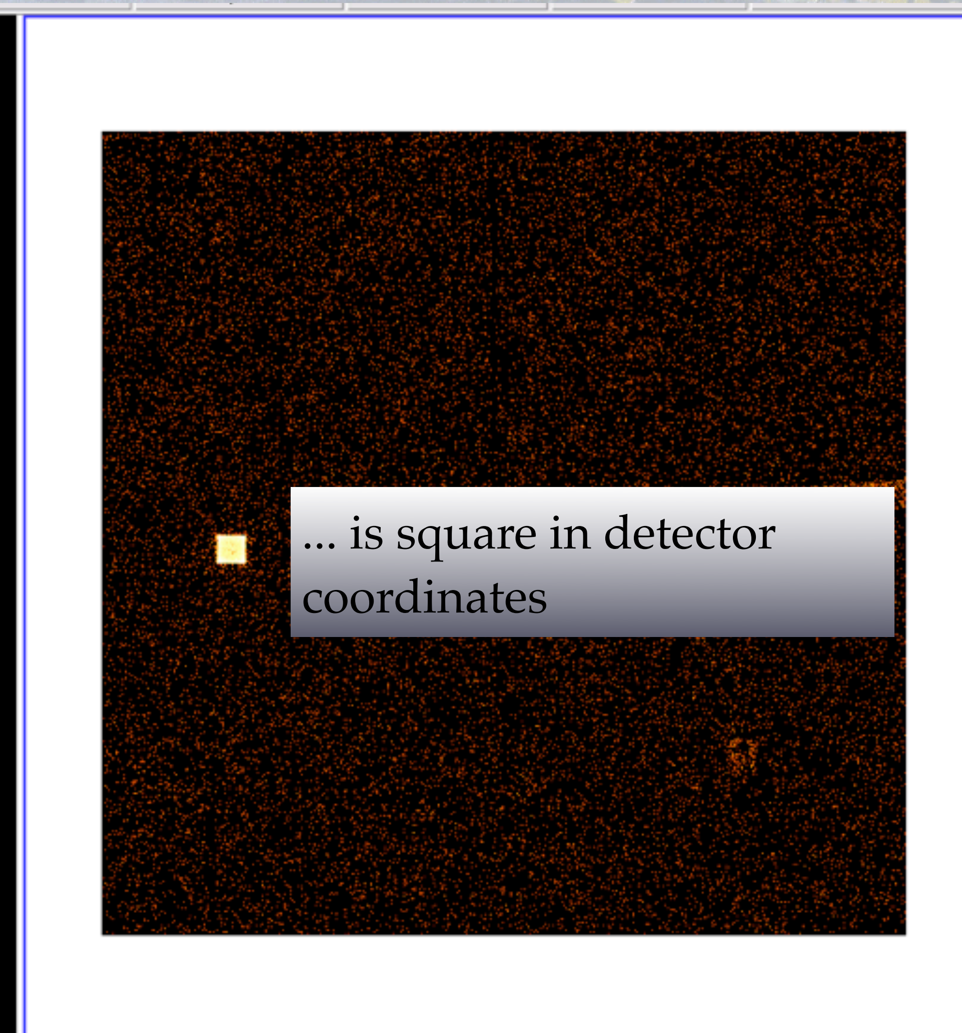

... is square in detector coordinates

11

... and Suzaku Wiggles

Thermal flexing of the Suzaku spacecraft leads to slow wobbling of the optical axis, and hence blurring of the image. Current Suzaku tools partially correct this effect by adjusting the spacecraft attitude file. `aeattcor.sl` further improves this correction for bright sources by using their detected image to create a new attitude file. Applying this attitude file via the FTOOLS `xiscoord` command leads to sharper PSF images, as shown below. (See `space.mit.edu/cxc/software/suzaku/`.)

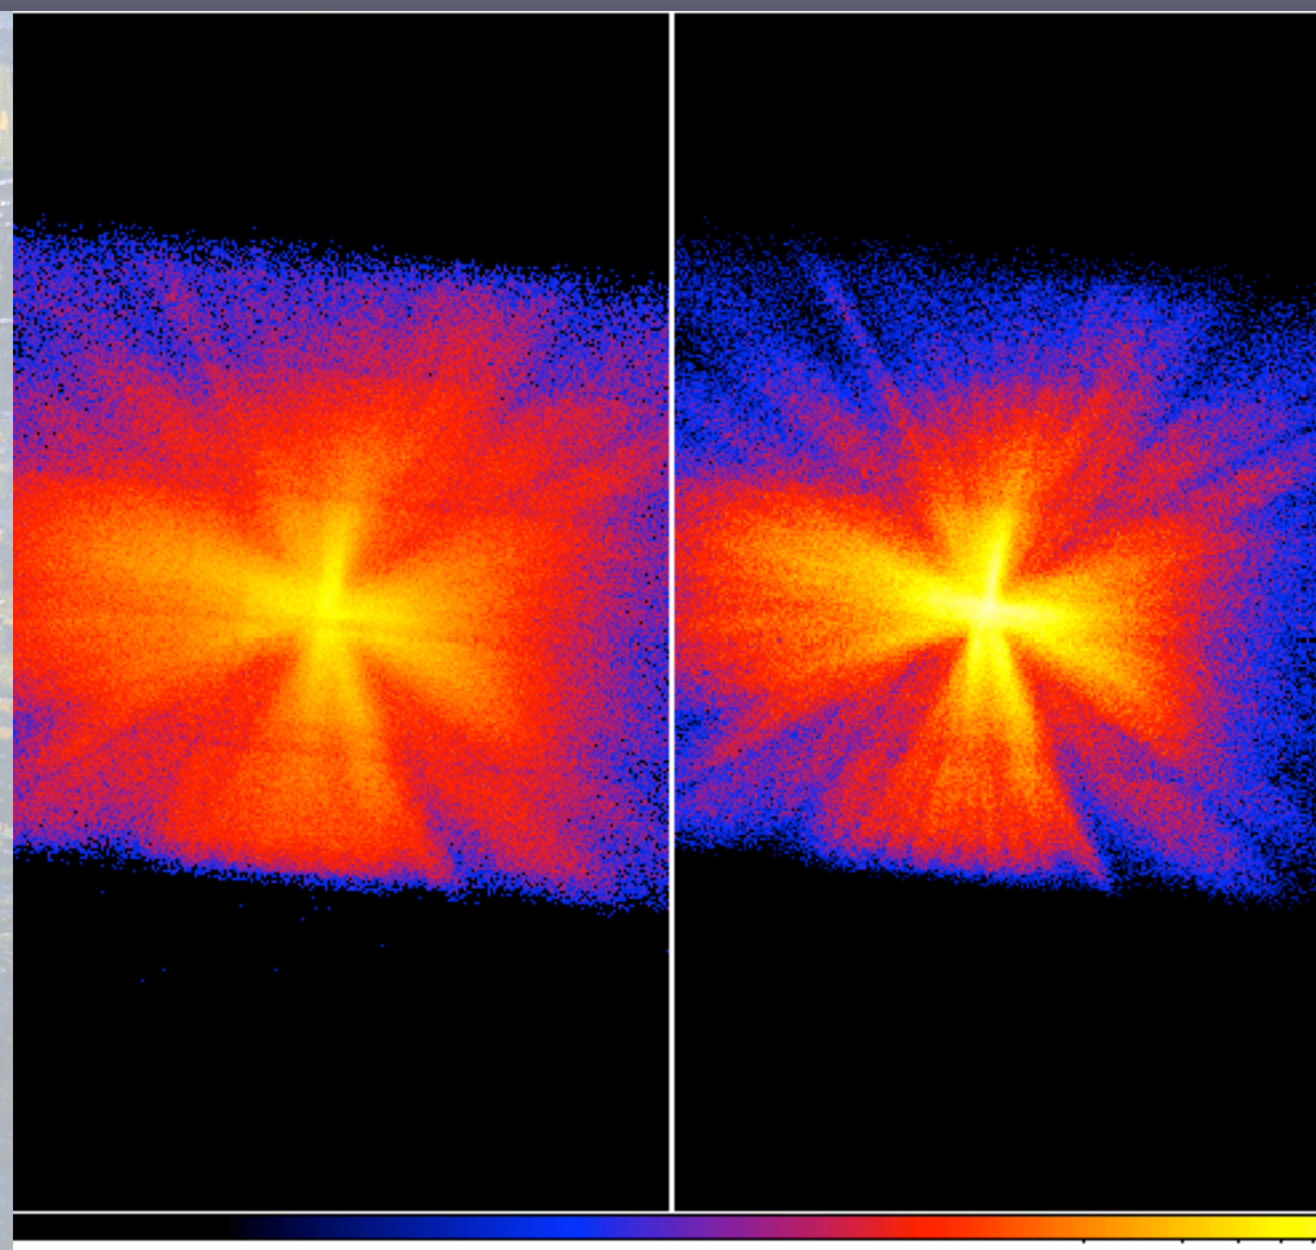

12

Aside: Inspecting Aspect with ds9:

`pcad*asol1.fits` file, extension ASPSOL, columns `ra` and `dec` – use "Binning Parameters" menu and fiddle with the blocking and center to get a view:

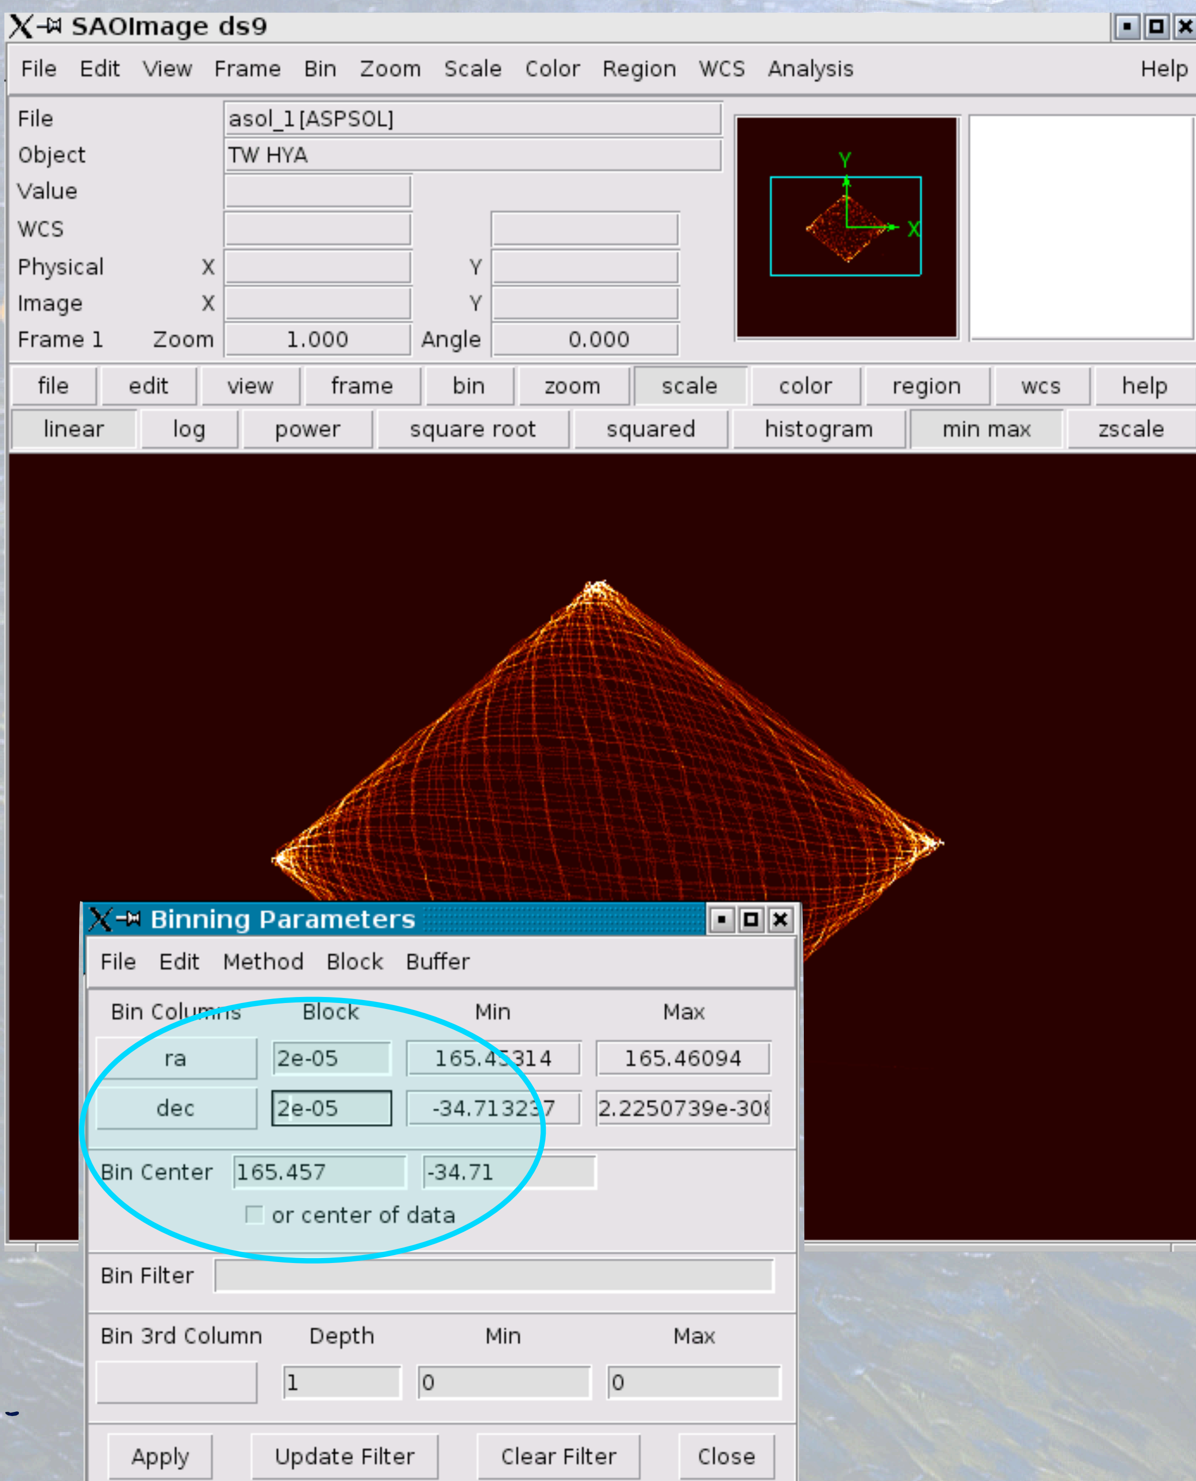

Filtered on `time:` (e.g., first determine time range, as from `dmlist pcad*asol1.fits subspace | grep time)`

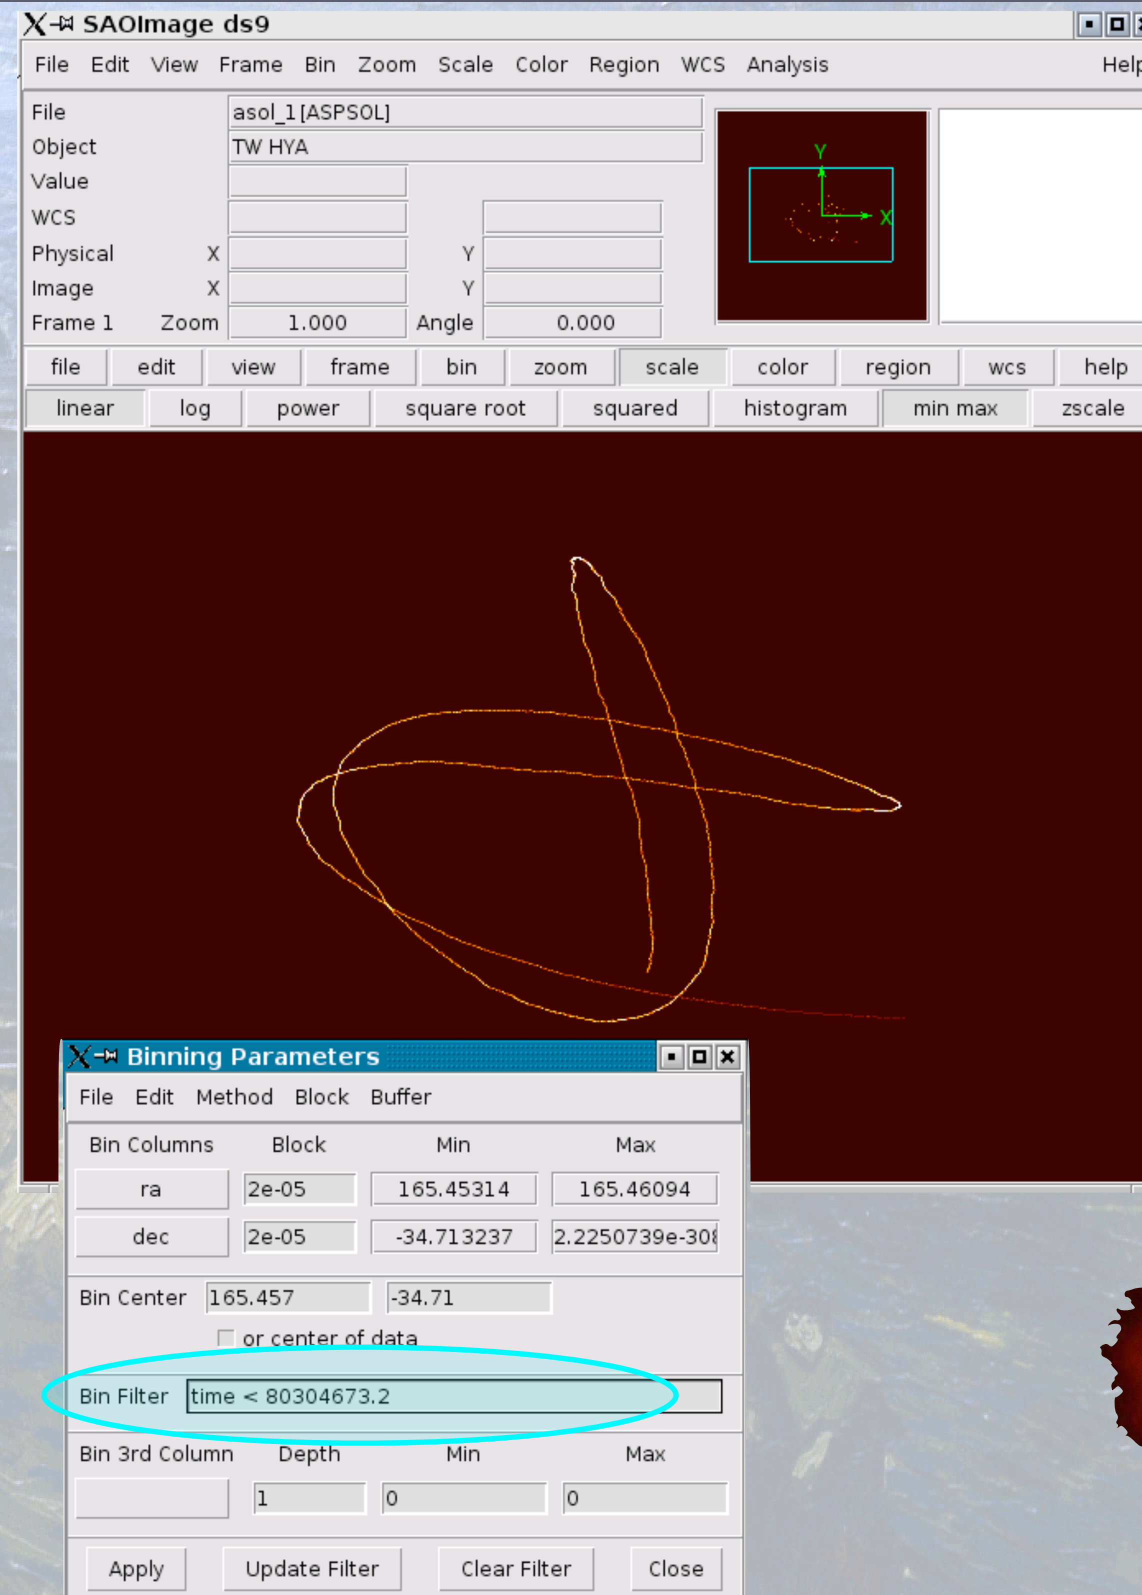

13

Working With Event Lists: Filtering

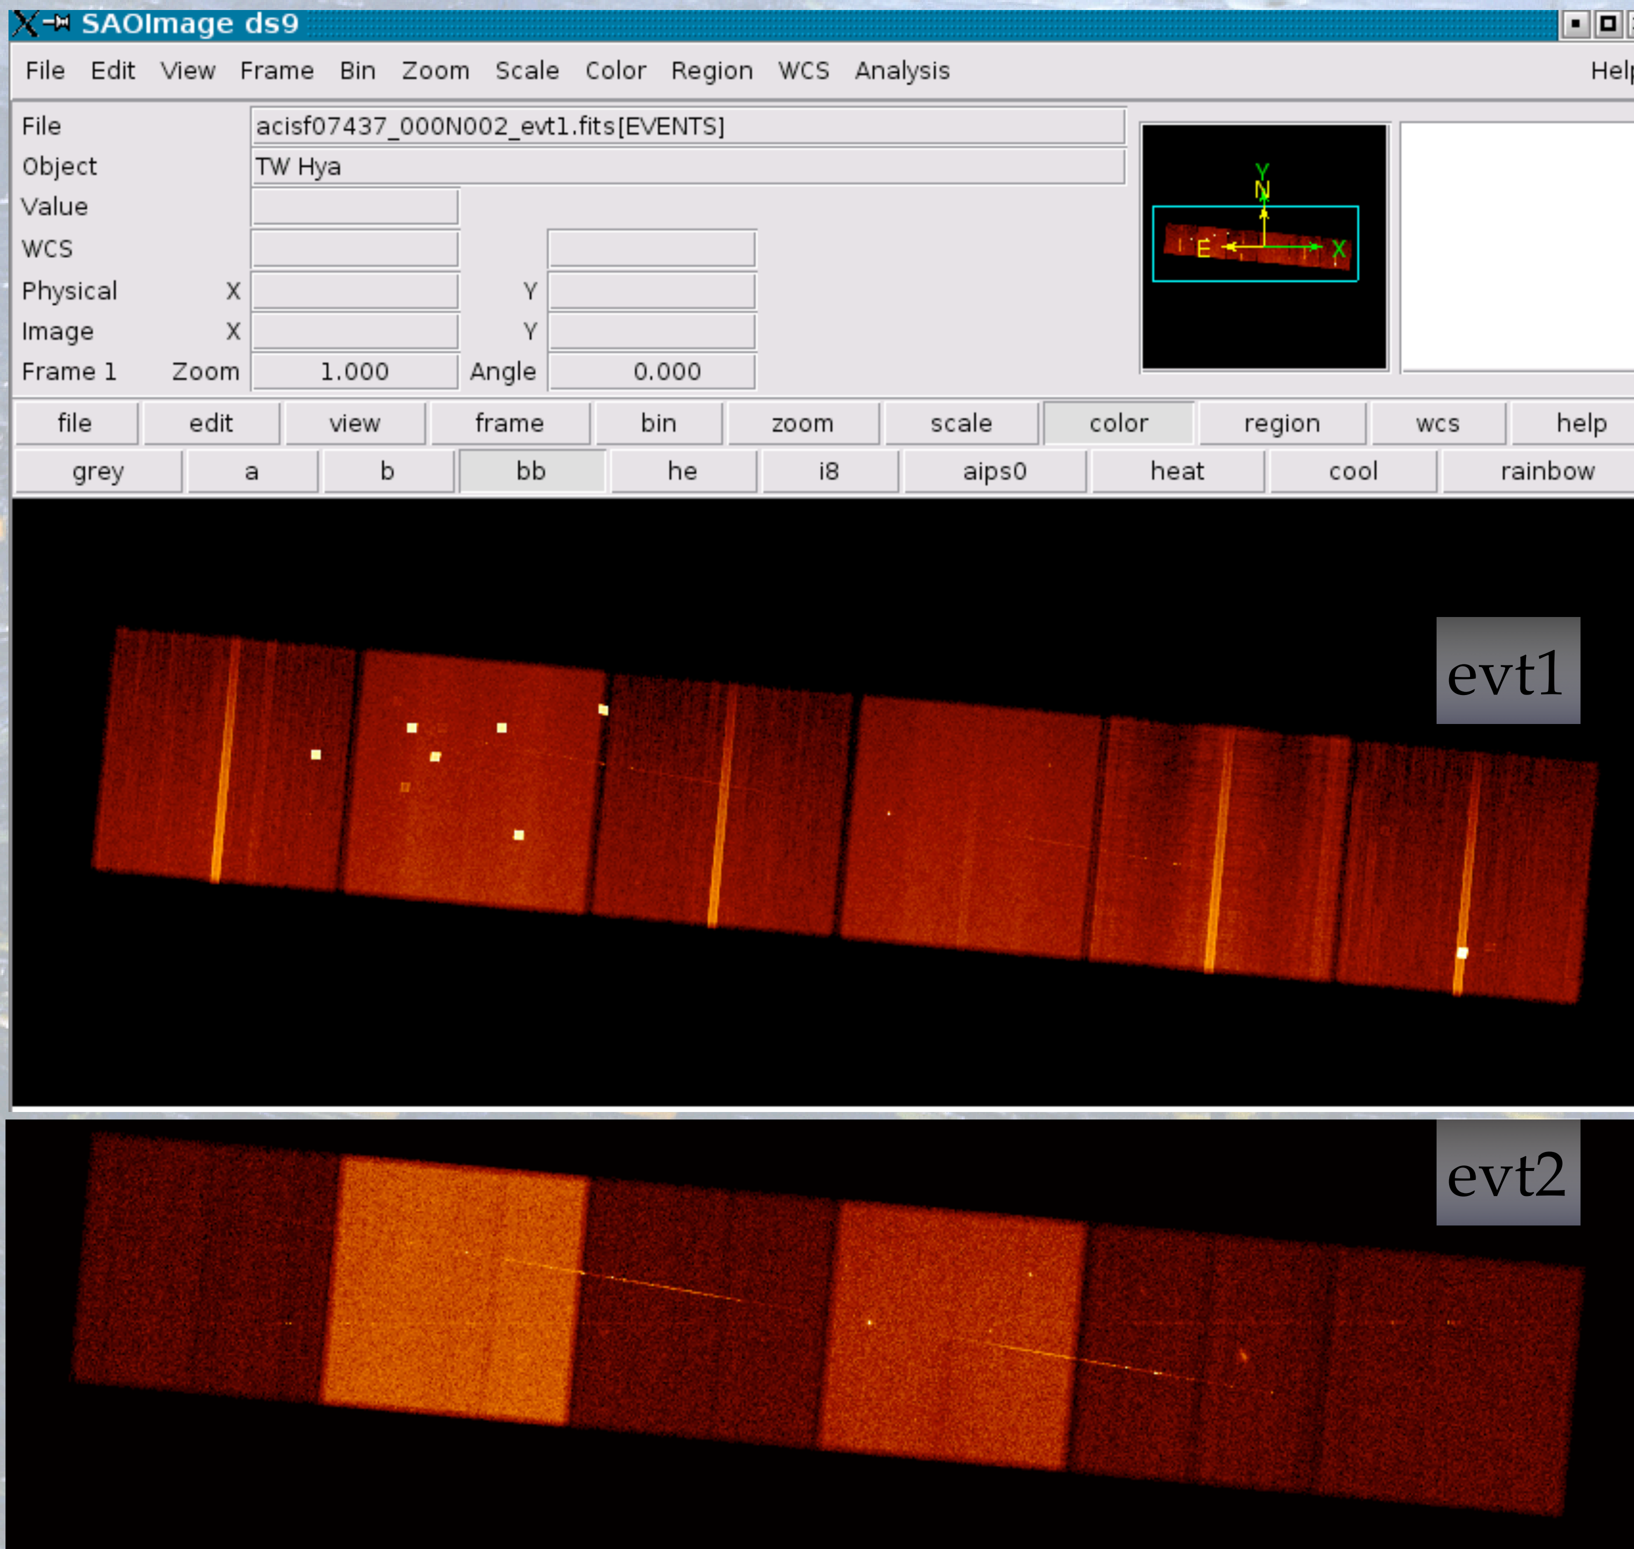

"Bad" data can be filtered from an event list, e.g., with

- `dmcopy` (ciao)
- `fselect` (ftools)

This example shows a Chandra HETG observation for "Level 1" (`evt1` file) and after filtering bad pixels, bad grades (`evt2` file).

Working With Event Lists: XMM Image Display w/ ds9

An EPIC/MOS event file displays in sky coordinates “out of the box” in ds9

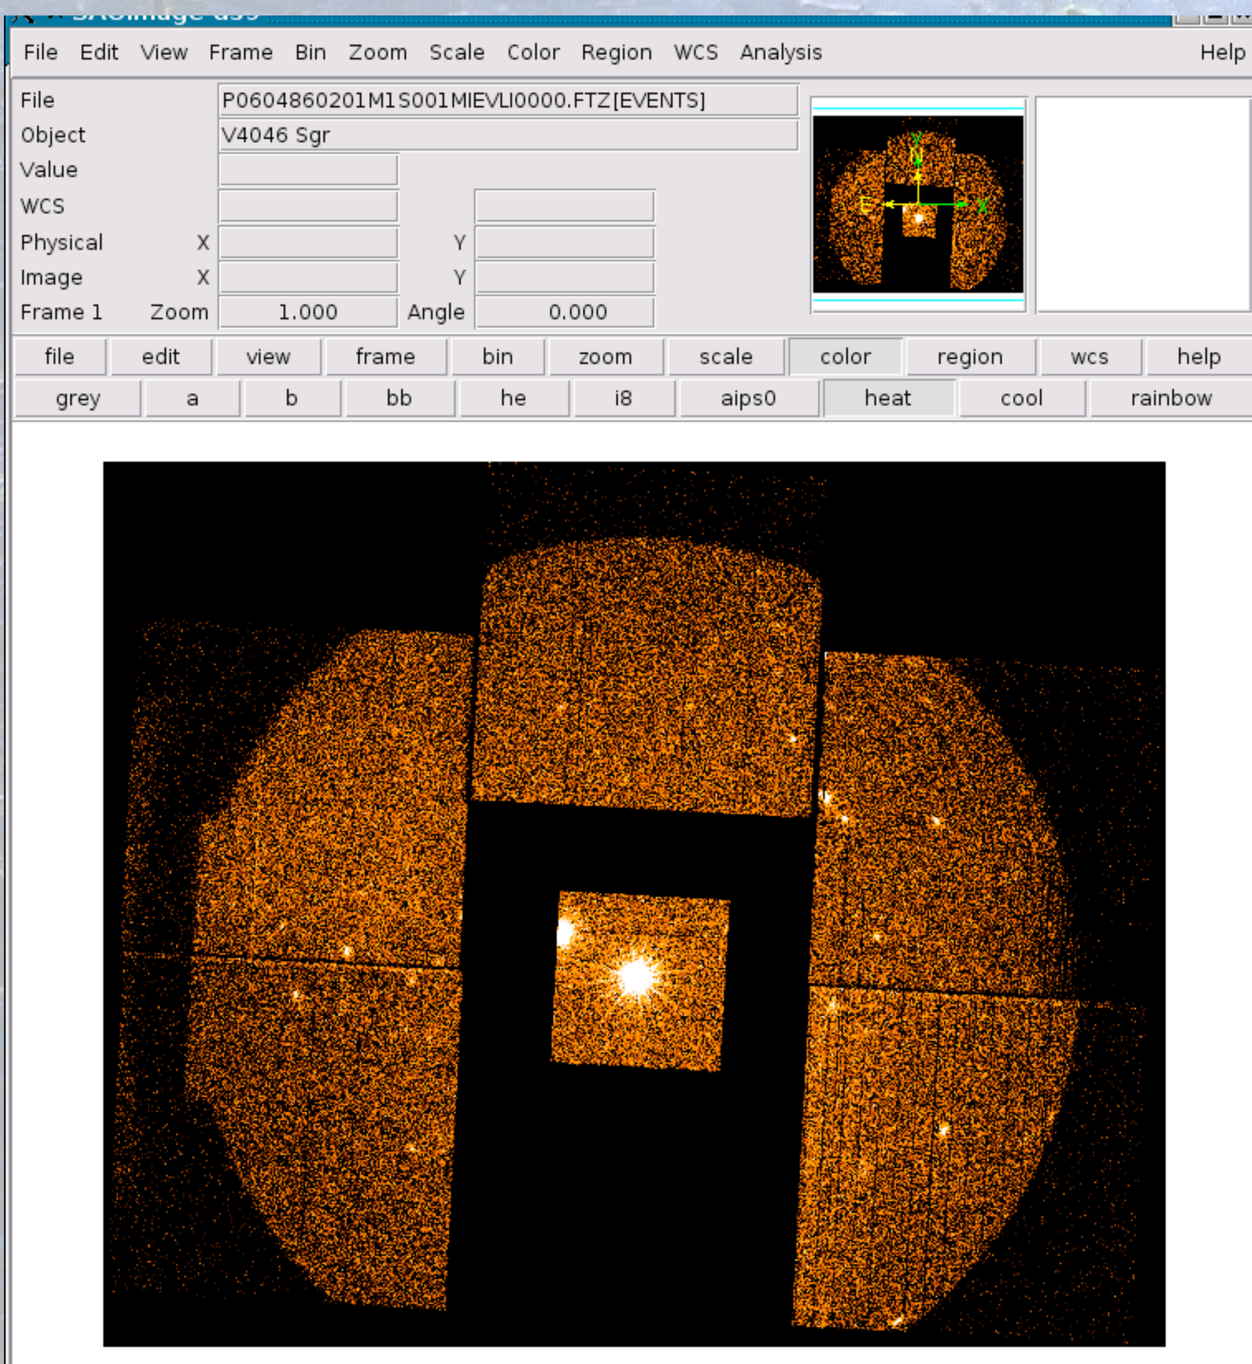

To display an RGS file in dispersion/PI coodinates, use the ds9 “bin->Binning Parameters” dialog:

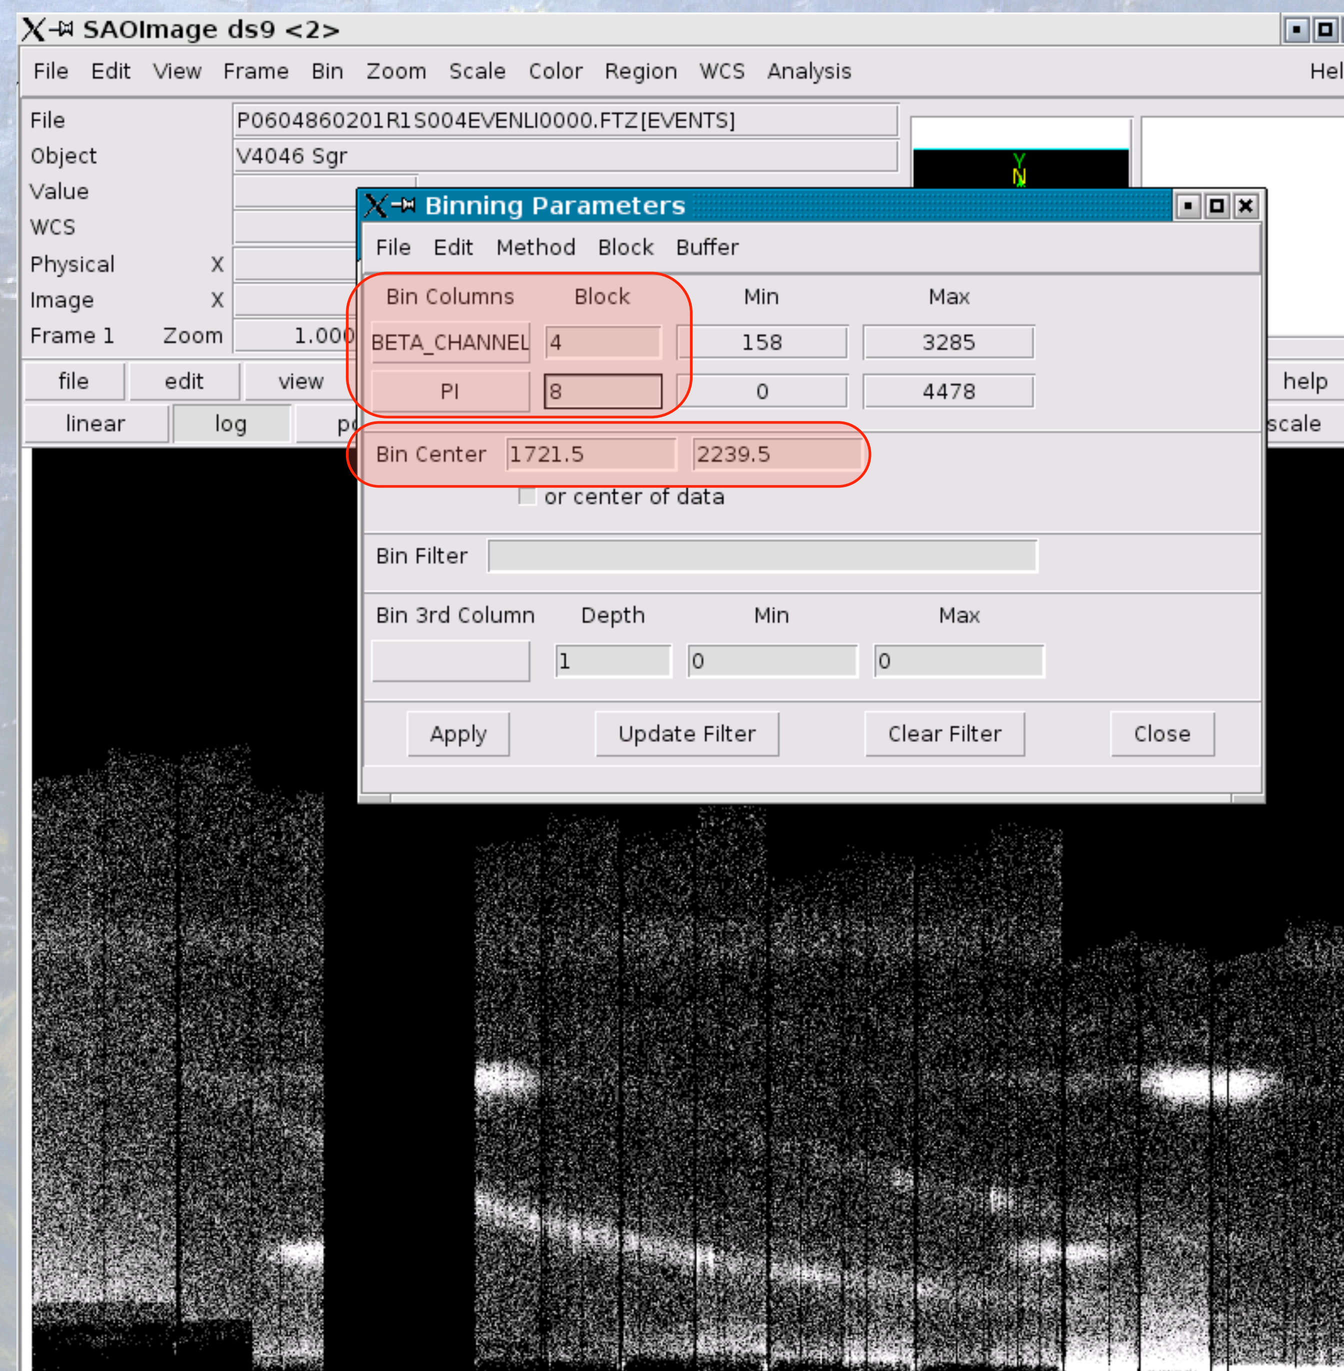

15

Working With Event Lists: Screening Light Curves

Background can be variable and needs to be examined. One method is to look at the light curve of all events in a region with no bright source. Below is an example use of `prism` to view the curve for CCD_ID = 4.

IF we have variable background, AND if it would be significant for our source region, THEN we need to exclude the affected times. IF we have a variable source, we may wish to filter by time on flux states.

We can check for source variability by supplying an appropriate filter, here the MEG first order events, but it could also be a circle about a point source. In this case, the source is variable.

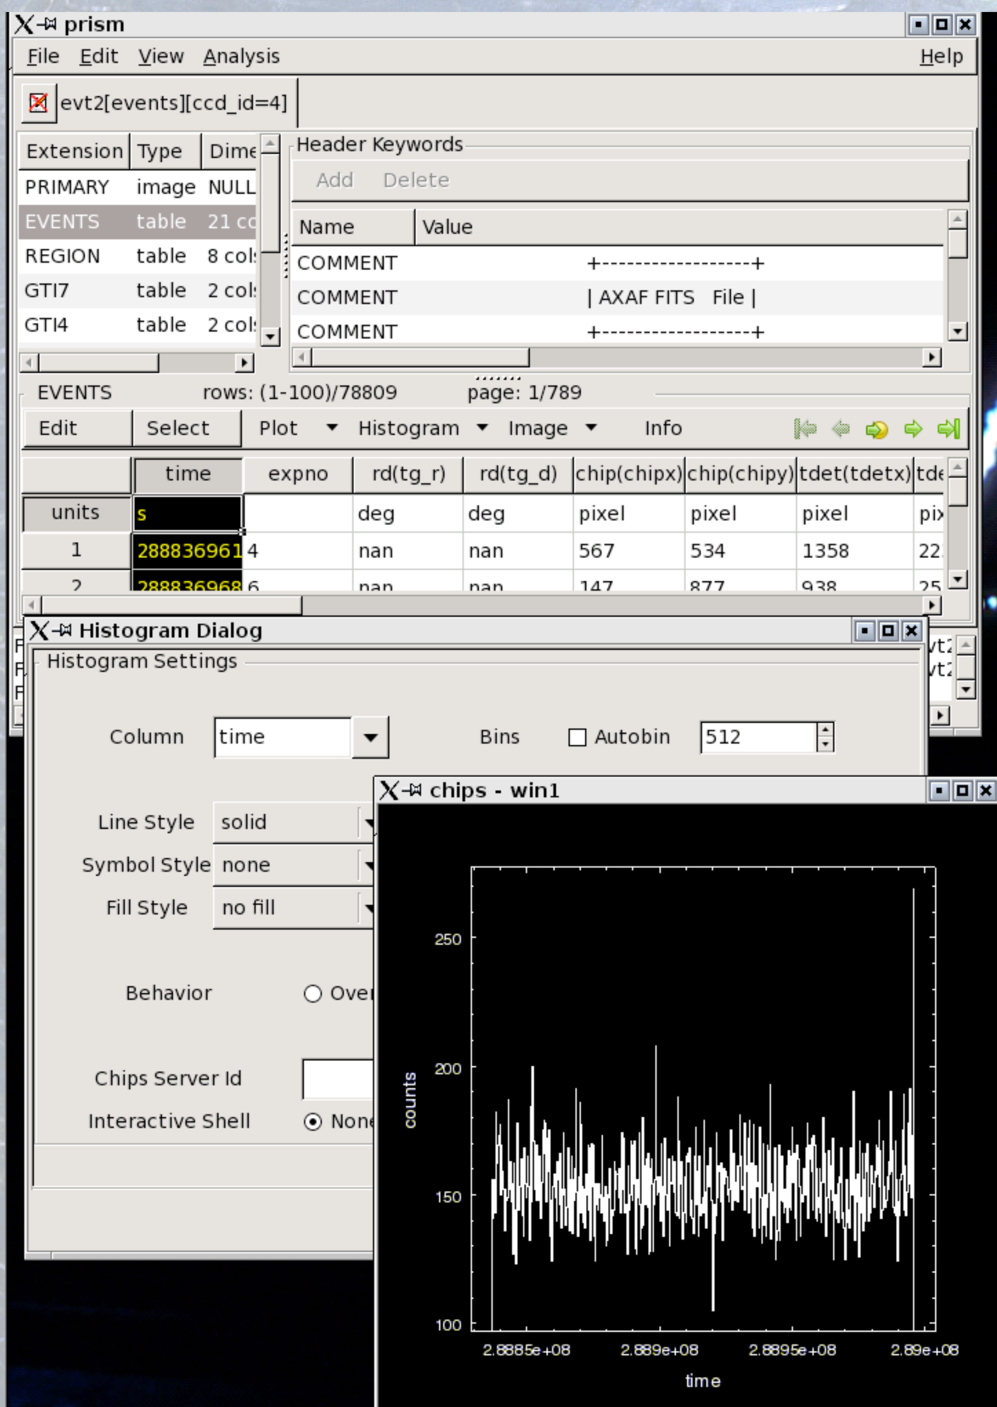

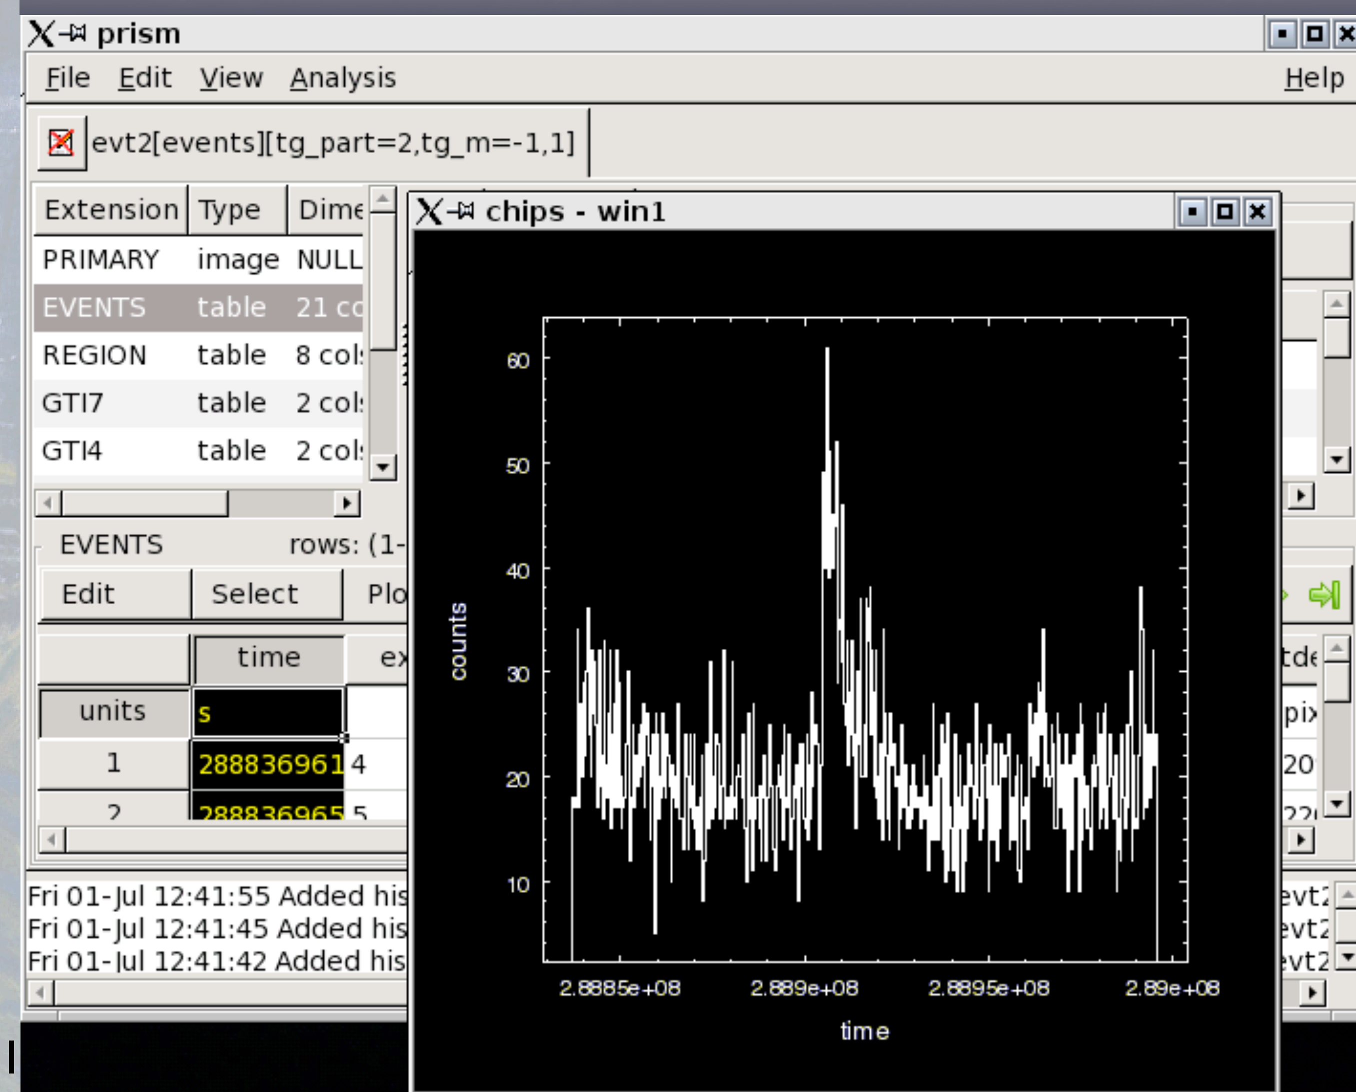

Constructing the "Interesting" Products

Once you have reviewed events, reviewed recommended filters (read the manual), applied any custom filters, you can then prepare your analysis products:

- Image (perhaps in "energy" bands)
- Source list, which is needed for extraction of:
- Spectra (low-resolution, or high-resolution)
- Light Curves (perhaps in "energy" bands)

Depending on the source count rate and the instrument, you may also need to prepare a background object.

The analysis products are usually stored as **FITS** files: FITS == "Flexible Image *Transport*". There are standard OGIP ("Office of Guest Investigator Programs") specifications for images, spectra, and light curves. Specific instruments will have customized fields (within the scope of the official format). But these products are in some sense, *generic*, and can be used in many analysis systems with a little care.

Calibration, *Calibration*, & ***Calibration***

Calibration is an ongoing process for any observatory. Some terms are time-dependent, and even if they are not, they are improved as more data are obtained. Only when the observatory ceases to function can the calibration be frozen.

You need an up-to-date Calibration Database ("CALDB").

But *beware*: **Calibration is not perfect**; it contains both statistical and systematic uncertainties. And be *aware* of the statistical limits of your data AND the accuracy of the calibration AND the possibility of systematic uncertainties (particularly when comparing one observatory's data to another's).

CALDB files generally contain low-level, observation-independent quantities giving component responses (mirror effective area by shell, detector QE, filter transmission, grating efficiency, system geometry). Some components may be time-dependent.

Calibration, *Calibration*, & ***Calibration***

For each analysis product, there are *observation and extraction dependent* high-level calibration objects which need to be prepared.

Imaging:
- Exposure map
- Point Spread Function

Spectroscopy:
- ARF (*Ancillary Response File*; effective area)
- RMF (*Response Matrix File*; spectral redistribution)

Timing:
- (it depends: window functions, dead time corrections..., ARFs if you need flux instead of count rate)

Exposure Map: Formal Definition

An *exposure map* is an observation-dependent quantity which allows you to convert from a sky-coordinate image counts to physical flux --- *approximately.*

A formal definition: borrowing from
“An Introduction to Exposure Maps”
(John Houck; `cxc.harvard.edu/ciao/download/doc/expmap_intro.ps`):

$$C(h, \hat{\mathbf{p}}) = \tau_{\text{eff}} \int d\lambda\, E(h, \lambda, \hat{\mathbf{p}}) S_{\mathcal{F}}(\lambda, \hat{\mathbf{p}})$$

Counts vs channel and sky position

exposure time

wavelength (or energy)

Exposure Map

Source model, vs wavelength and position (PSF-smeared)

Exposure Map: Weighted

After some algebra, binning, and specification of spectral weighting:

$$\int_{\Delta\lambda} d\lambda S_{\mathcal{F}}(\lambda) \approx \frac{C(\Delta h)}{\tau_{\mathrm{eff}} W}.$$

counts image in a channel range (corresponding to the bandpass)

Source model image, integrated over a bandpass

exposure time

weighted exposure map

Exposure Map, and Instrument Map

Conceptually, the exposure map is an *instrument map* convolved with the aspect solution.

An *instrument map* is an image in detector coordinates of the instrument sensitivity, including mirror area, detector QE, ... (the units are [cm^2 counts/photon])

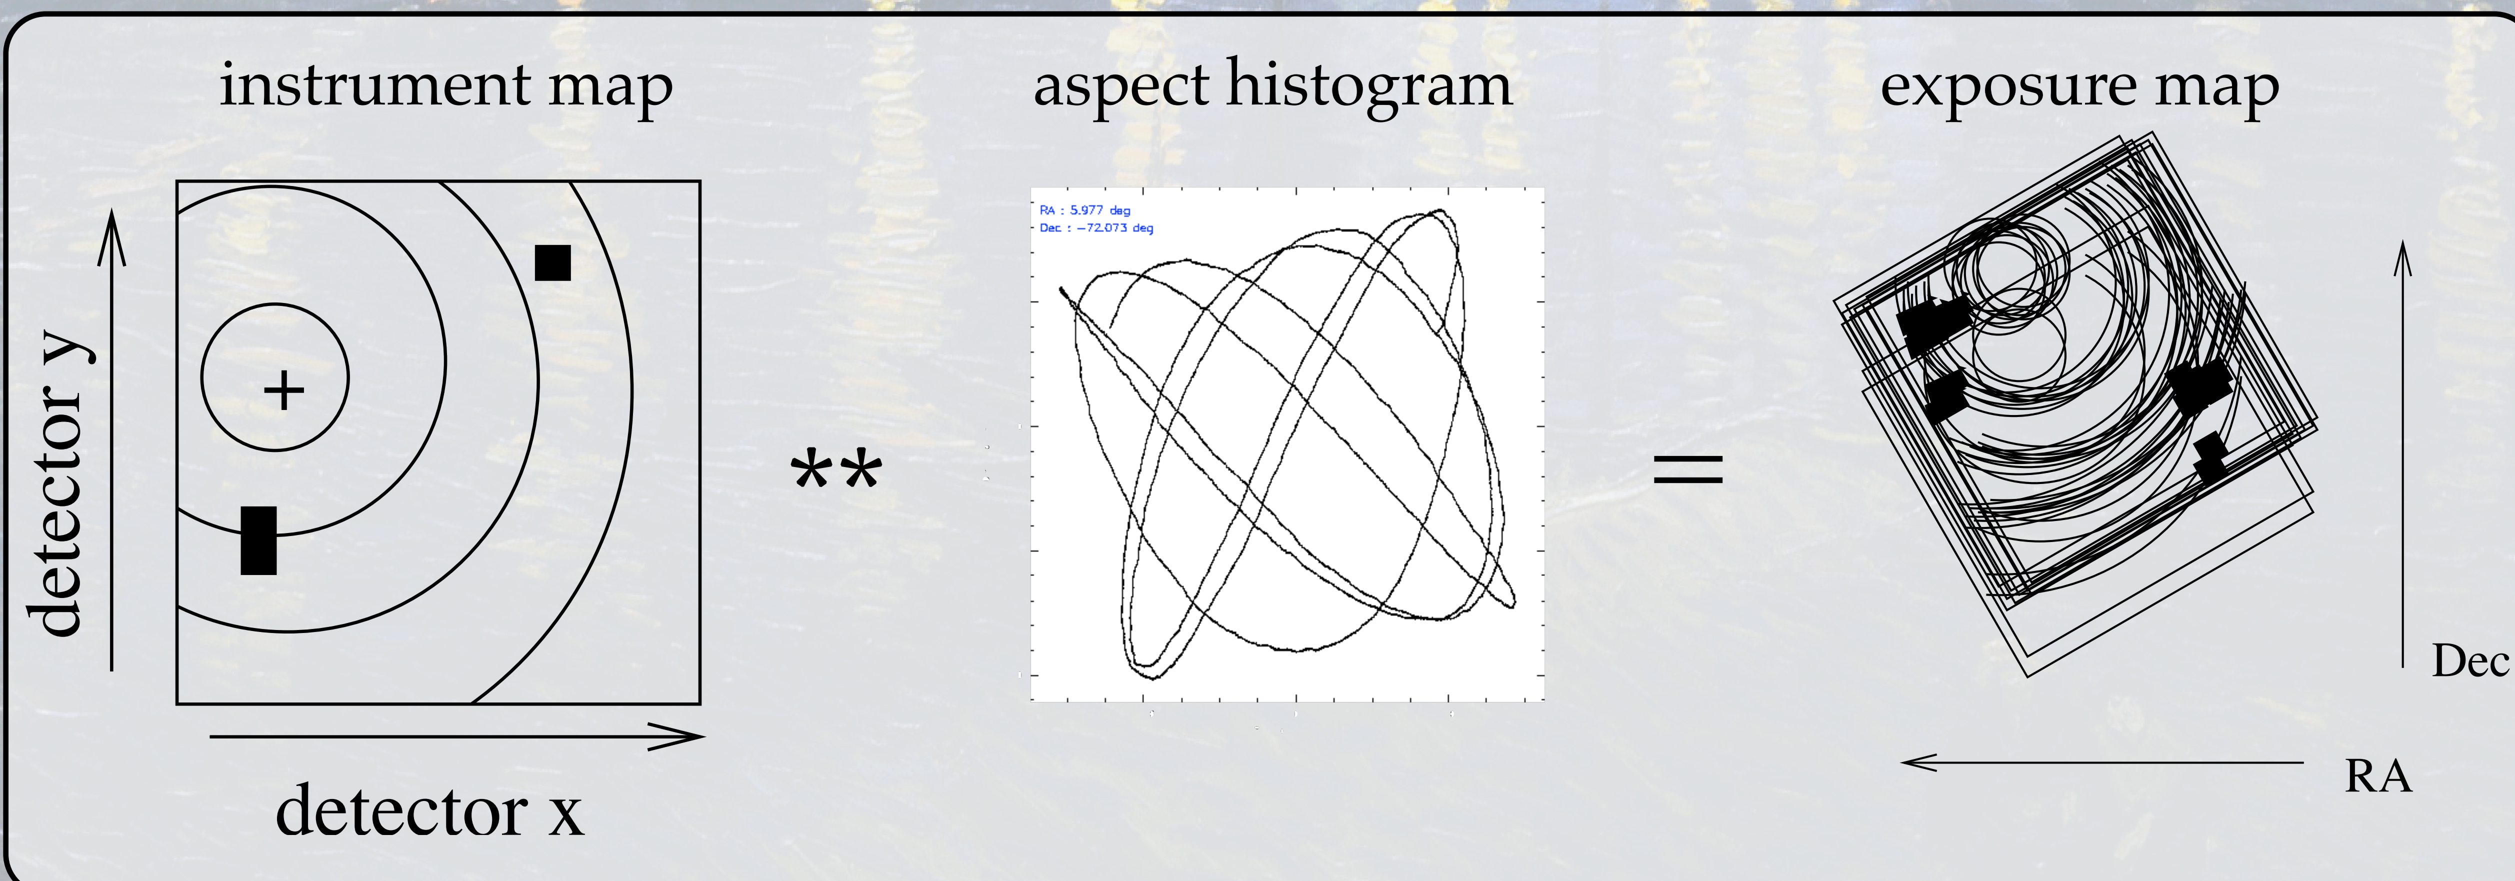

Exposure Map: Examples

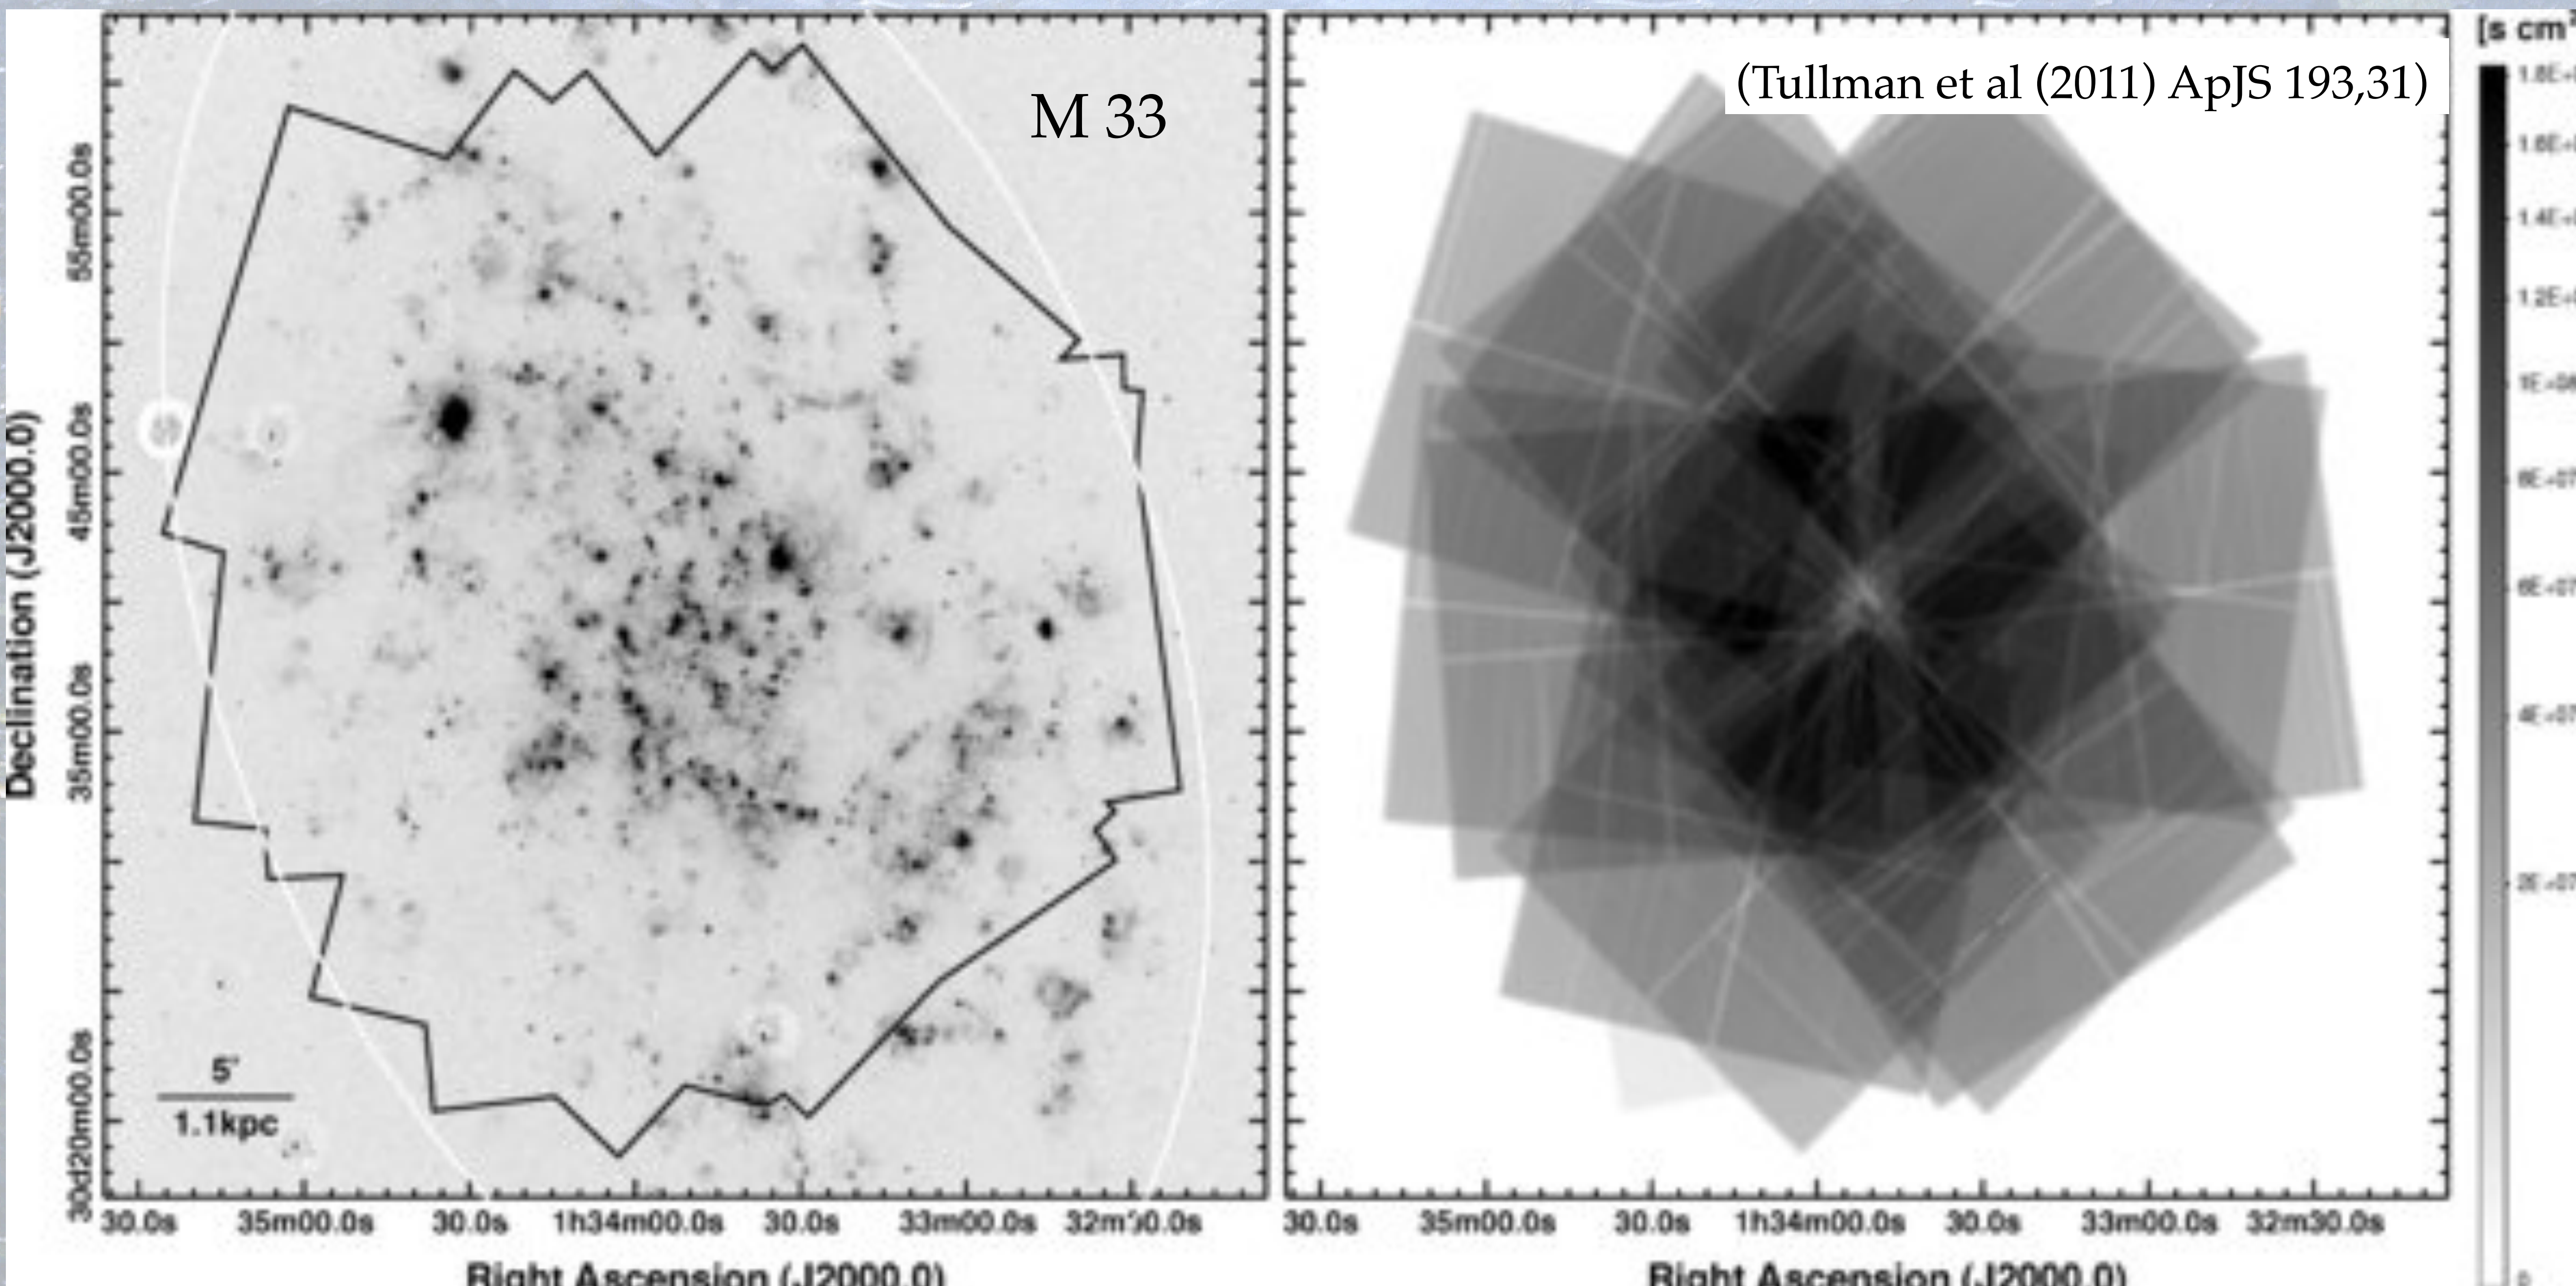

Exposure Map: Examples

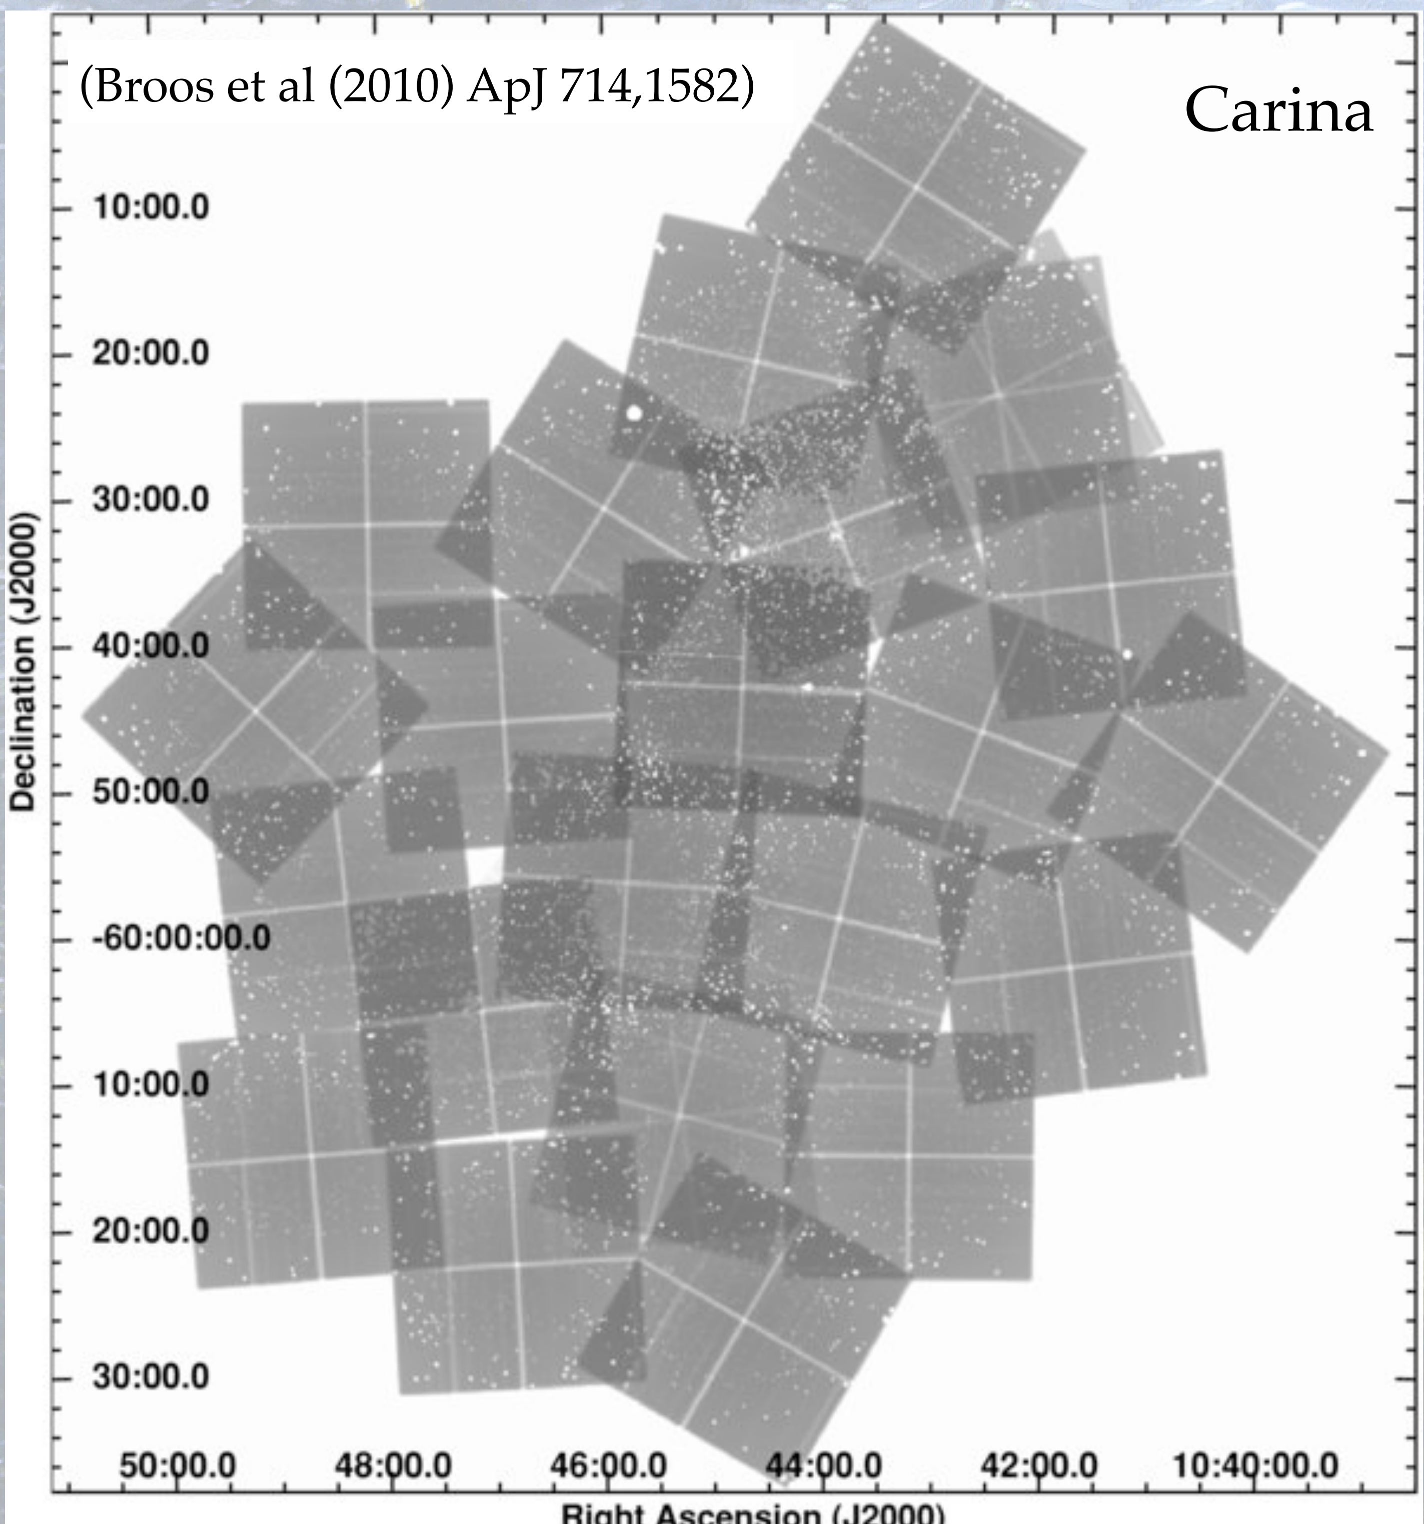

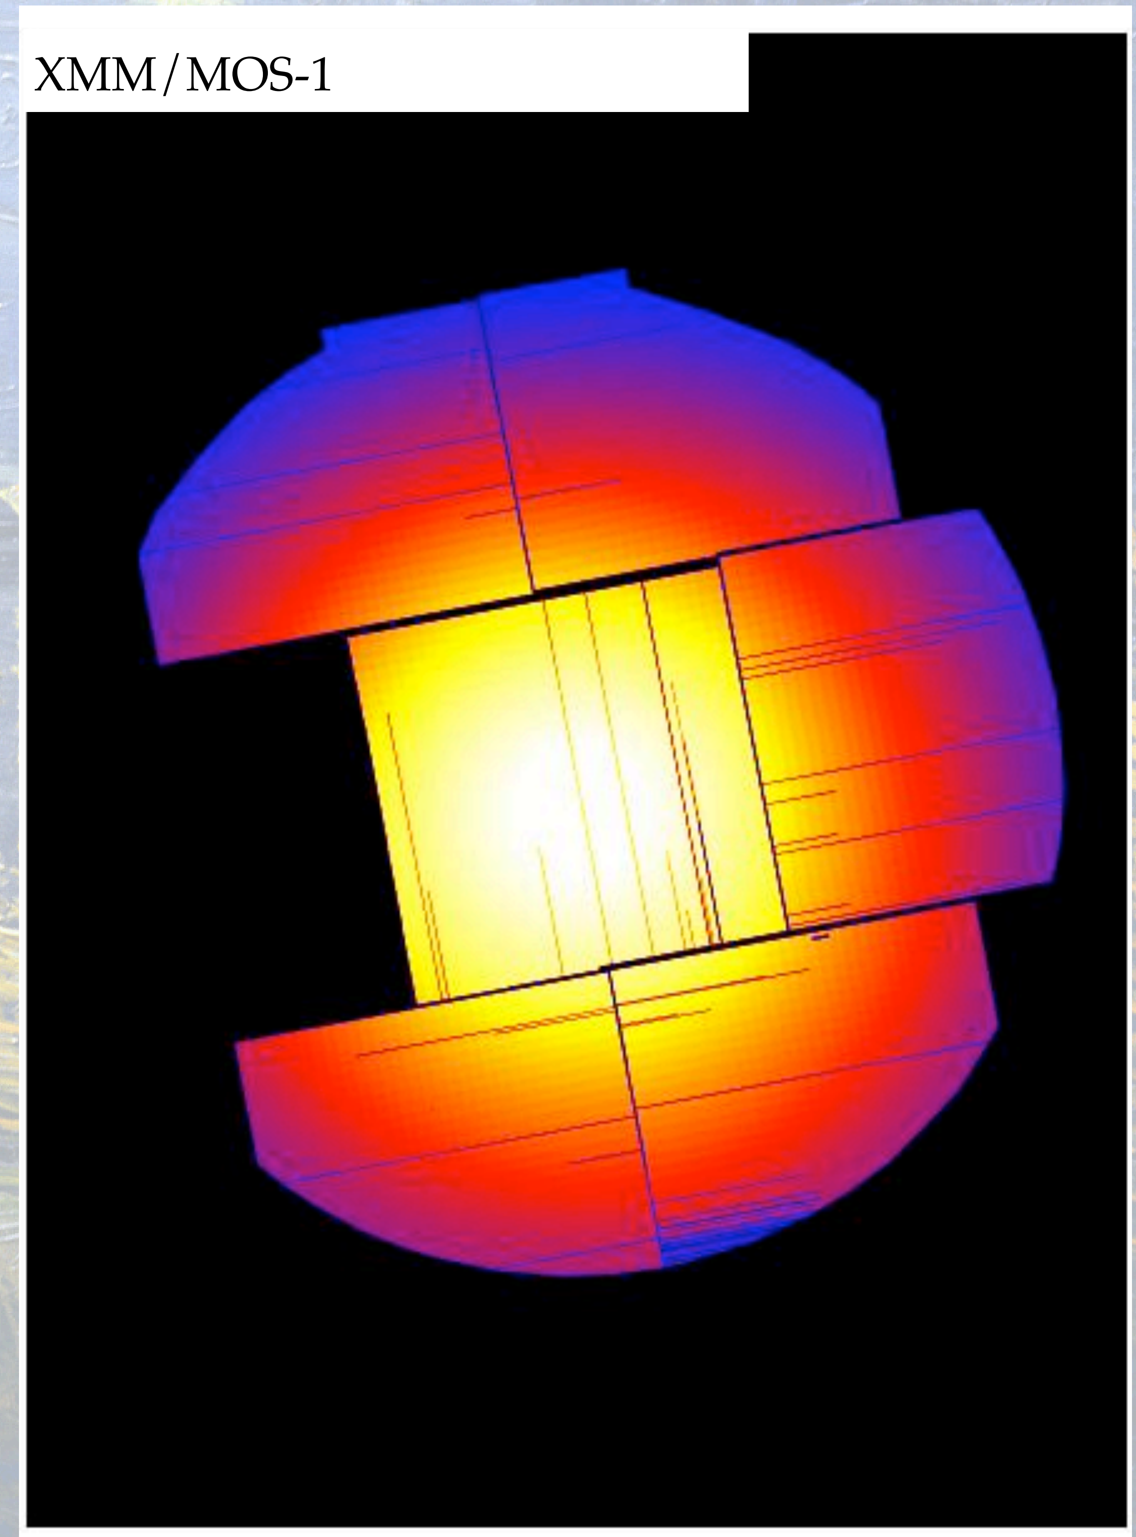

Point Spread Function (PSF)

The *PSF* refers to the redistribution of the direction of rays from a point source. It depends on energy, off-axis angle, aspect reconstruction accuracy, and detector sampling (pixelation). You need to know the PSF to determine *source extent* (morphology), or *flux* (enclosed energy fraction).

Chandra, vs radius and energy

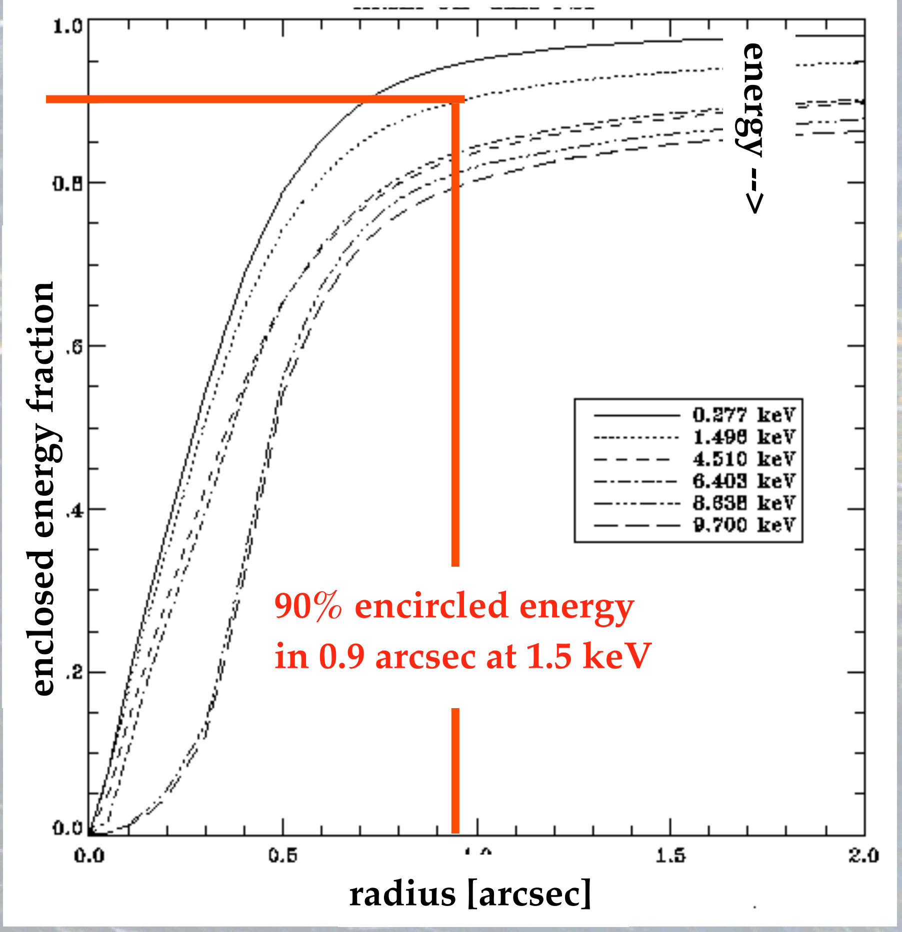

Chandra, vs off-axis angle

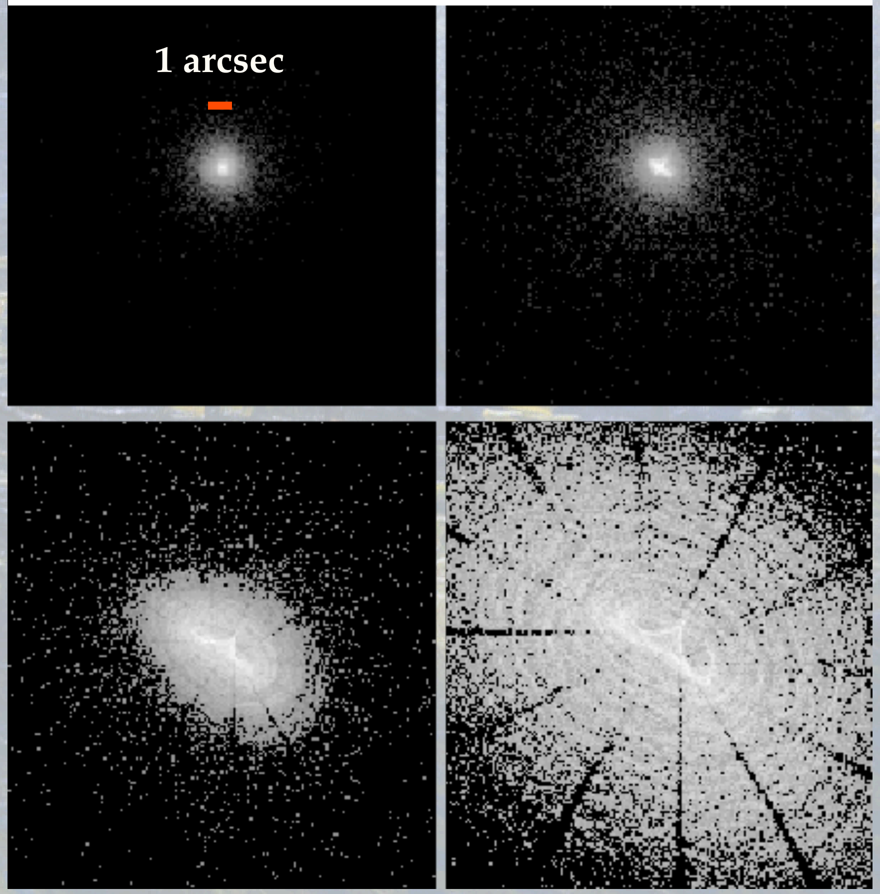

Suzaku

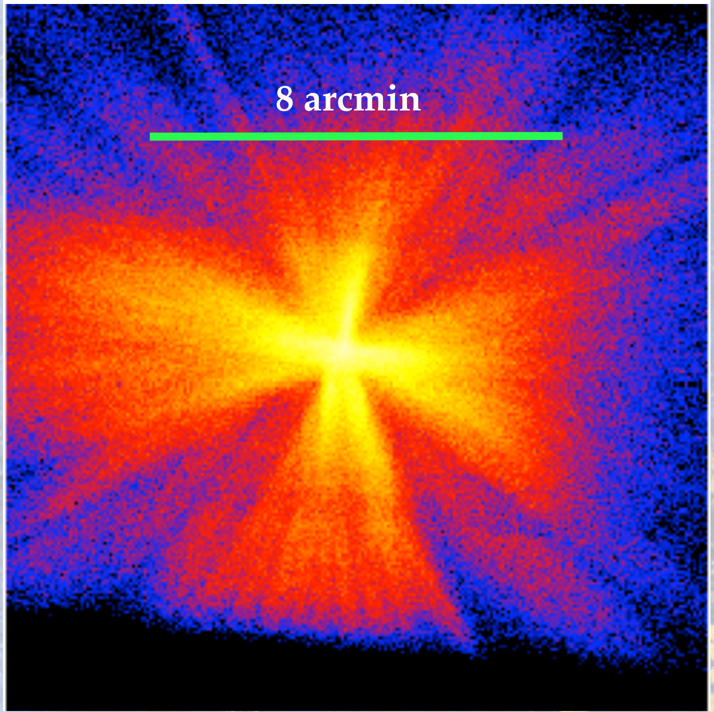

XMM/Newton

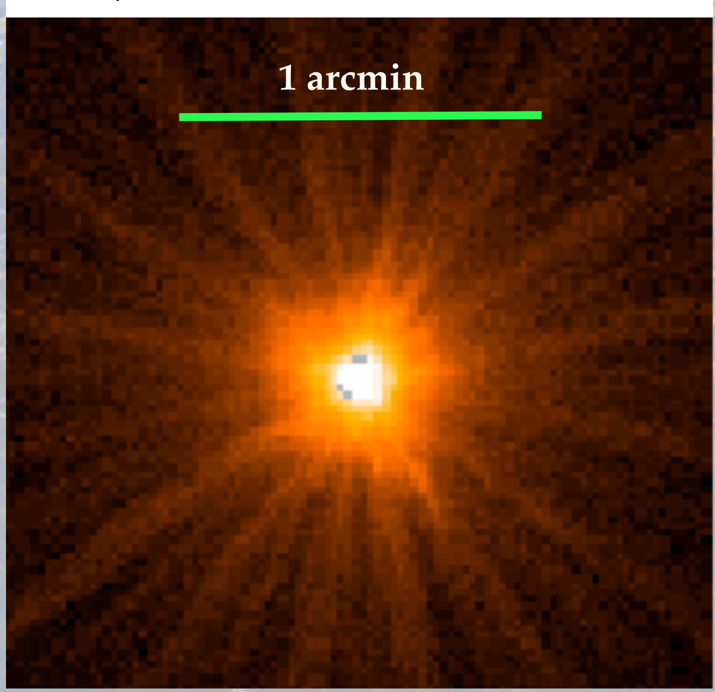

PSF for Analysis

PSF computation can be difficult and time-consuming, since the structure is not easily parameterized, and there are strong dependencies on energy and angles. Accurate PSFs for analysis generally require use of ray-trace models.

Ray-trace:

- Chandra end-to-end: MARX
- Chandra mirrors: CHART + MARX
- XMM/SciSim
- Suzaku: xisrmfgen (for ARFs)

ARF (Effective Area)

ARF: Ancillary Response File -- this contains the effective area as a function of energy for an extraction (*source dependent!*) of an observation (*observation dependent!*) The ARF includes the efficiencies of:

- mirrors (angle dependent)
- gratings
- filters (possibly epoch dependent)
- detector (epoch and readout mode dependent)

Since there are local non-uniformities on detectors, the ARF depends on pointing history (aspect) and on extraction regions.

Units: [cm^2 counts/photon]

ARF generation is mission-specific. ARF use (as a standard FITS format) is generic.

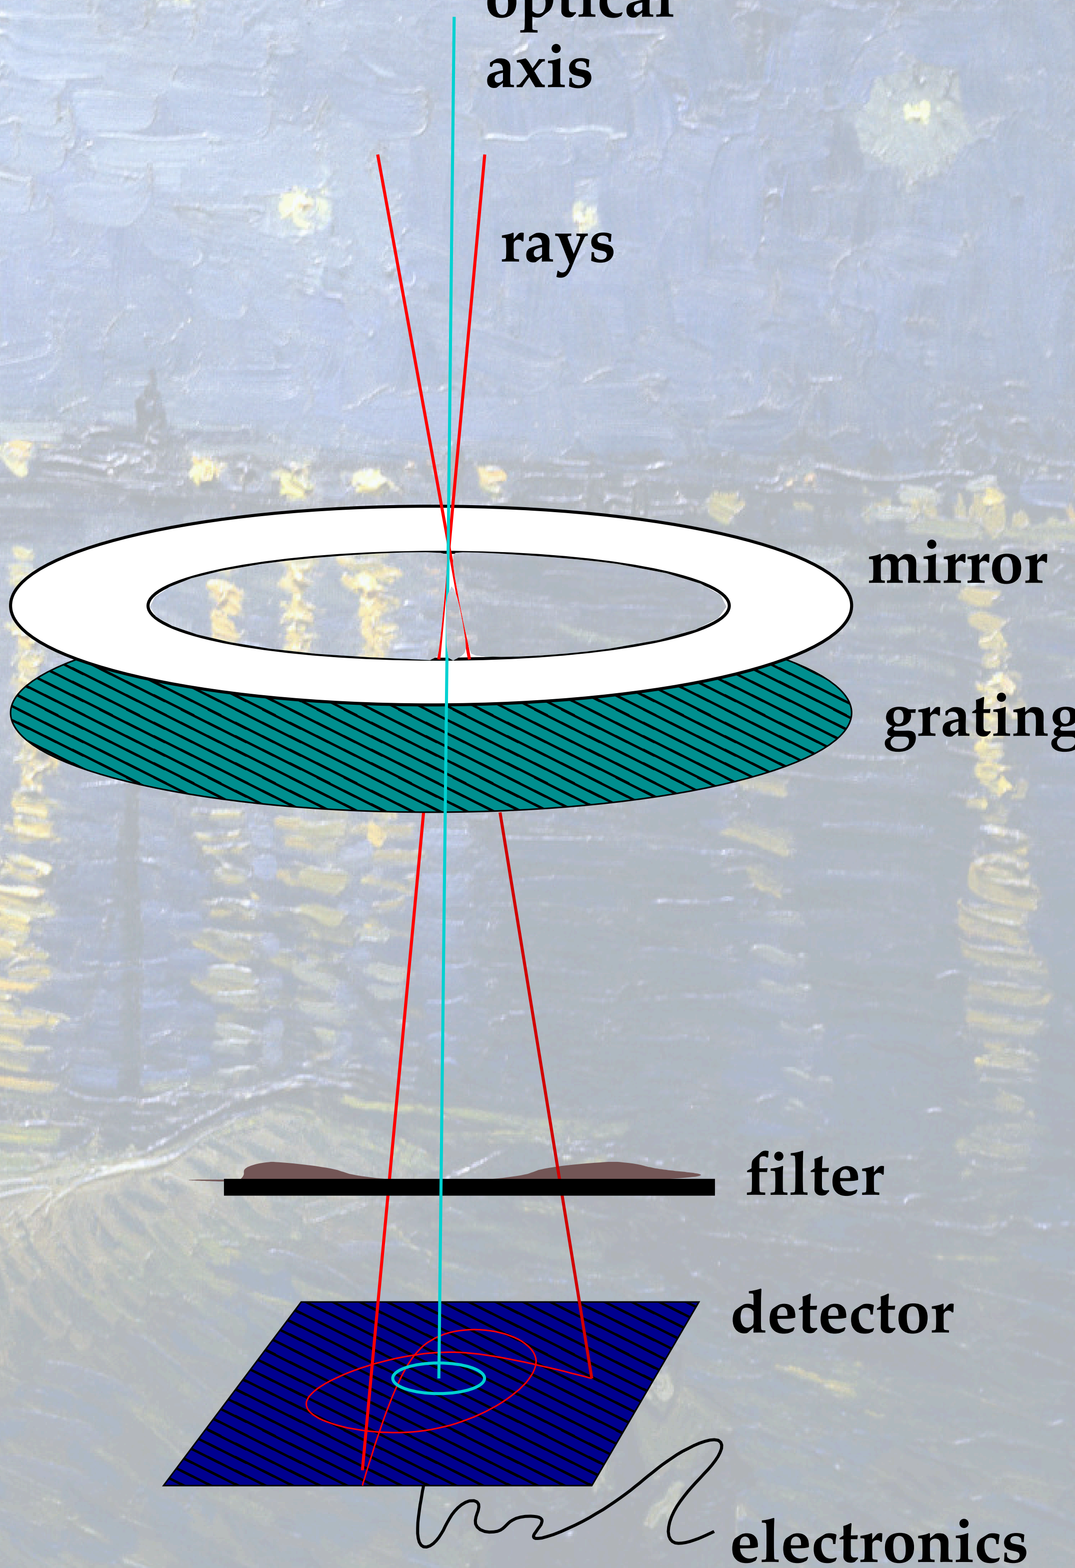

ARF: examples

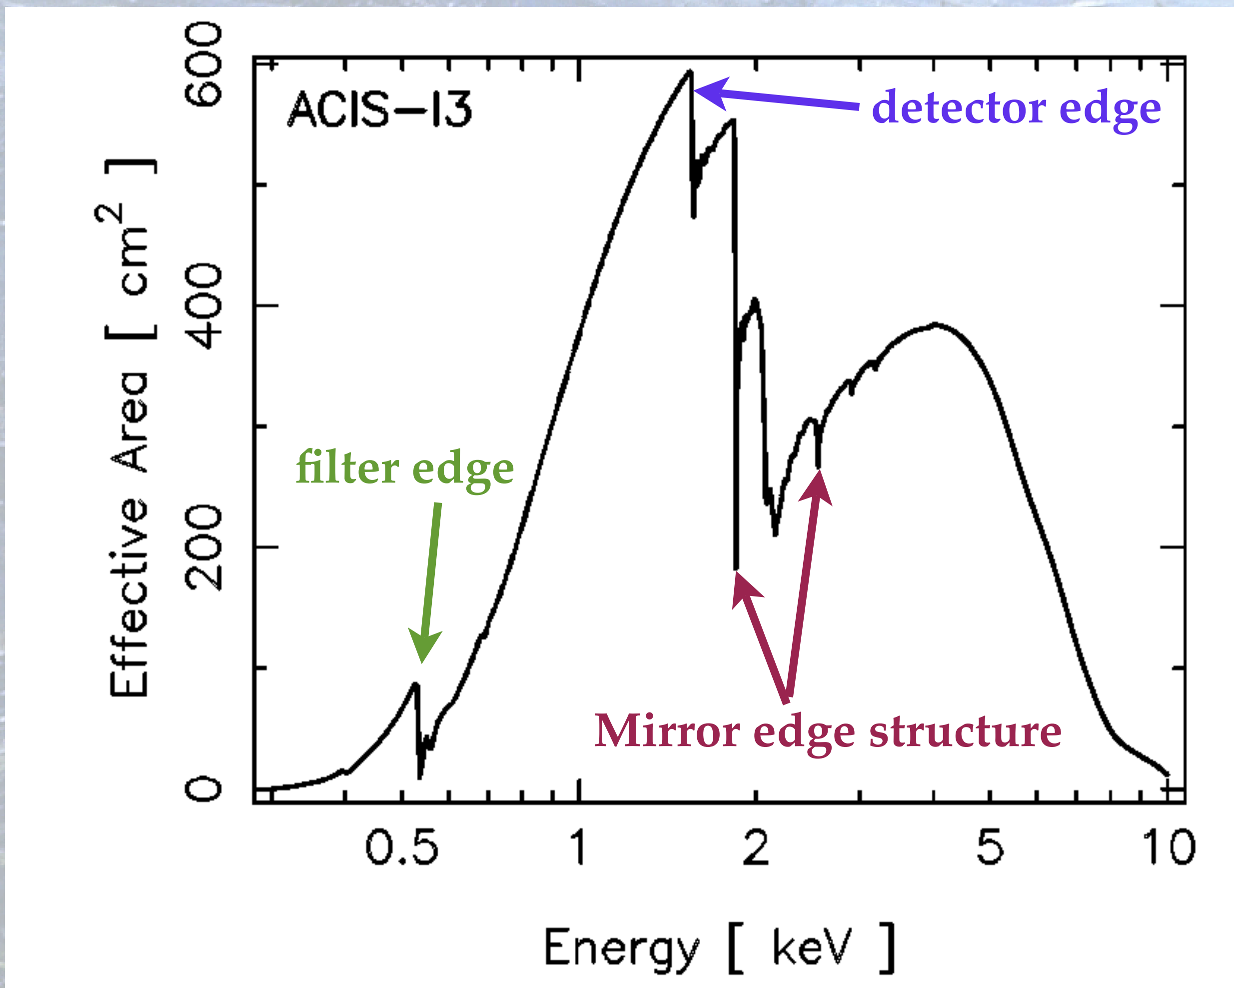

ARFs have strong spectral structure.

The amplitude depends on the extraction region (through the PSF).

ARF shape and amplitude depend on the aspect history (through pointing/dither and detector non-uniformities; with diffraction grating, position maps to energy).

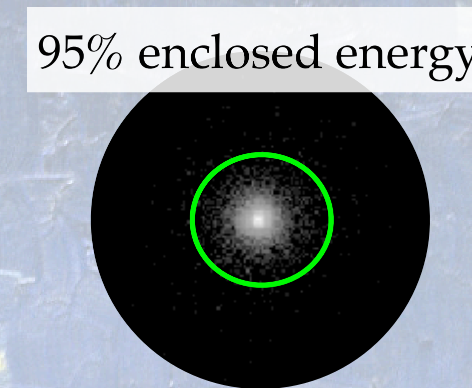

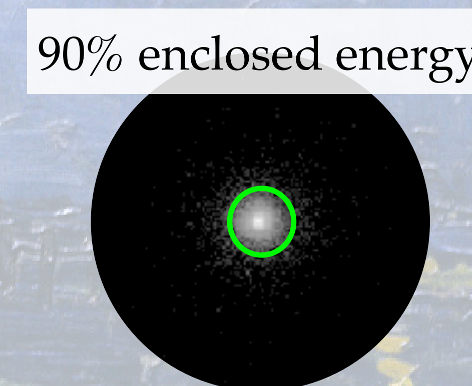

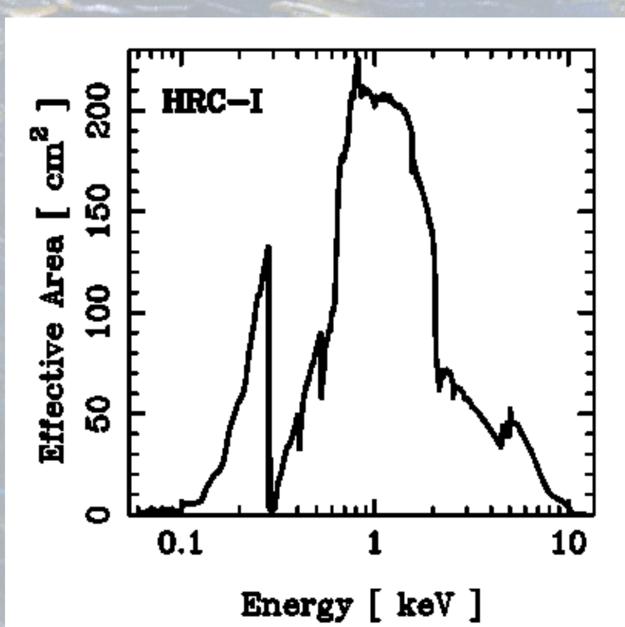

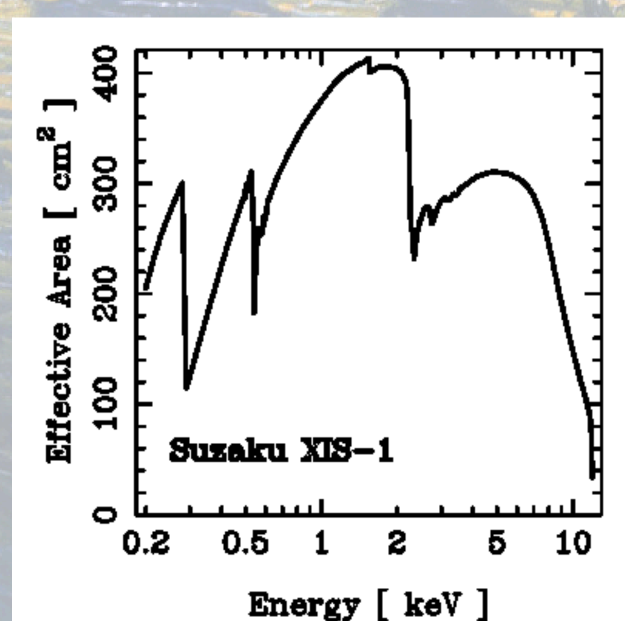

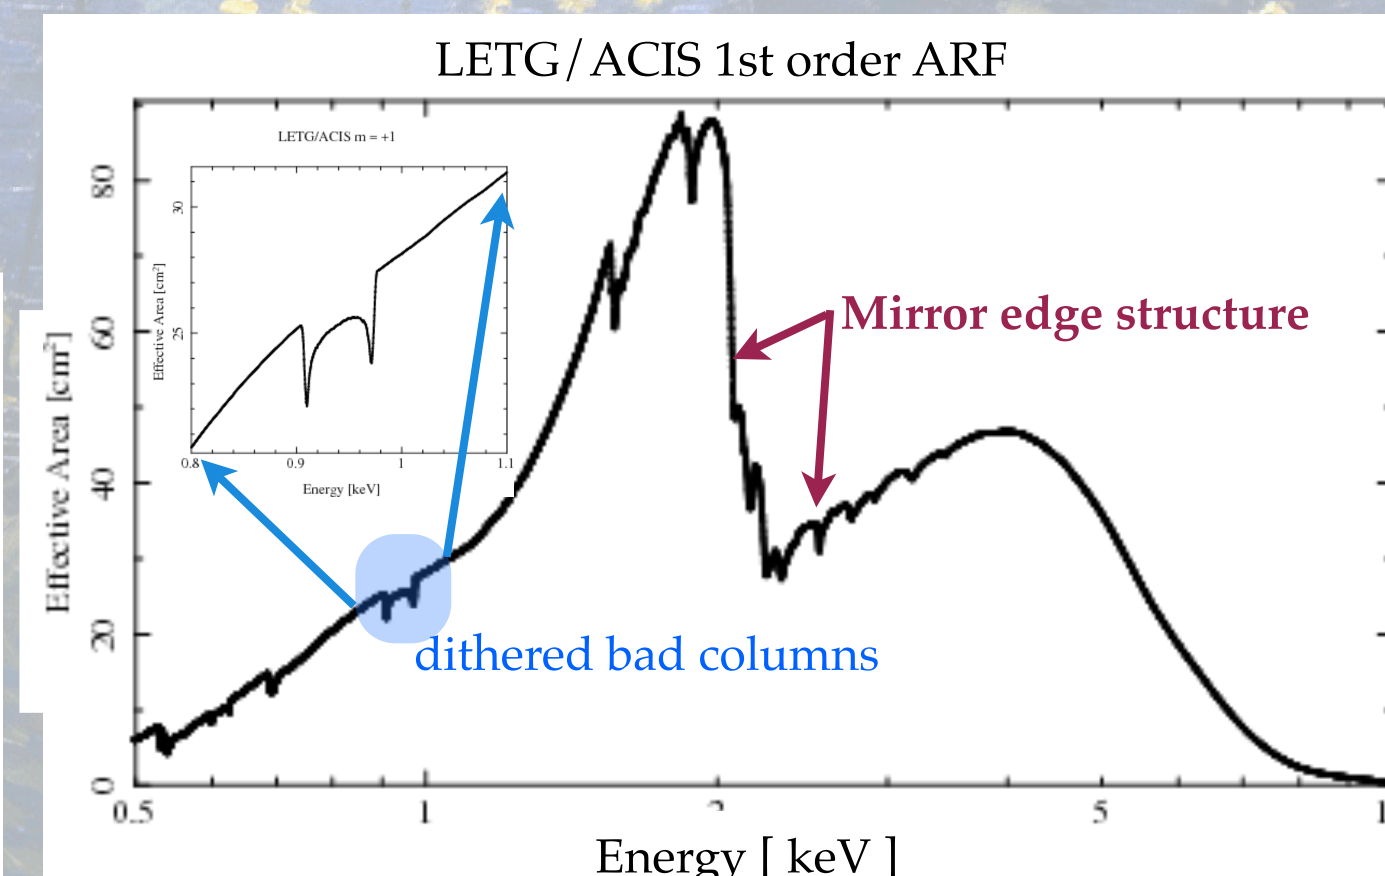

28

Response Matrix File: RMF

The RMF (a standardized format FITS file) encodes the relationship between the *incident photon energy* and the output signal's *distribution over channels.*

PHA*:* originally, "Pulse Height Analyzer", now "Pulse Height Amplitude"

High energy photons interact with the detector material (e.g., silicon for CCDs, xenon gas for a proportional counter, or glass pores for a micro-channel plate). Some number of electrons are released in proportion to the photon energy, but not without statistical and other physical effects.

CCD example PHA generation.
A photon of some energy generates an electron charge cloud. The PHA is the sum of the charge (after conversion to a digital signal) in the neighborhood (blue bars).

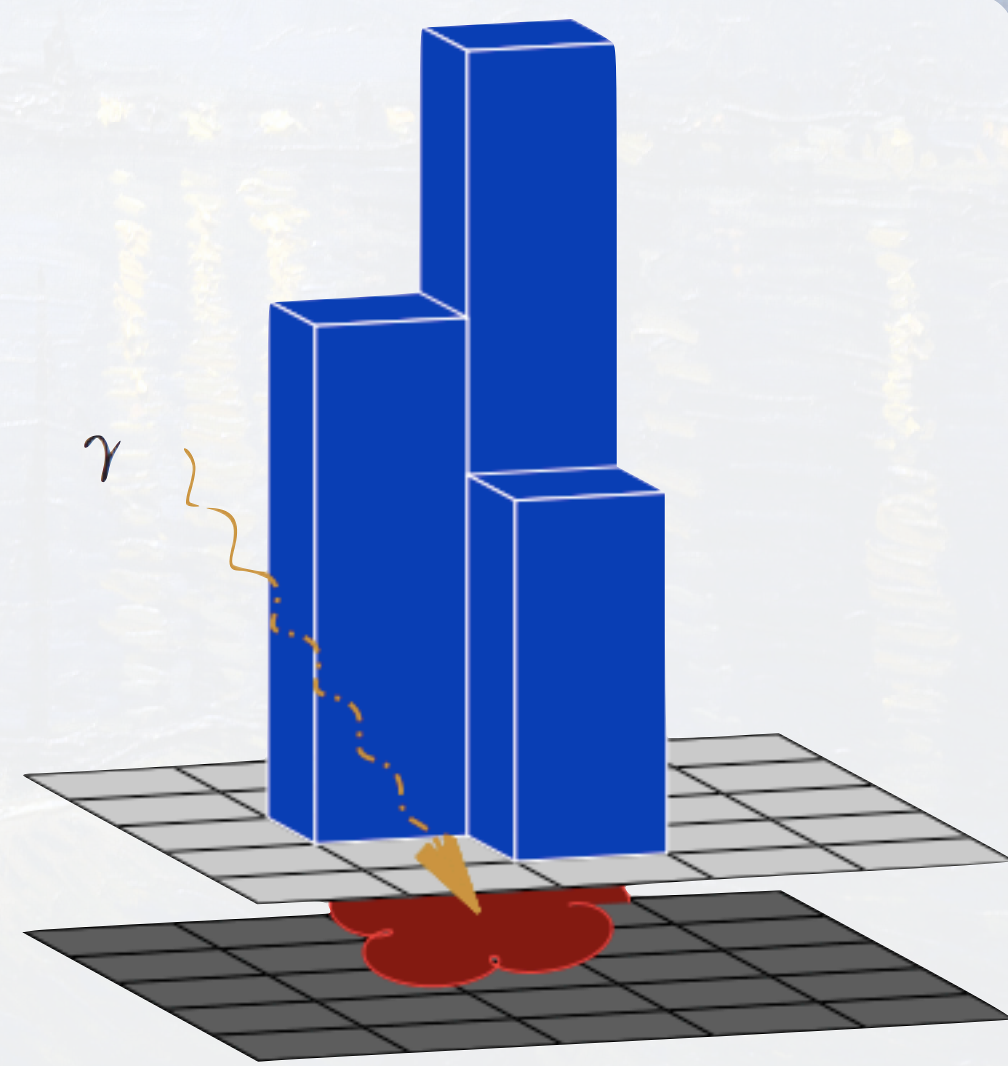

From J.Davis, "Pileup Modeling"
http://www.jedsoft.org/fun/slxfig/pileup2008.pdf

RMF: distribution

An example PHA distribution of Chandra CCDs for 7.5 keV photons. This represents ***one column*** of an RMF (remember, the response function is *R(E, h)*, represented by a *2D* matrix; *h* is the output channel, or *PHA*, for input photon energy *E*).

NOTE: this pseudo-energy is a convenient scaling of the output PHA (or channel). All input photons had a single energy!

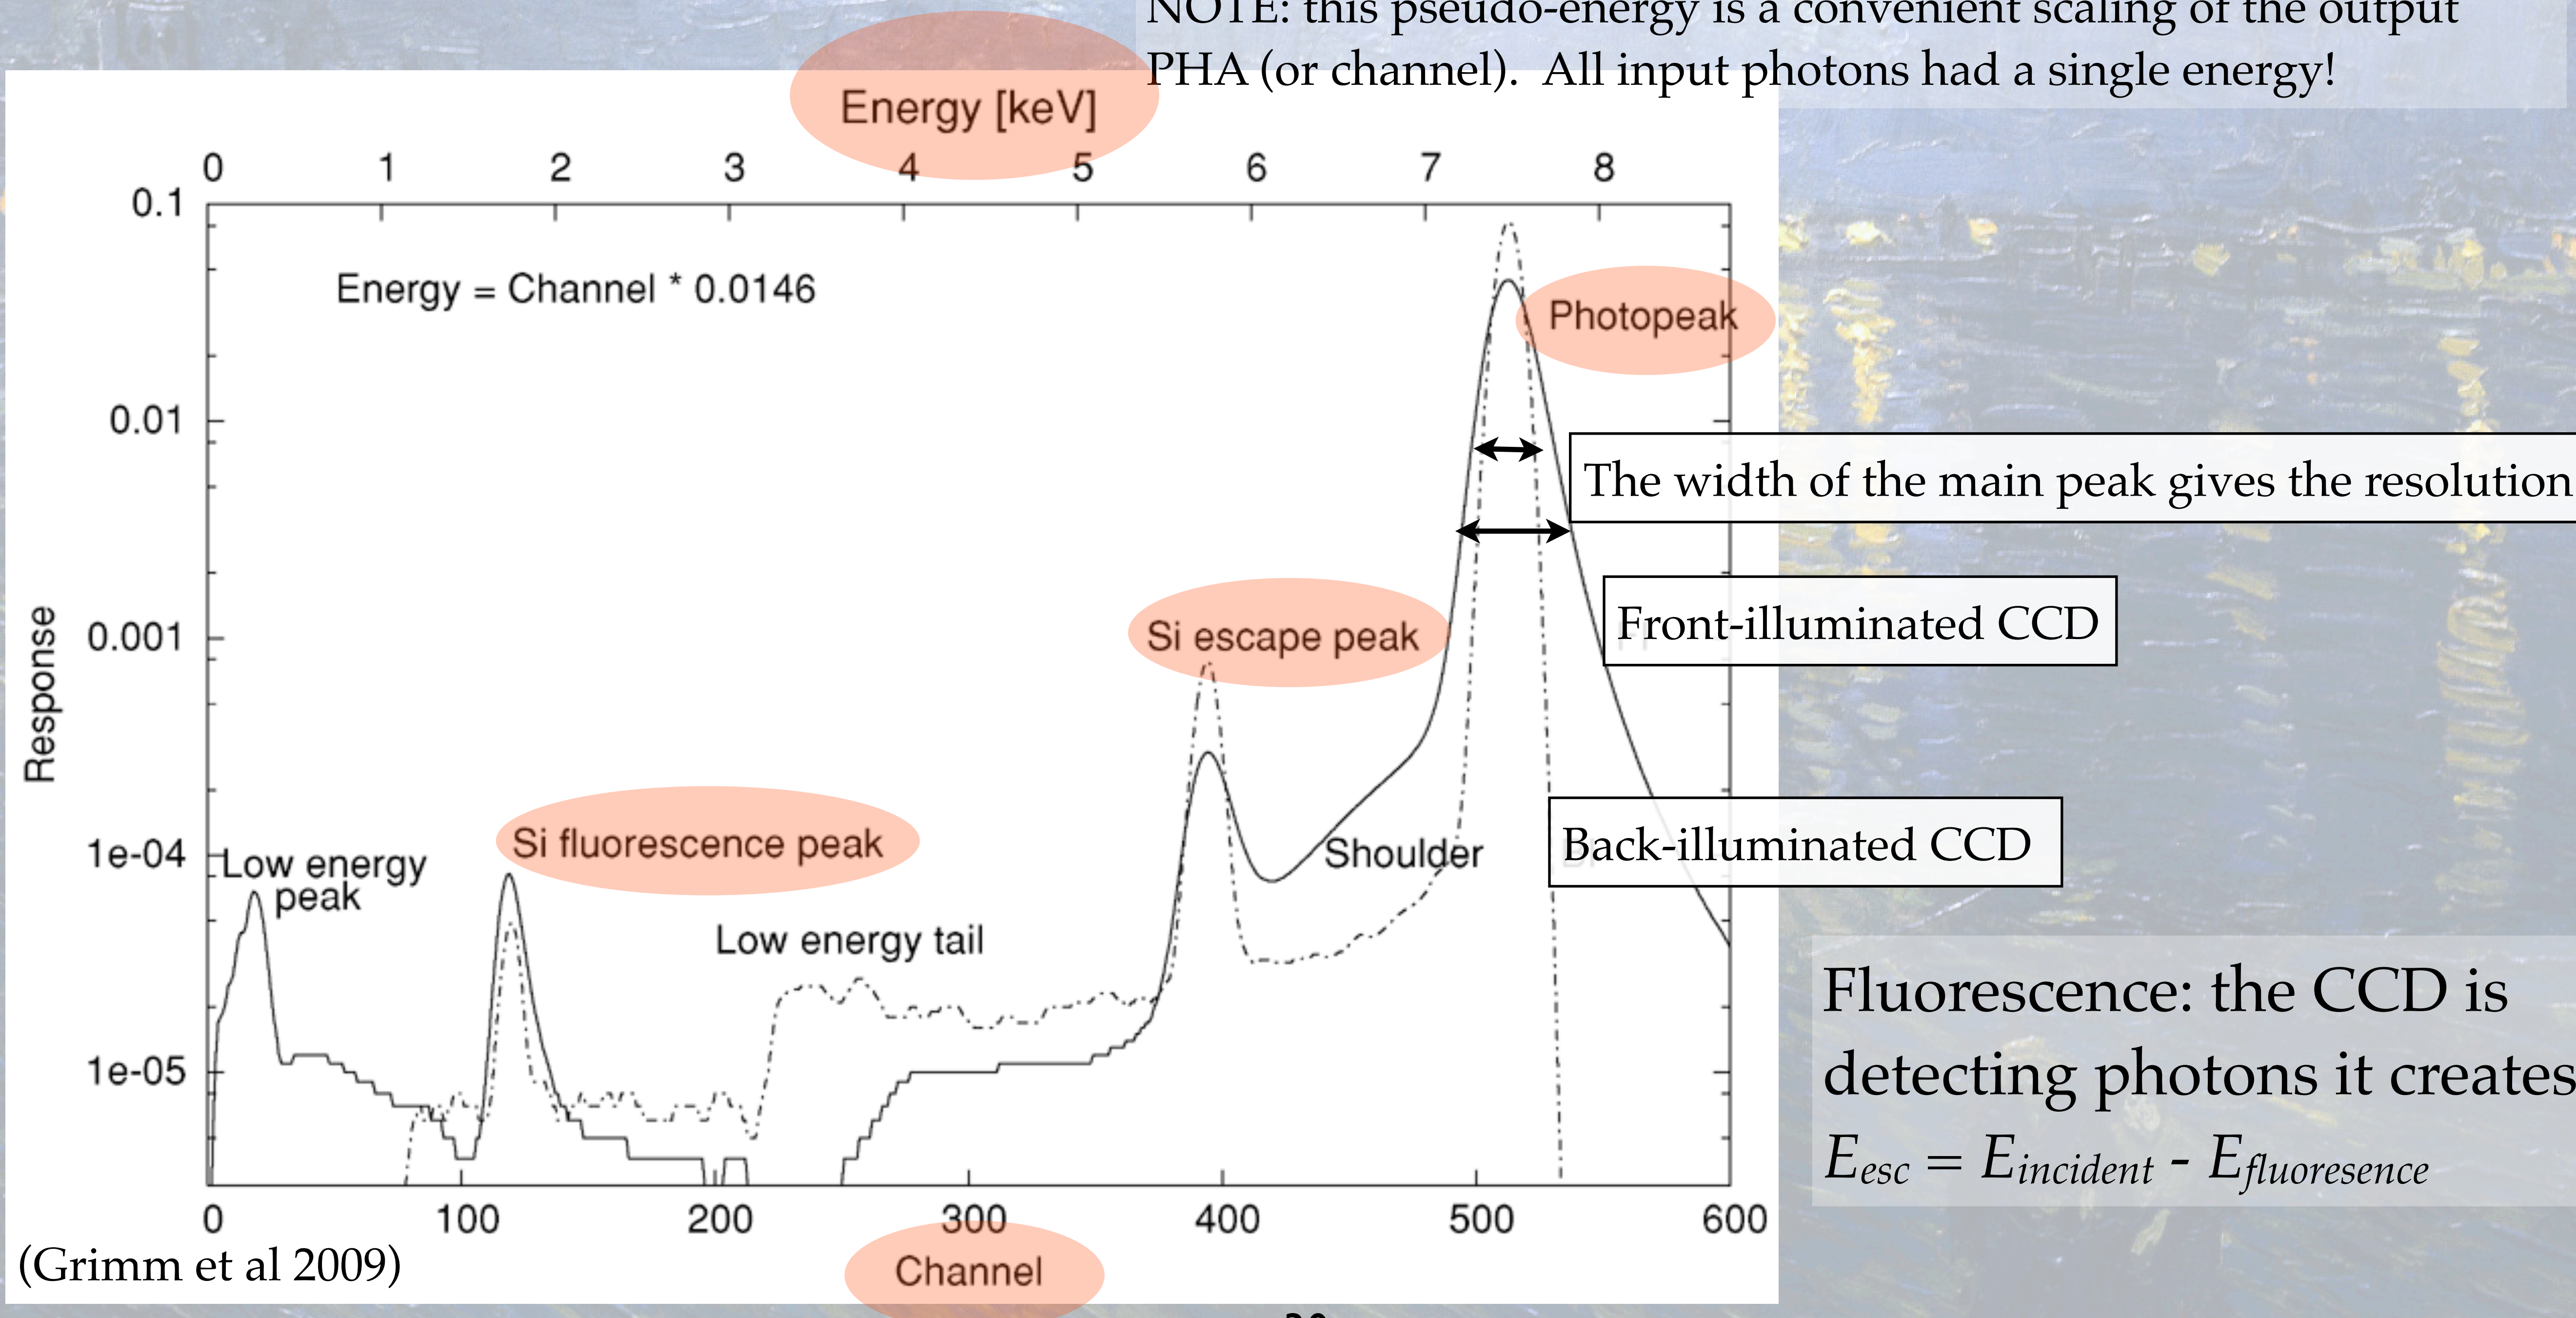

(Grimm et al 2009)

Fluorescence: the CCD is detecting photons it creates:

$E_{esc} = E_{incident} - E_{fluoresence}$

RMF: distribution

Visualized as a matrix, we see the photopeak scales nearly linearly with input energy. It's width is determined by Poisson processes, and roughly scales as $E^{1/2}$.

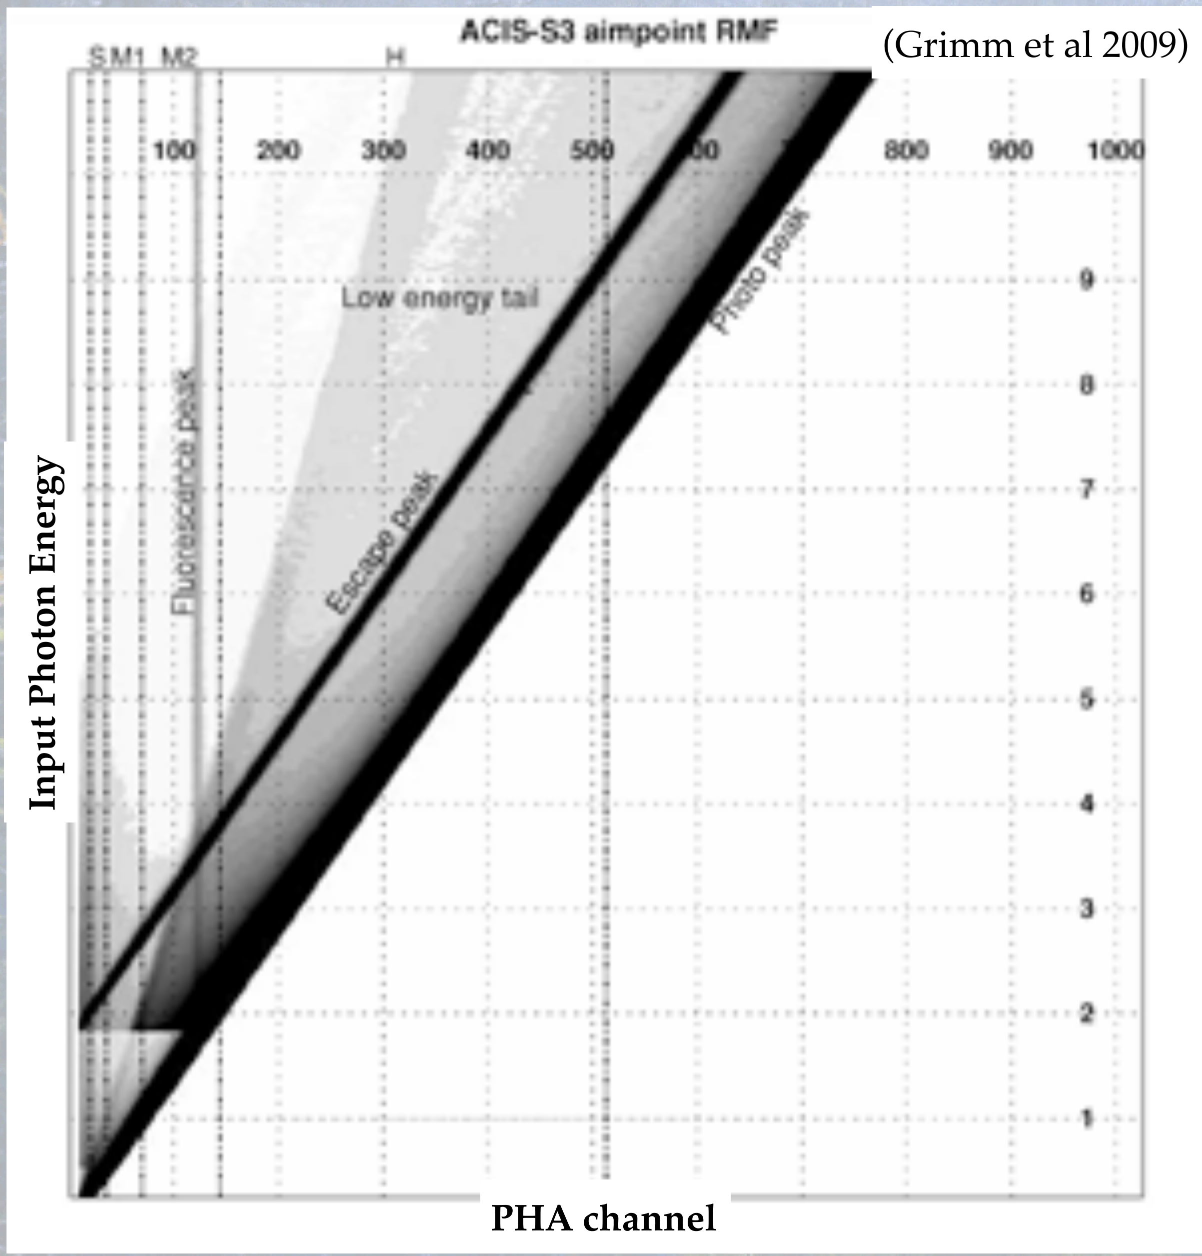

PI (pulse invariant) is a linearized *PHA,* to provide a common relation for all input energies.

Output "Energy" is a scaled *PI,* which maps the photopeak channel to the input photon energy; the relation between channel and energy is called the *gain*.

Gain generally varies with detector position (but is *usually* handled well by the ARF and RMF).

RMF: inspection, visualization

There are a number of options for visualizing RMFs. As FITS images which you can view with ds9:

- ciao: `rmfimg`
- ftools: `dmprmf`

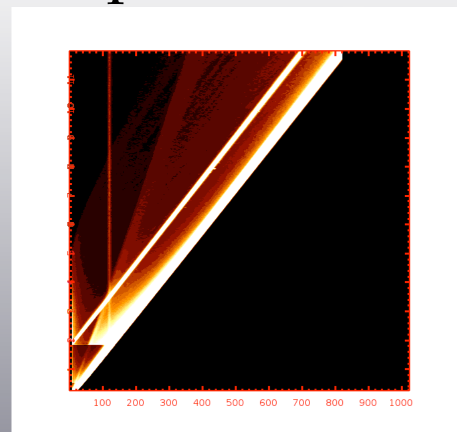

For distributions, use a spectral modeling system, such as ISIS, Sherpa, or Xspec.

Example, in ISIS: load an RMF, define a delta-function model, evaluate model counts, and plot:

```
isis> load_rmf( "aciss_aimpt_cy13.rmf" );     % load the RMF
isis> assign_rmf(1,1);                        % assign to new (fake data ) index

isis> fit_fun( "delta(1)");                   % use a delta function model
isis> set_par( "delta(1).norm", 1 );          % some large flux
isis> set_par( "delta(1).lambda", _A(2) );    % wavelength for 2 keV

isis> eval_counts;                            % forward fold the model

isis> xlog; xrange(0.1,10); ylog; yrange;     % setup plot

isis> plot_model_counts( 1 ) ;                % plot the convolved model

isis> set_par( "delta(1).lambda", _A(4) );    % do it again, for comparison
isis> eval_counts;
isis> oplot_model_counts( 1 ) ;
```

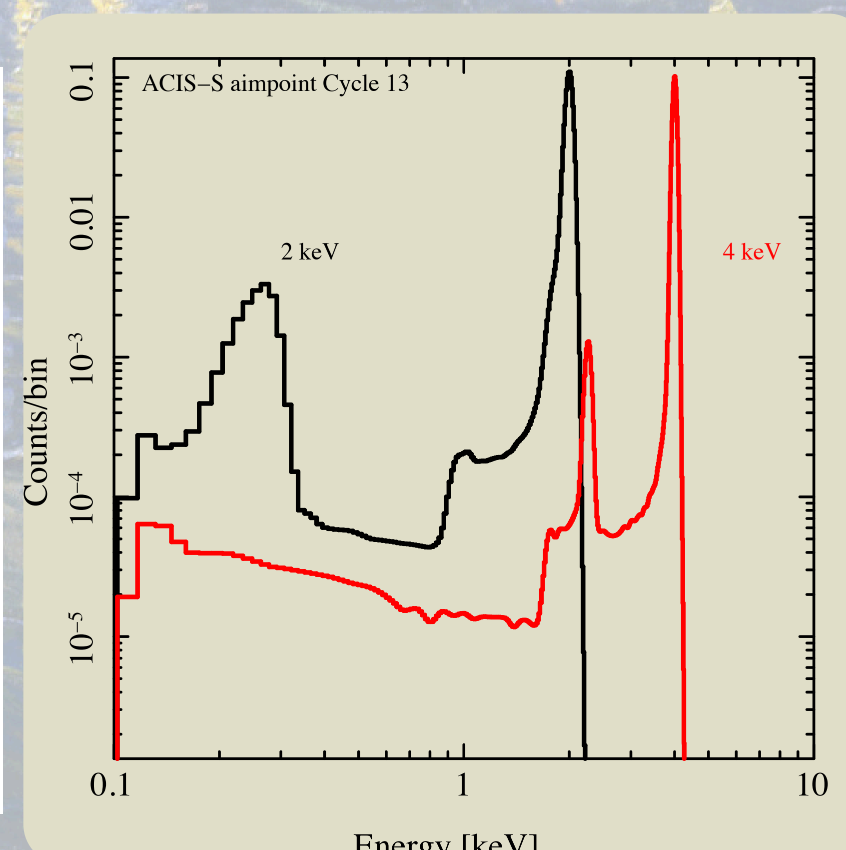

RMF: the Gratings Case

For spectra obtained with a diffraction grating (Chandra HETG or LETG; XMM/Newton RGS) the RMF defines the Line Spread Function (*LSF*). This is a combination of the imaging PSF and the grating diffraction profile. The *LSF* is also extraction-region dependent, via the cross-dispersion selection window which truncates the wings of the *LSF* in that direction.

Looking at a portion of the HETG/ACIS-S spectrum of a mildly extended source makes the spatial-spectral nature clear ...

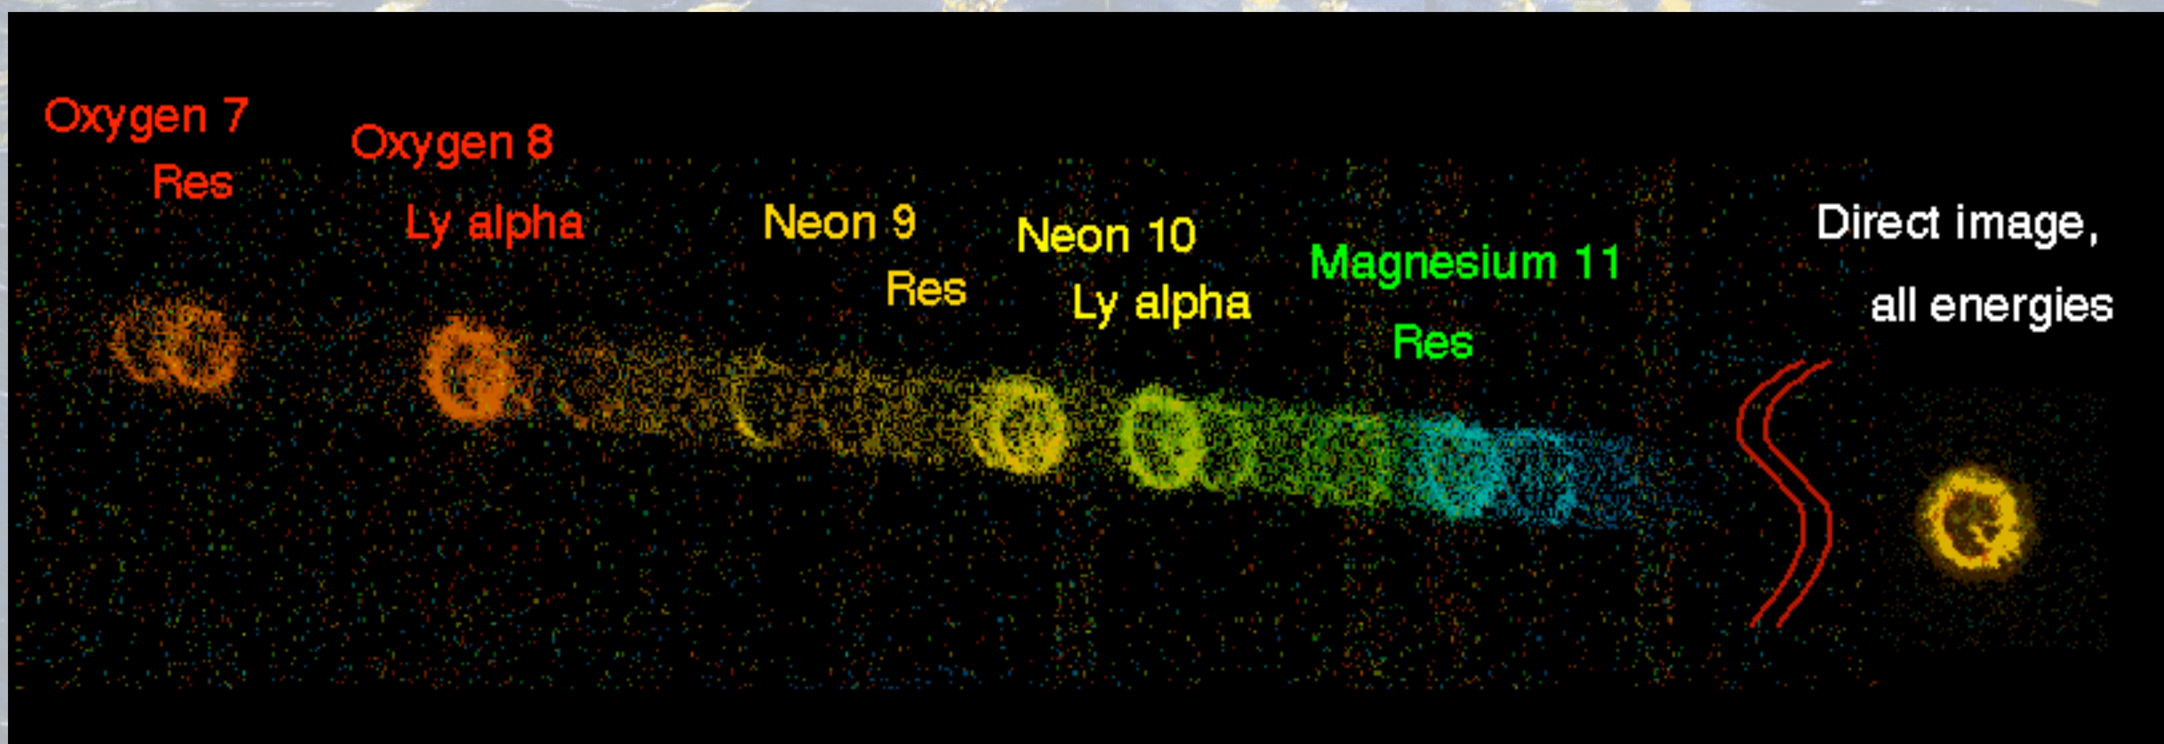

RMF: the Gratings Case, in detail

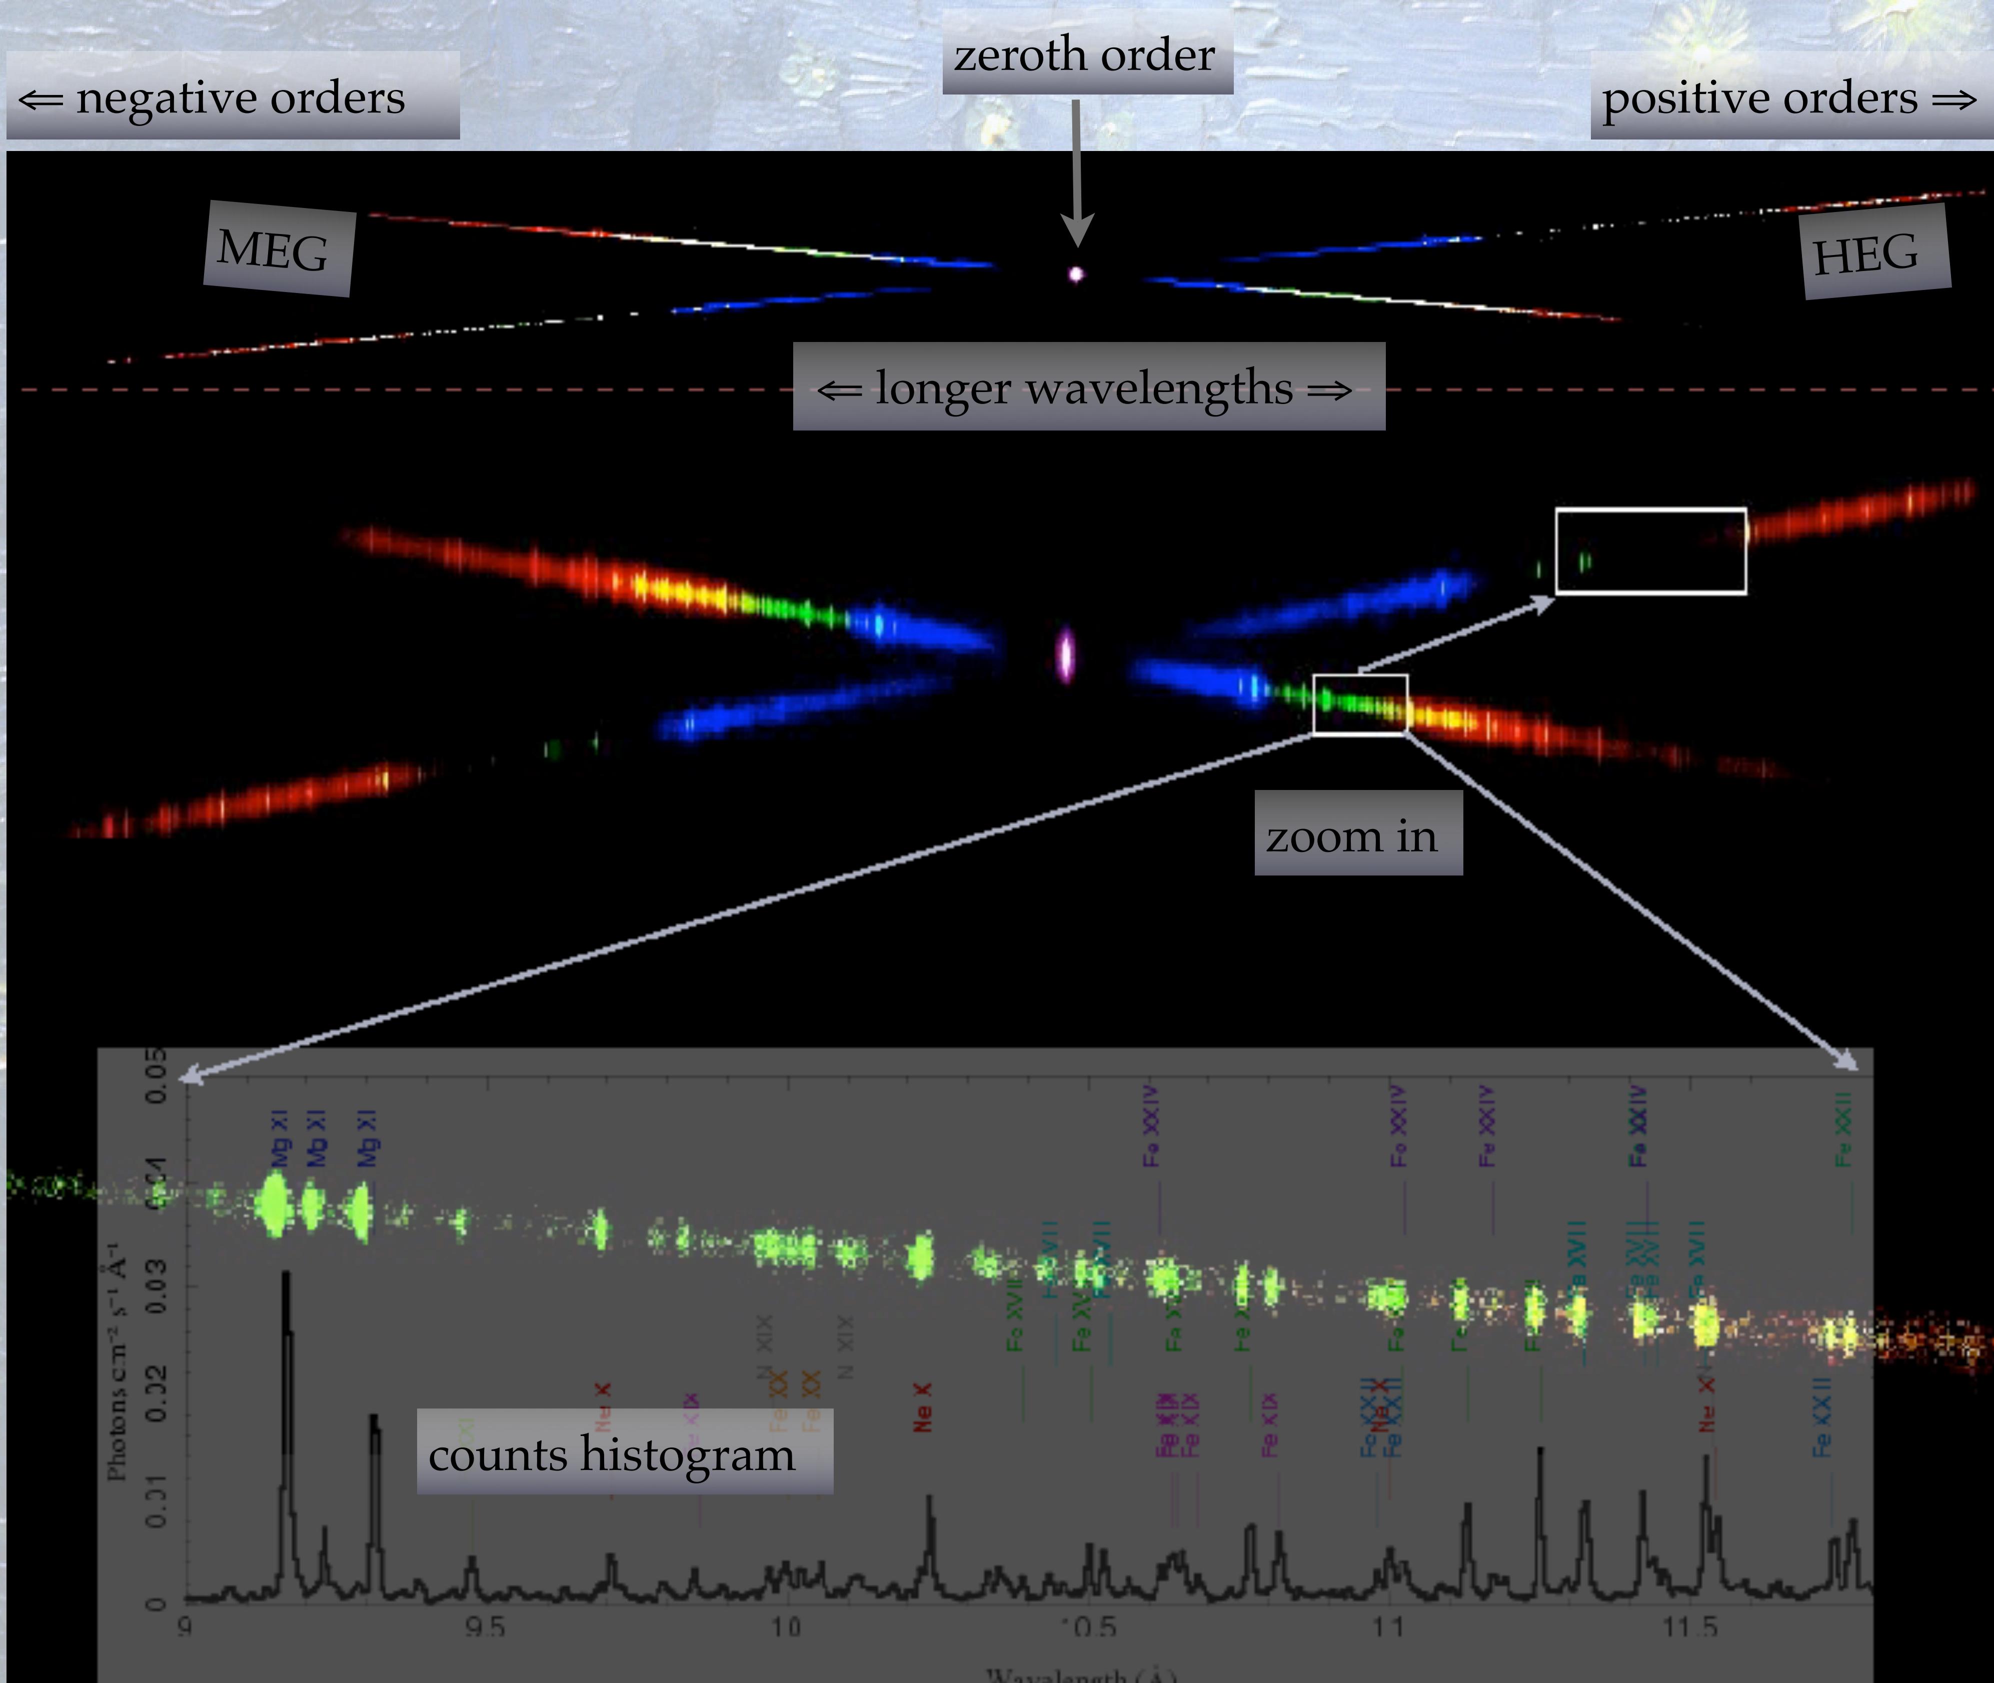

HETG spectrum, true aspect ratio

Stretched vertically

The zeroth order is a delta function and so represents the imaging PSF (for HETG).

The emission lines are also delta functions (at this resolution, for this star), so their shape is the *LSF*.

Aside: the intrinsic CCD resolution is used to sort overlapping diffracted orders: $m\lambda = P\sin(\theta)$.

RMF: the Gratings Case, in detail

Grating RMFs (and LSF) depend on the extraction width in the cross-dispersion direction. The effective area (ARF) also depends upon how much of the wings are clipped.

Extreme case: LETG/HRC-S: the spectrum becomes very wide at long wavelengths, requiring a “bow-tie” extraction region.

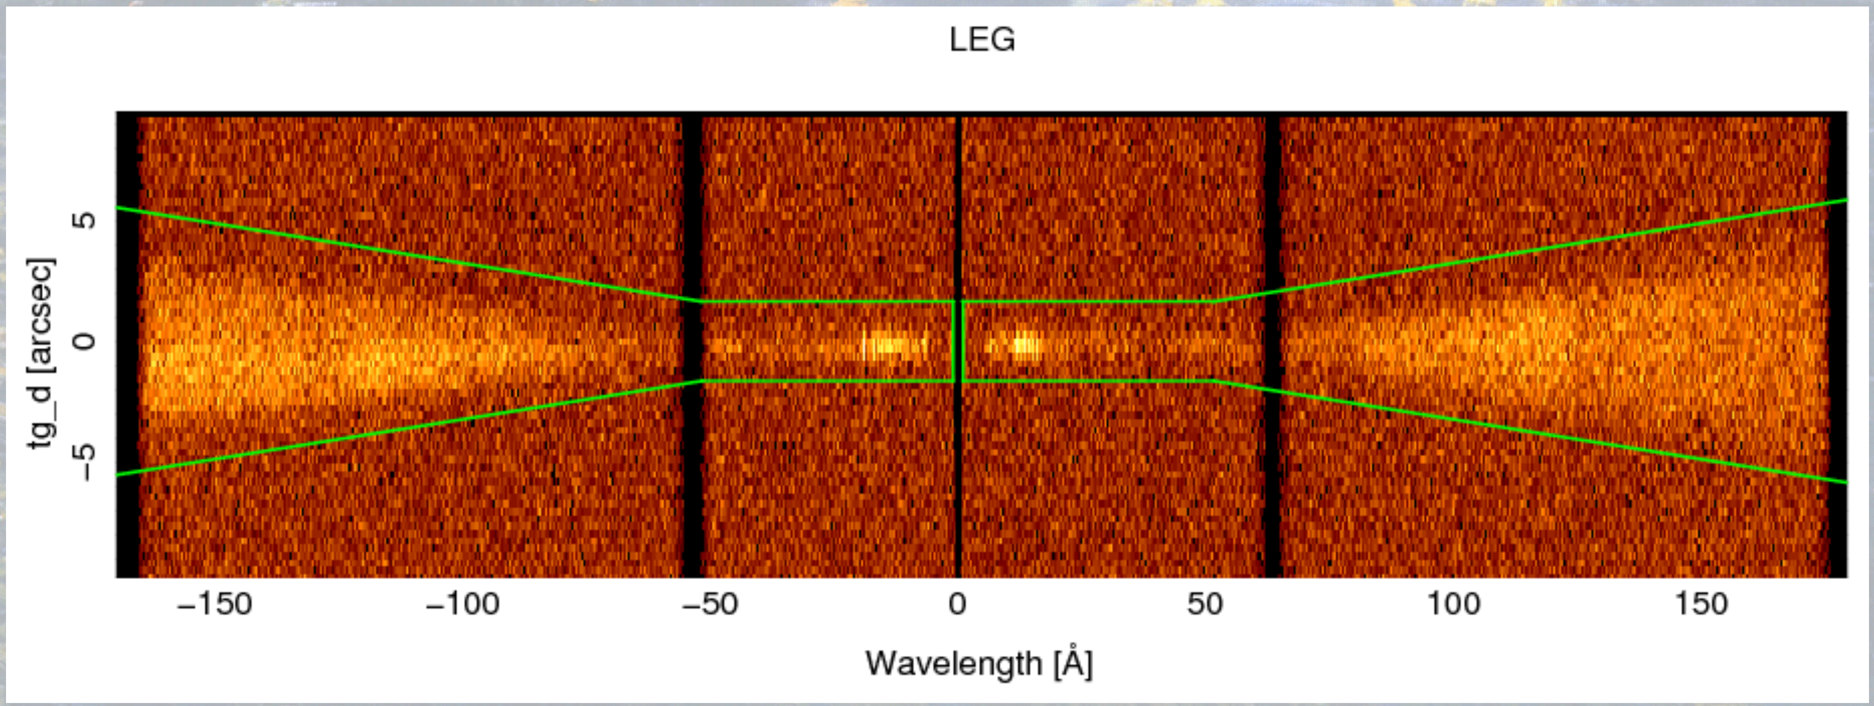

RMF (continued)

RMF computation is very instrument-specific. But once you have one, it is generic. An RMF is generally constructed in conjunction with an ARF, so grids match.

Since ARF and RMF pairs are almost always used together, they are sometimes multiplied together into the “response” file, or “RSP” file.

It is sometimes necessary to apply multiple responses to a single counts spectrum. Examples are LETG/HRC-S unresolved orders (the counts is a sum of orders, each with unique ARF*RMF), or overlapping sources.

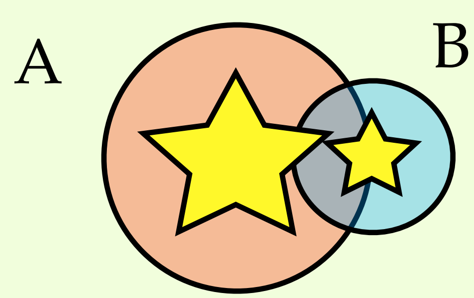

$$C_A = S_A R_A + S_B R_B f_B(A)$$

$$C_B = S_A R_A f_A(B) + S_B R_B$$

Overlapping sources: you can extract counts from regions A and B, but each contains a bit of the other source. This can be expressed through multiple response assignments to each region’s counts. C_x are counts in region x; S_x is the source model for region x. R_x is the “RSP” for region x, and $f_x(y)$ is the fraction of source x in region y. Two equations, two unknowns: solve for S_A and S_B*.

* See the ISIS help on “`combine_datasets`” for an implementation.

Spectral Modeling & Fitting

Given corresponding PHA, ARF, and RMF files we are nearly ready to solve for the source model spectrum. This is done by *forward-folding*. Due to the form of the RMF, the basic integral cannot be inverted directly. (It is sometimes OK for grating data - at least for visualization once a model has been obtained.)

There are two terms in the spectral counts equation we haven't discussed: those of the background, *B(h)*, and source model, *S(E)*.

$$C(h) = T \int_0^\infty \sum_i R_i(E,h)\, A_i(E)\, S_i(E)\, dE + B(h)$$

PHA Exposure RMF ARF source model "background"

Spectral Modeling & Fitting: the Background Term

The background term, *B(h)*, depends strongly on the instrument, and its importance also depends on your source ("bright" is relative to background).

Backgrounds can come from:

- diffuse or unresolved cosmic sources
- confusing sources
- detector electronics
- cosmic rays
- local effects (scattered light, filter leaks)

While some components are affected by the system effective area, *others are not.* Hence, it is generally sufficient (and sometimes necessary) to prepare a background as a counts histogram (PHA) term which includes whatever counts signal *deemed* to be background!

If you can reliably distinguish external and internal sources, then the external components could be included in *S(E)*.

Spectral Modeling & Fitting: More about the Background Term

Providing the background term:

- current observation (e.g., counts spectrum extracted background region in the field, or adjacent to the high-resolution spectrum); added to *B(h)*
- some other observation (e.g., Chandra background events file collected from many observations); added to *B(h)*
- modeled, possibly with free parameters to be fit with *S(E)*; included in *S(E)*, or added to *B(h)* (or both).

Some issues (not covered in detail here):

- Don't subtract a background! (Add to the predicted model counts)
- Backgrounds need to be scaled by exposure and extraction area to match the source exposure and region.
- Backgrounds can alter the statistics.

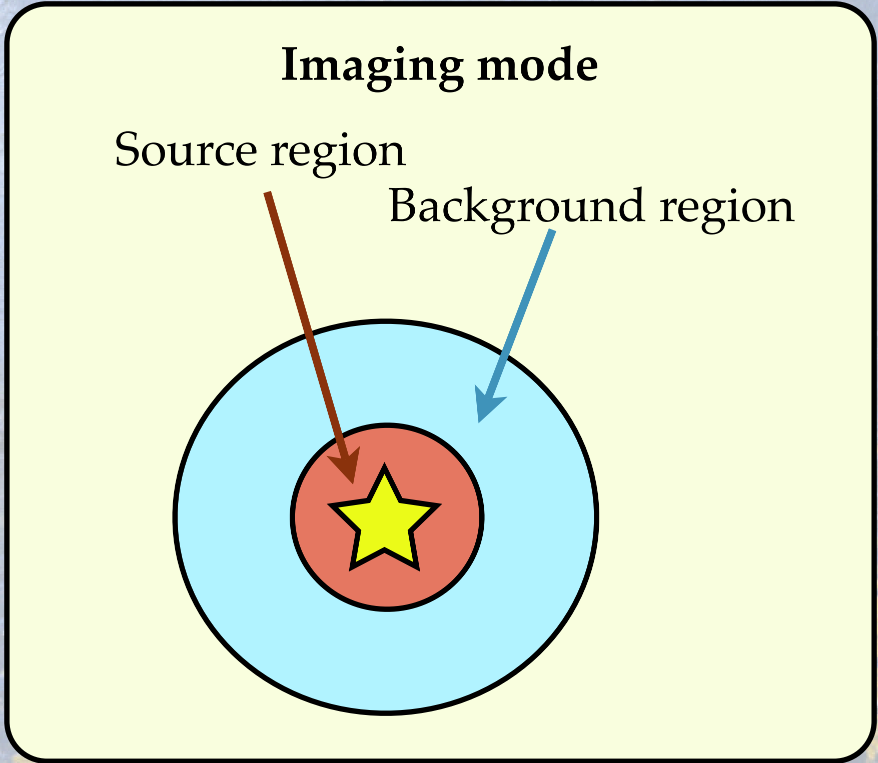

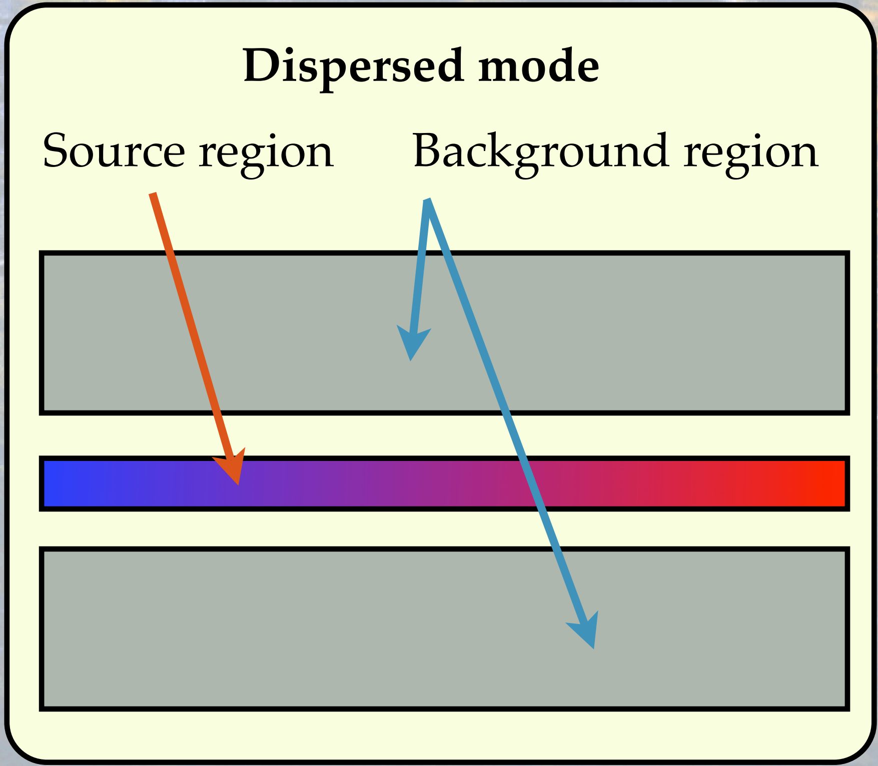

Spectral Modeling & Fitting: the Source Term

S(E) is your source spectrum, discretized on some fine energy grid (usually finer than *h*), having units of [photons cm^{-2} s^{-1} bin^{-1}]*. (Note the "per bin"; if you want "per keV" or "per Å", divide by the bin-width; if you want ergs, multiply by ergs/photon.)

Models may be simple parametric functions (Gaussian, power law), simple physical models (black body), or complicated physical models involving extensive code and tables (`warmabs`, `apec`) or be combinations thereof, linear (additive, multiplicative) or non-linear (convolutions). Models might have one parameter, or 200.

Models are subject to their own "calibration, *calibration*, & <u>*calibration*</u>" issues. They have systematic uncertainties from underlying atomic data (20%-60%), physical assumptions, numerical approximations, and bugs (*trust no one*).

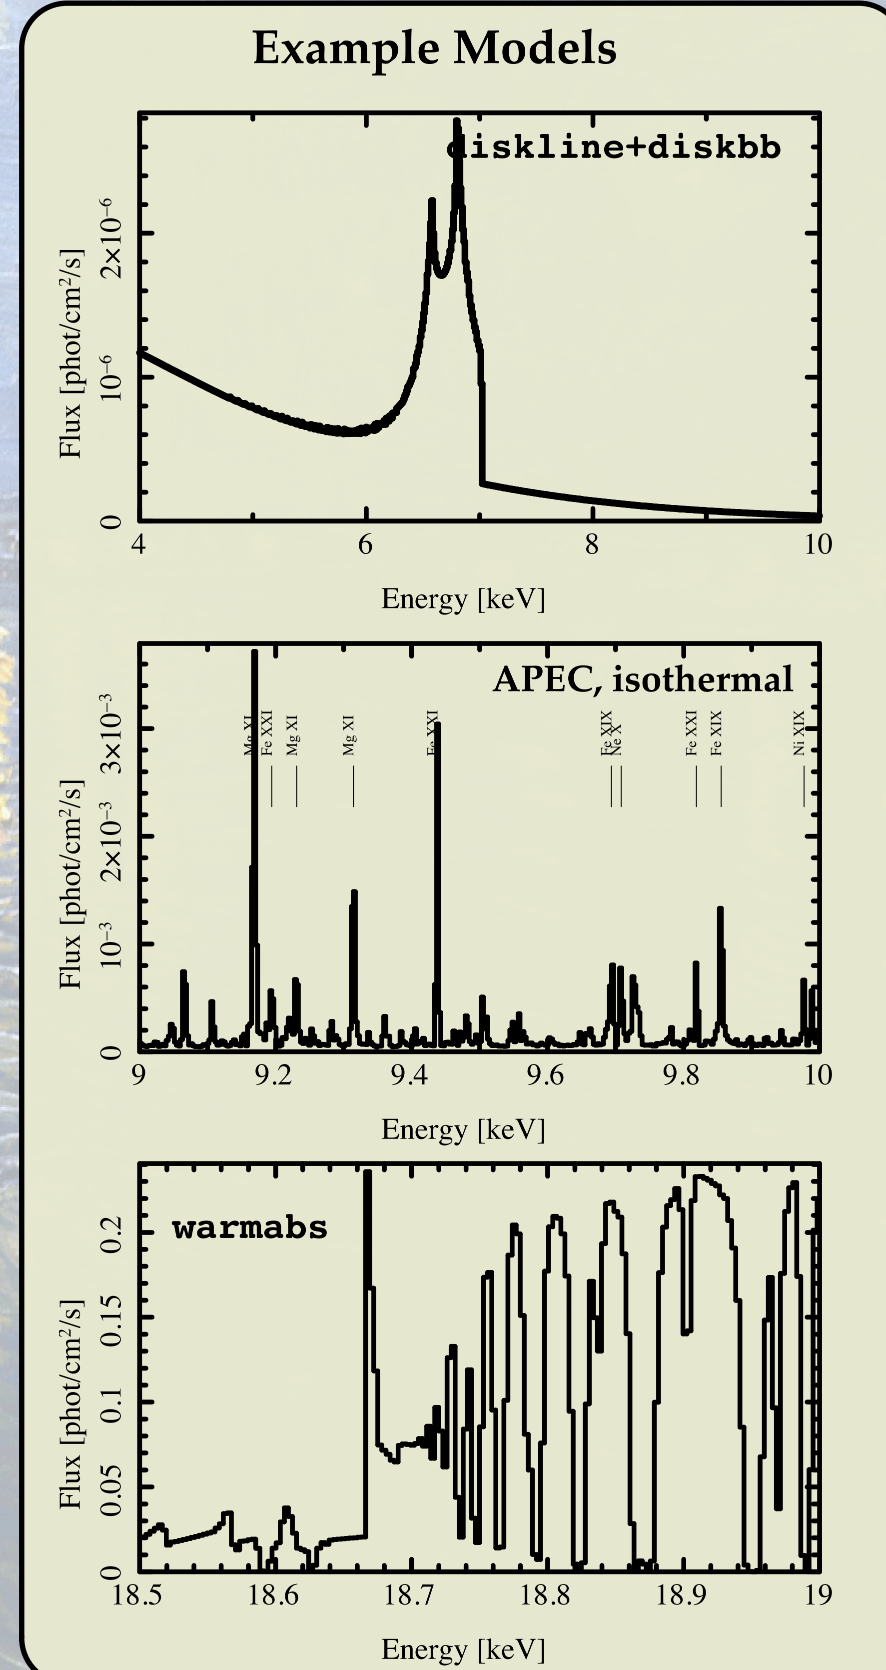

* Remember, photons are *not* counts.

Spectral Modeling & Fitting: the Source Term

I repeat: ***trust no one.***
The models are the subject of your research. You need to be able to dissect them, validate against published work, extend them. You need to know their limits of applicability. Which means you need to be able to examine them in detail as a function of their parameters. It might mean *years* of work (e.g., a Ph.D. thesis).

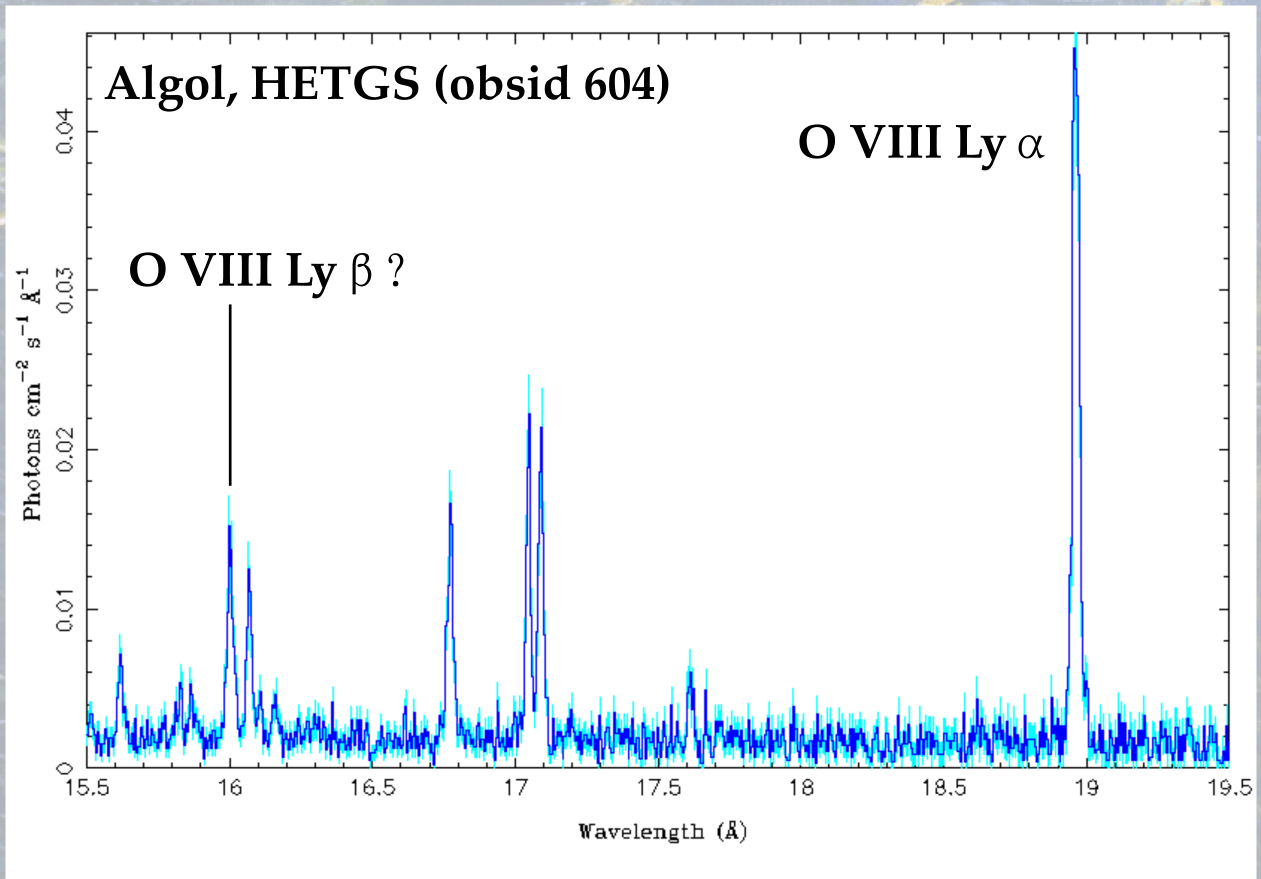

Example: getting inside the "black-box" of a thermal plasma model (or, what's in that feature at 16Å?)

Procedure:

- define a model
- define a wavelength grid
- evaluate the model
- plot it
- plot/page line identifications
- plot the emissivity vs temperature

(see next page for a detailed example using ISIS and APED.)

41

Spectral Modeling: Collisional Plasma, detailed example of "opening the black-box"

```
isis> plasma(aped);
isis> create_aped_fun( "Aped_1T", default_plasma_state);
isis> fit_fun( "Aped_1T(1)" );
isis> set_par( "Aped_1T(1).temperature", 10^6.8);

isis> (wlo, whi) = linear_grid( 15, 20.0, 16384 );
isis> y = eval_fun( wlo, whi ) ;

isis> xrange( 15.95, 16.05 ) ;  yrange(0);
isis> hplot( wlo, whi, y ) ;
isis> plot_group( brightest( 10, where( wl( 15.95, 16.05) ) ) );
```

```
isis> page_group( brightest( 5, where( wl( 15.98, 16.02 ) ) ) );
```

```
isis> t = 10^[ 6.0 : 7.5 : 0.1];
isis> e1 = line_em( where(trans( O, 8, [6,7], 1) ), t );
isis> e2 = line_em( where( trans( Fe, 18, 5, 1) ), t );

isis> xlog; xrange;
isis> xlabel("Temperature"); ylabel("Emissivity");
isis> plot( t, e1, 2 ) ;   oplot( t, e2, 4 );
```

```
isis> page_group( brightest(5, where( wl(15.98, 16.02))));
#  index        ion    lambda    F (ph/cm^2/s)  A(s^-1)  upper lower   label
 421194 *    Fe XIX    15.990    2.887e-05   2.510e+10    40     6 2p3 - 2s1
 390247 *   Fe XVII    15.996    3.947e-03   7.543e+08    14     1 2p5 - 2p6
 404081 *  Fe XVIII    16.004    2.522e-02   1.580e+12     5     1 2p4 - 2p5
  50041 *    O VIII    16.006    1.278e-02   6.829e+11     7     1 3p1 - 1s1
  50034 *    O VIII    16.007    7.718e-03   1.367e+12     6     1 3p1 - 1s1
```

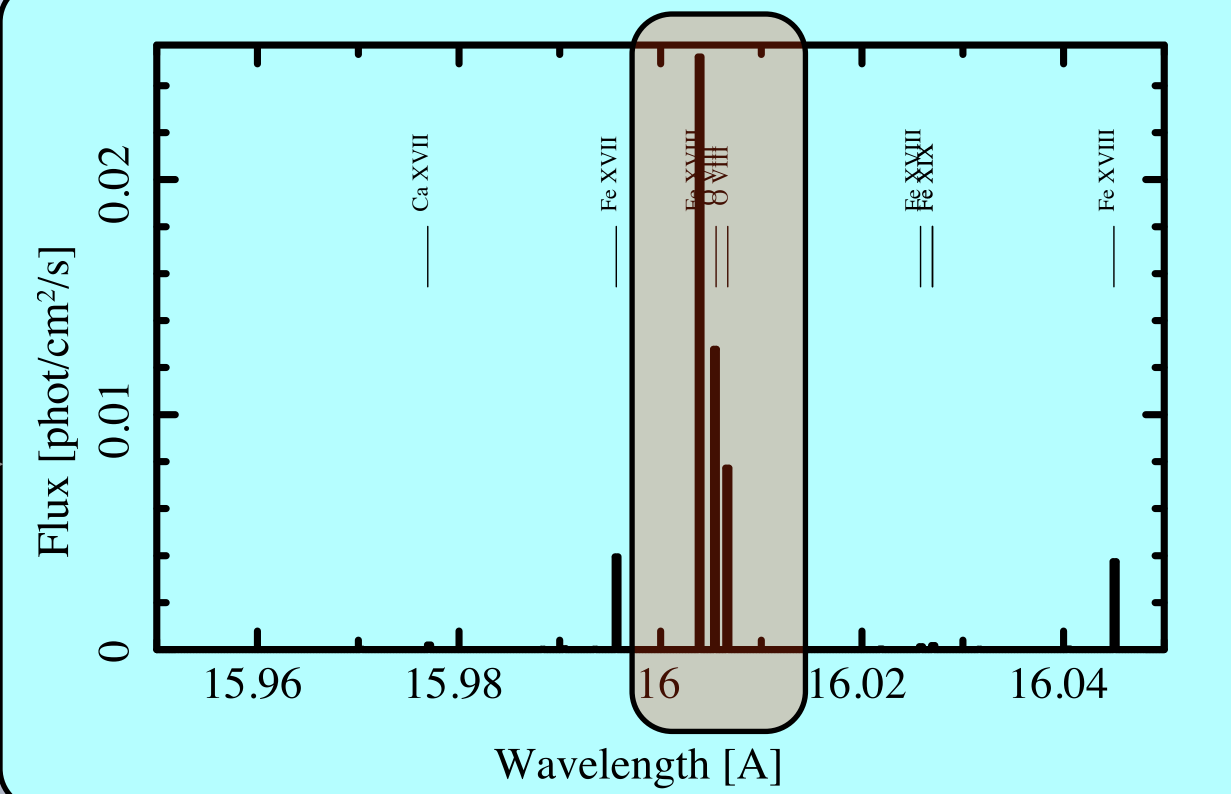

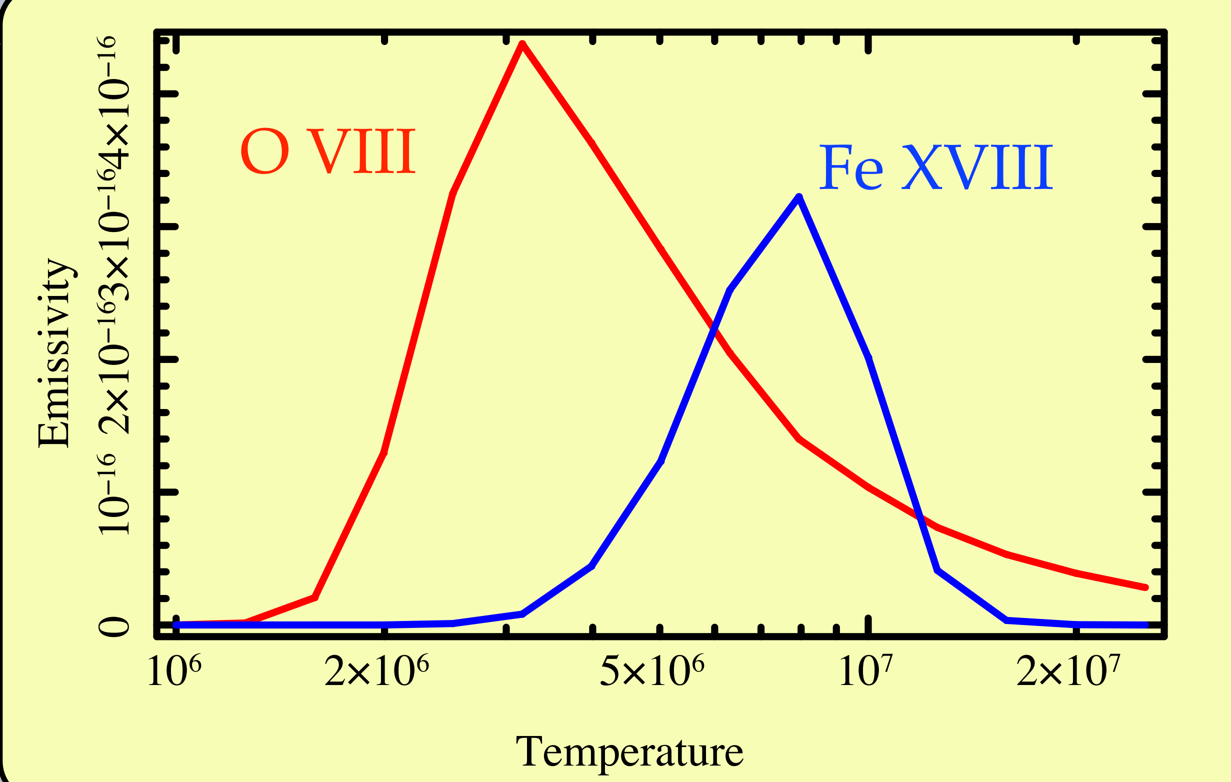

The 3 lines highlighted all fall within a single MEG bin of 0.005A. If you naively intended to measure a temperature from the ratio of O VIII Ly α to Ly β, you would be wrong.

42

"A more intimate analysis..matured my conjecture into full conviction."

Spectral Fitting

This should look familiar by now:

$$C(h) = T \int_0^\infty \sum_i R_i(E,h)\, A_i(E)\, S_i(E)\, dE + B(h)$$

It can be discretized and expressed as a matrix equation. Simplifying somewhat (by combining *R* and *A* and assuming that *i*=1), we could write in matrix form as:

$$\mathbf{c} = T\,\mathbb{R}\mathbf{s} + \mathbf{b}$$

So why not just invert it to solve for **s**?

$$\mathbf{s} = \mathbb{R}^{-1}(\mathbf{c} - \mathbf{b})/T$$

Because: there is *noise* in (**c-b**), and because the form of **R** (generally a rather broad redistribution) does not mathematically permit a *unique* inversion -- it is very unstable to small perturbations (noise), even *if* there is a unique solution (look up Fredholm equation, Type I).

Spectral Fitting: "unfolding"?

A reasonable way to define an "unfolded" spectrum is to divide the observed counts (minus background) by the counts per unit flux (i.e., by model counts for $S(E) = 1$):

$$\mathcal{F}(h) = \frac{C(h) - B(h)}{T \int A(E) R(E,h) dE}$$

This is risky - spurious features can appear, particularly where there are strong gradients in A with redistribution by R. It is good for a qualitative view only, but *not for quantitative analysis*.

(Note: "unfolded" spectra often appear in the literature, but most often ***NOT*** with this definition; beware - they have an even more dubious relation to truth.)

Spectral Fitting: a bad unfolding

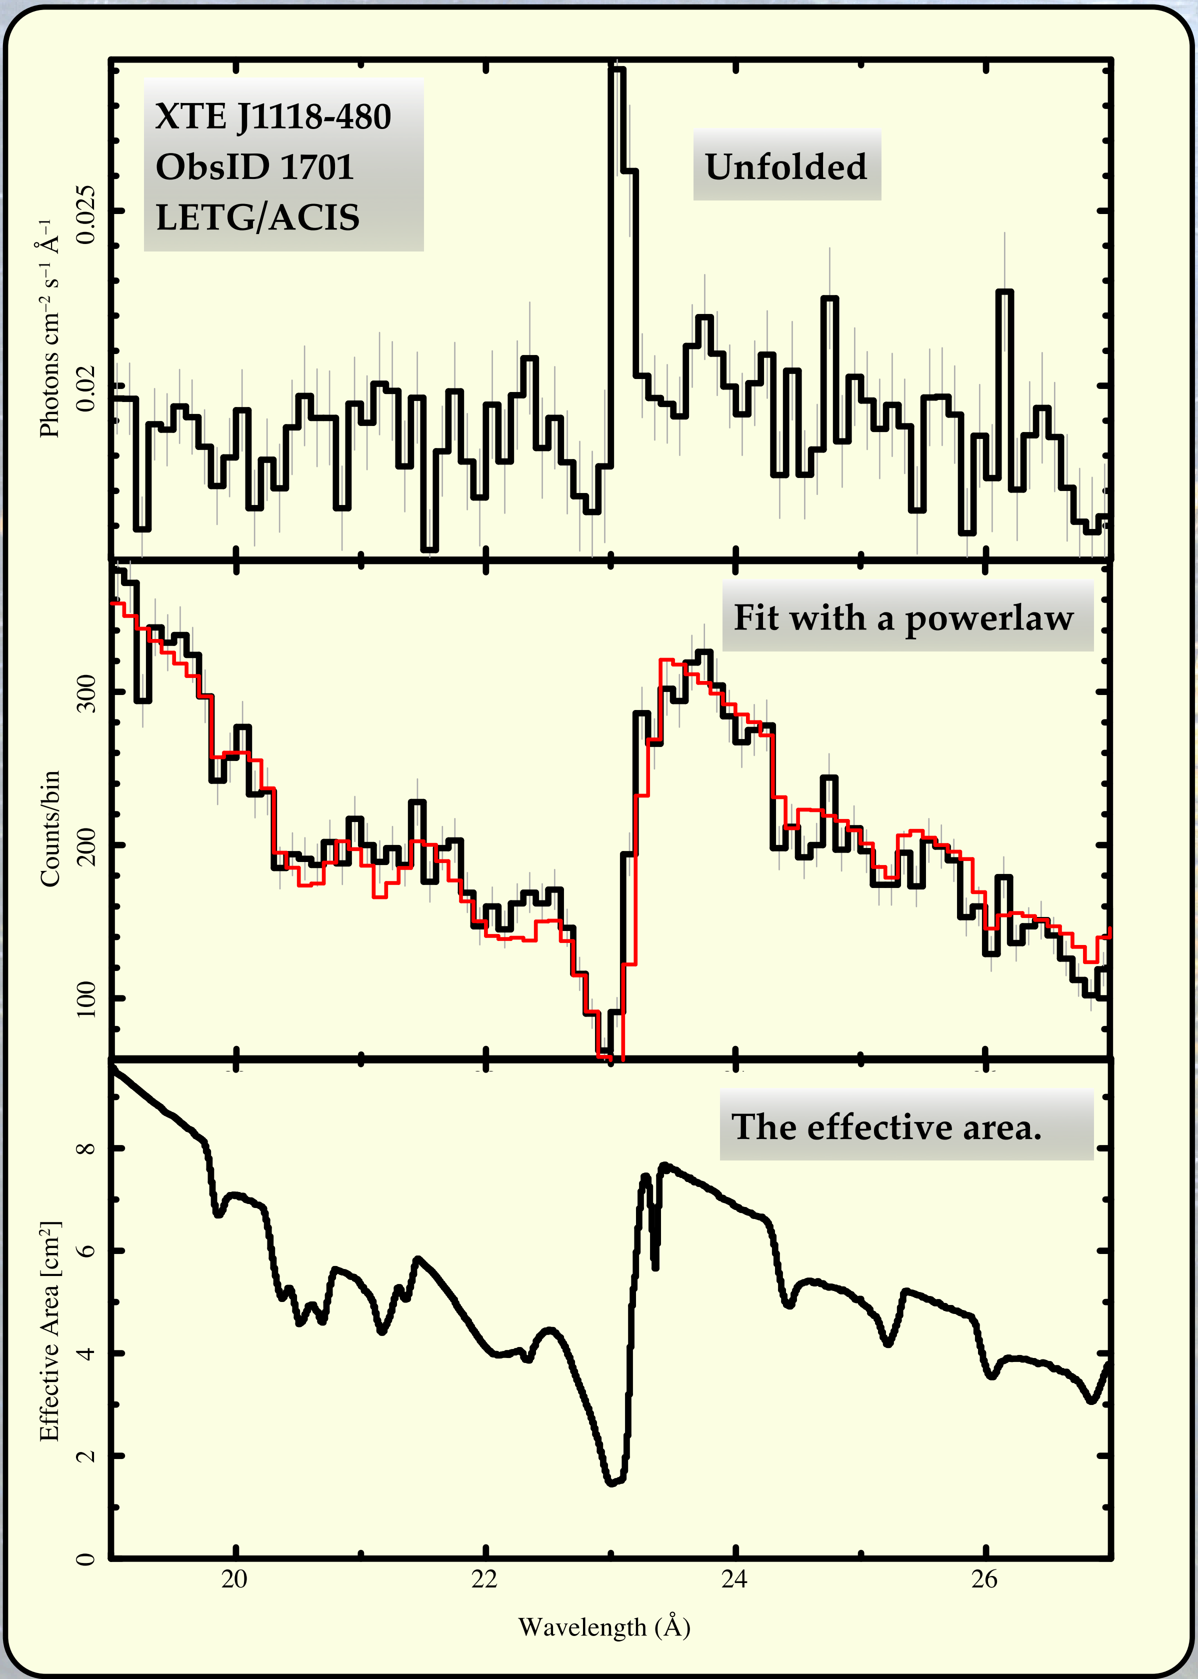

The unfolded grating spectrum of the low-mass X-ray binary, XTE J1118-480 (a.k.a. KV UMa) seems to show a strong emission line at 23Å, not a location known for any prominent feature.

In fact, a fit with a featureless power-law does pretty well.

The spurious “line” is due to the sharp edge in the effective area. “Unfolding” is sometimes a poor inversion technique.

Spectral Fitting: the forward-folding technique

When we forward-fold, we make an educated guess for the model, then minimize a statistic (typically χ^2) formed from the difference between the data and the model. An optimization routine iteratively varies parameters to search for a minimum in the statistic.

statistic

observed counts

model counts

+ background counts

$$\chi^2 = \sum_h \left(\frac{C(h) - \left(T \sum_i \int R_i(E,h) A_i(E) S(E;\mathbf{p})\, dE + B(h)\right)}{\sigma(h)} \right)^2$$

minimize this

uncertainties

model parameters

Note: the uncertainties, σ, may be derived from the statistics of both *C(h)* and *B(h)*.

Spectral Fitting: forward-folding

Forward-folding is also non-unique: this observed counts spectrum (duplicated in top and bottom) has been fit with two-temperature component thermal models. Panels (a) and (b) represent equally good fits, but the two components are quite different.

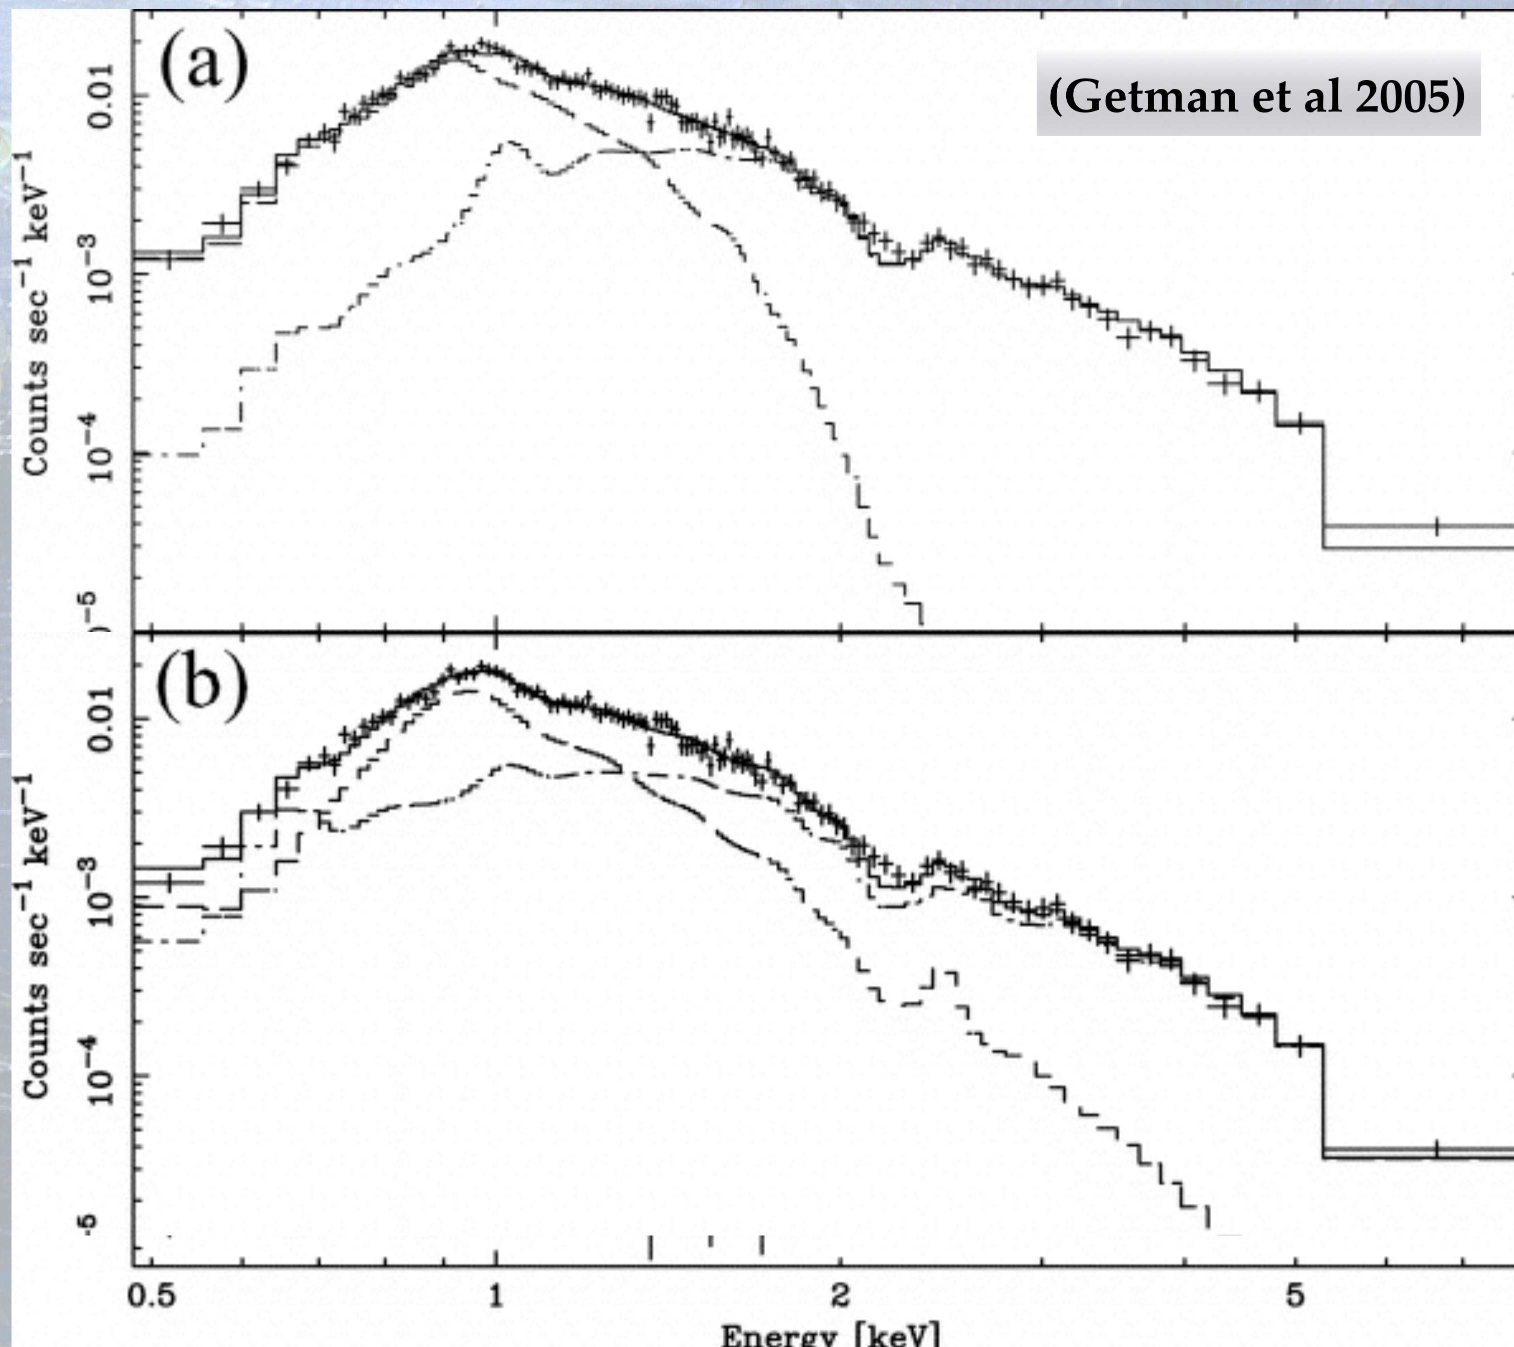

Spectral Fitting: being practical

Be practical, reasonable, and be careful:

1) Load your data and responses: make sure they match
2) Inspect your data; use redundancy, if possible (multiple orders, multiple observations)
3) Define a model
4) Set reasonable starting parameter values and ranges (*think - don't use the defaults!)*
5) Define grouping and regions of interest (and plot again)
6) Evaluate the model, and inspect it, relative to data; adjust your parameters (e.g., renormalize)
7) Choose a statistic and fit method (they are sometimes coupled). Some methods are faster than others. (examples: χ^2, Cash; Levenberg-Marquardt, subplex, Monte-Carlo)
8) Do a fit, plot data and residuals, review parameters and ranges.
9) If the fit model and residuals look reasonable (not just a low statistic, but sensible, physical, plausible, *regardless* of the absolute value of the statistic),
10) then compute 1D error-bars, or 2D confidence maps (remember: parameters can be strongly correlated)

You will need to iterate, e.g., steps 4-6, or 6-8. You should also repeat from step 3 with alternative models. You might be able to use different methods for fitting, and for error-bars (for either efficiency, or for accuracy). When in doubt, randomize, simulate, fit something with a known answer.

Timing Analysis

The primary object for timing analysis is usually the *light curve*: the count rate per time bin. Usually in a light curve table,:

- time
- counts
- error
- count rate
- exposure

Exposure is important: it incorporates

- GTI - Good Time Intervals
- Dead time (saturation)
- Any other exposure effects (e.g., dither of source to dead regions)

Also important, but generally implicit, is the instrument natural sampling interval, and any periodic instrumental signal (e.g., dither).

Light curves from background regions can also be very important for interpretation.

Example (from `tgcat.mit.edu` preview products for obsid 2523, V471 Tau):

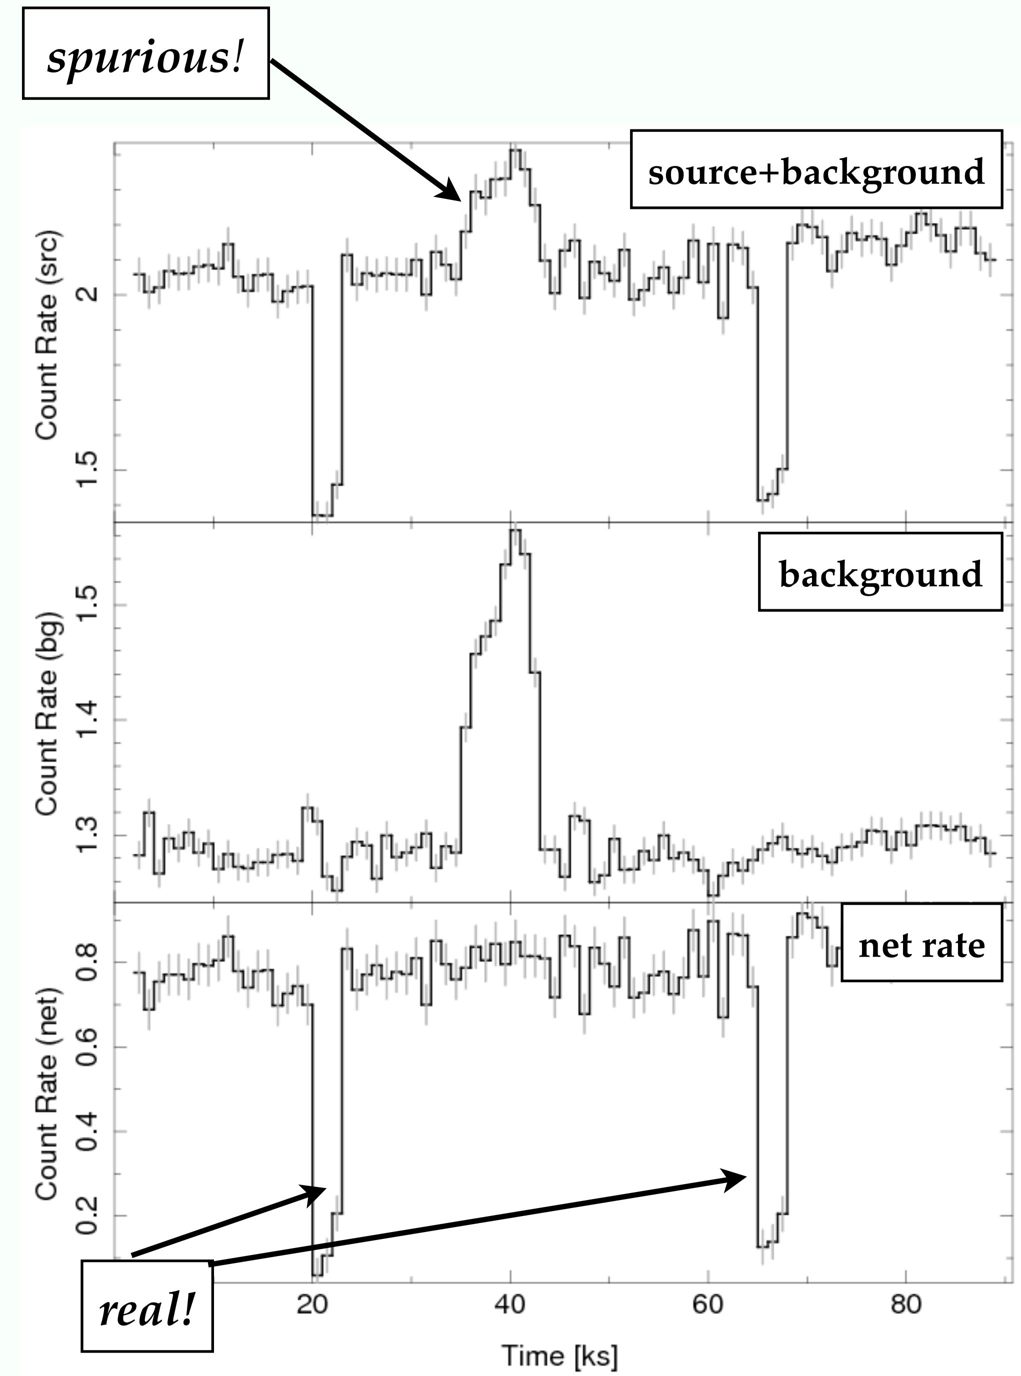

Timing Analysis

Other types of timing analysis involves are power spectra, epoch folding period folding, statistical analysis of variability (K-S test, Bayesian Blocks).

NGC 3783 - an erratic light curve

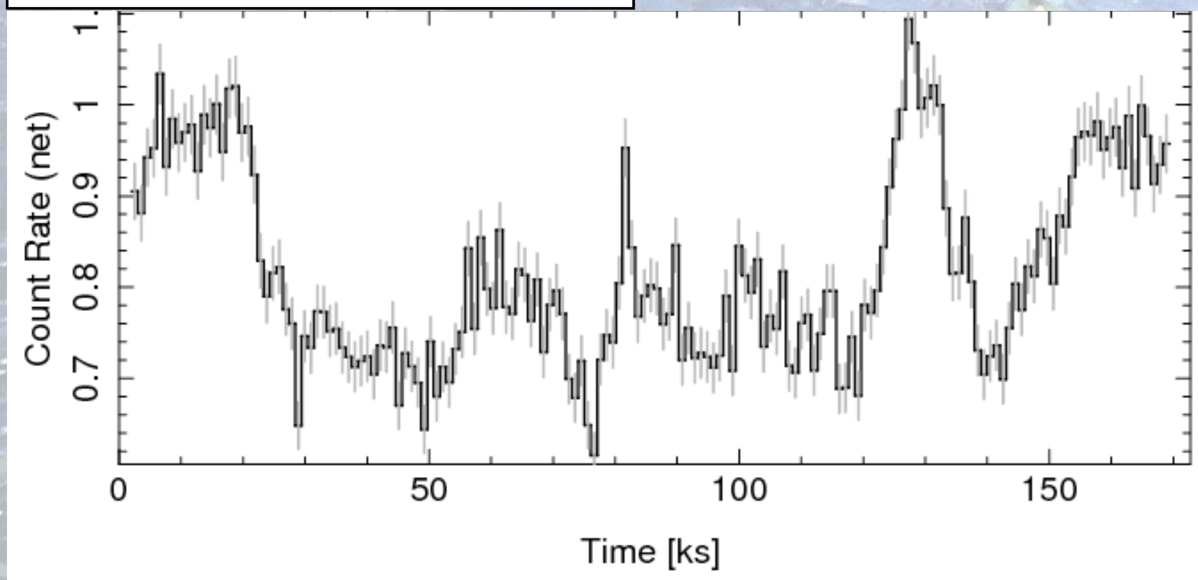

EX Hya - a multi-period curve

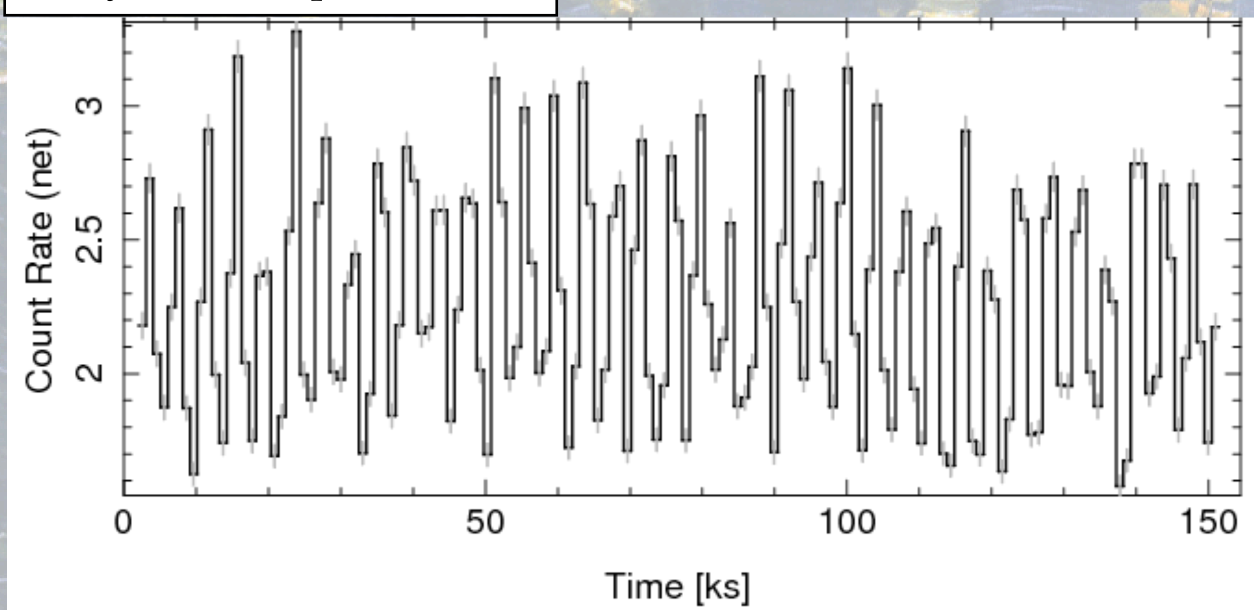

II Peg undergoes a stellar flare

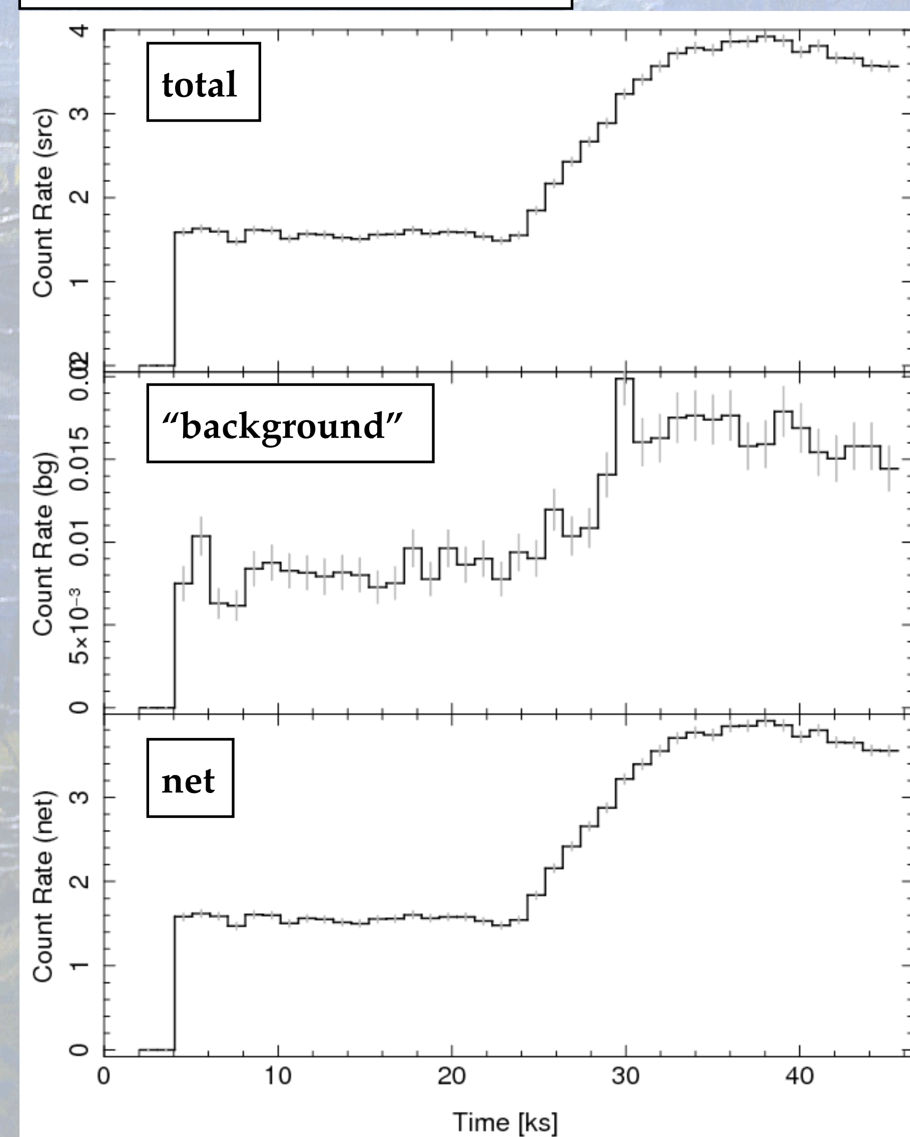

Note the trend in the "background". This is instrumental, due to the positional aliasing of ACIS events during the CCD frame readout time.

Timing Analysis

Reading and folding RXTE/ ASM data on Her X-1 (from M.Nowak's SITAR web page examples, space.mit.edu/cxc/analysis/SITAR/)

Light Curve

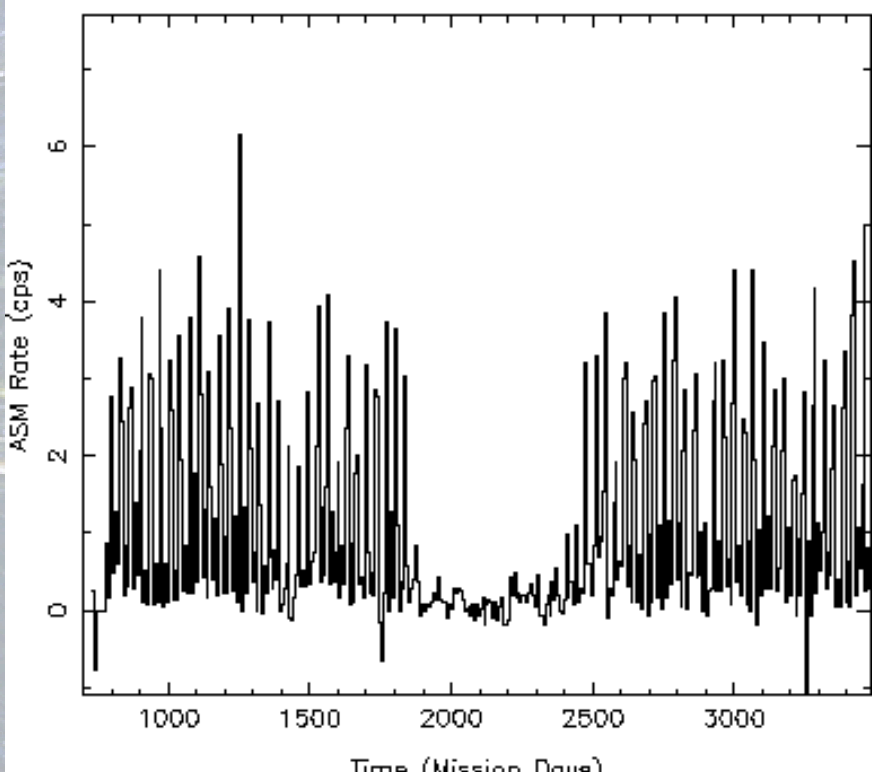

Period search via epoch-folding

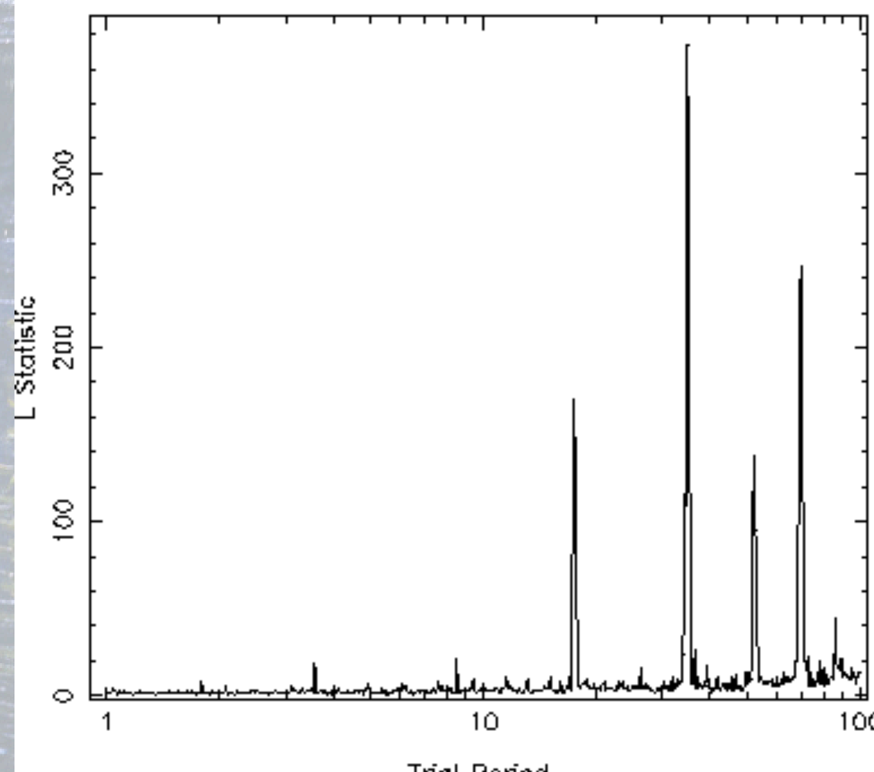

Period-folded on strongest peak

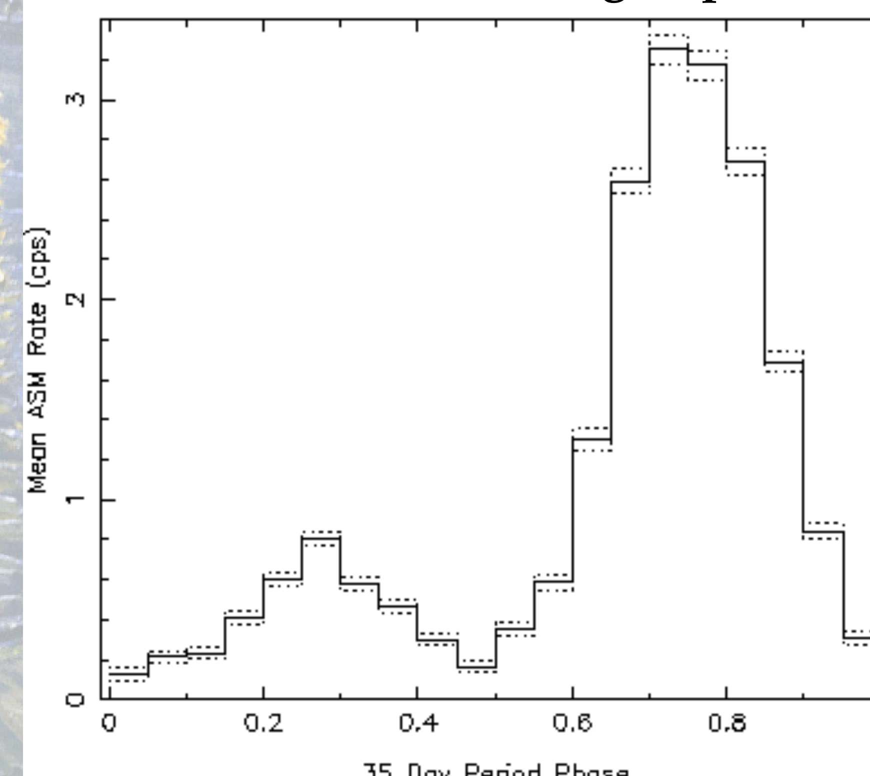

More About High-Resolution Spectroscopy

Your motivation:

this:

becomes

this:

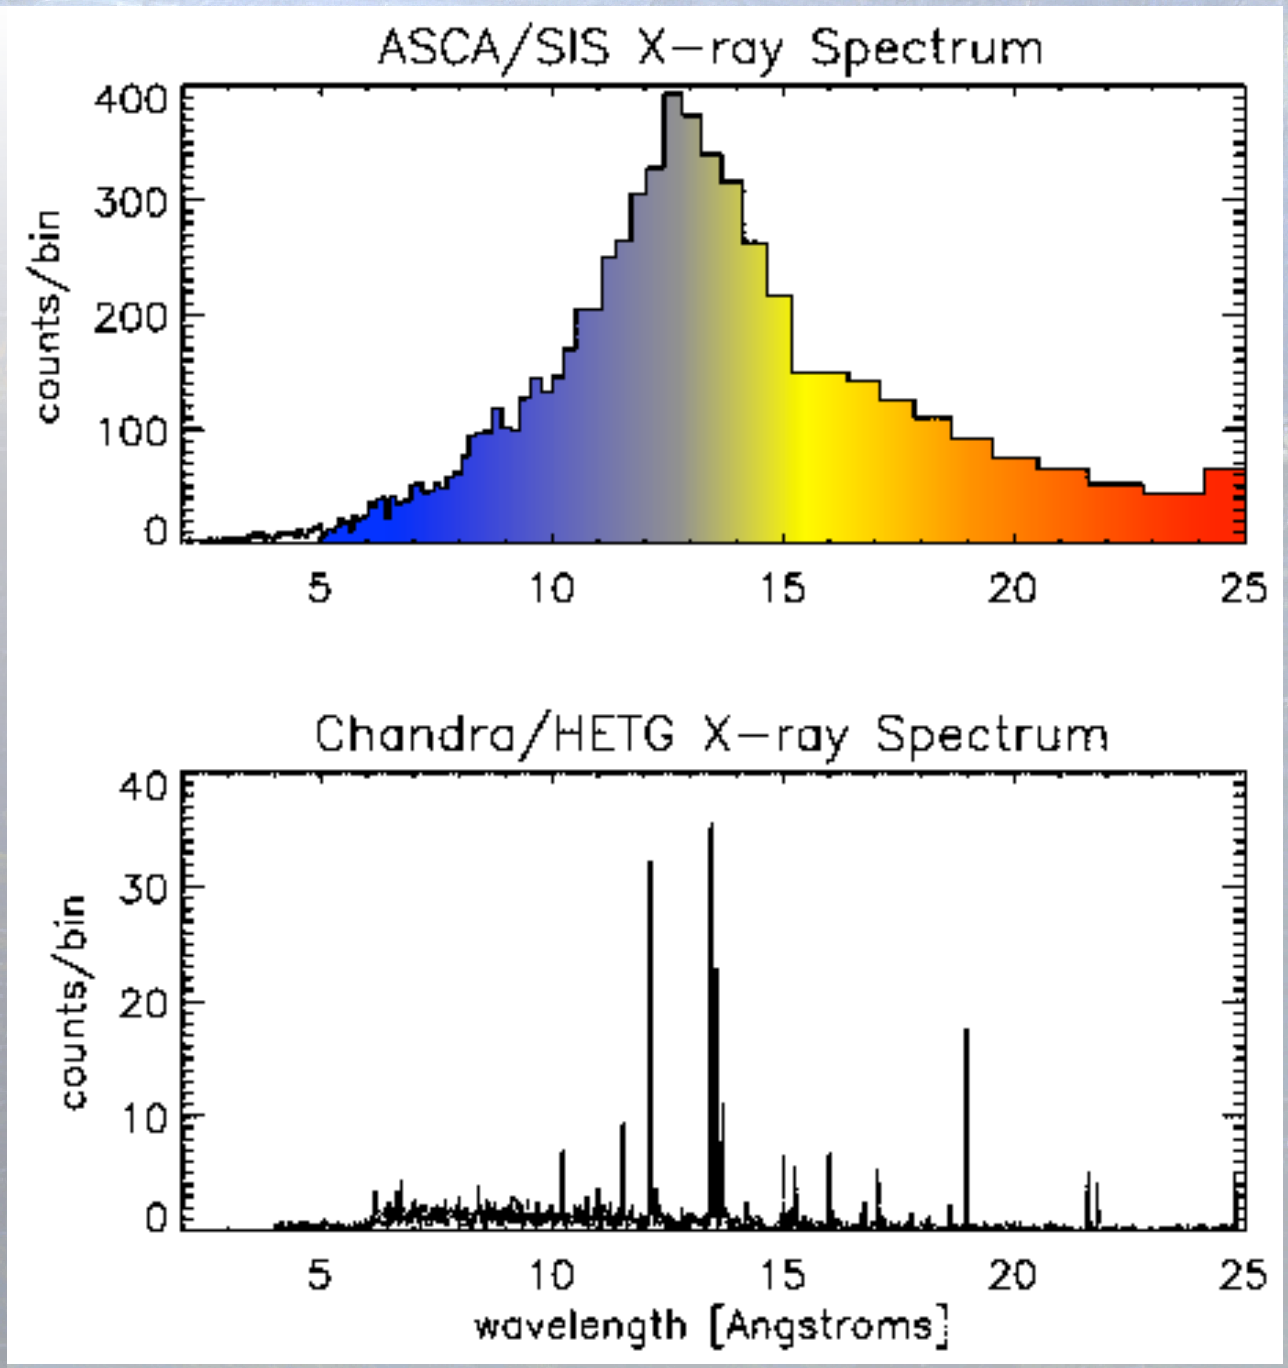

High-Resolution Spectroscopy: more motivation

Emission line spectra ...

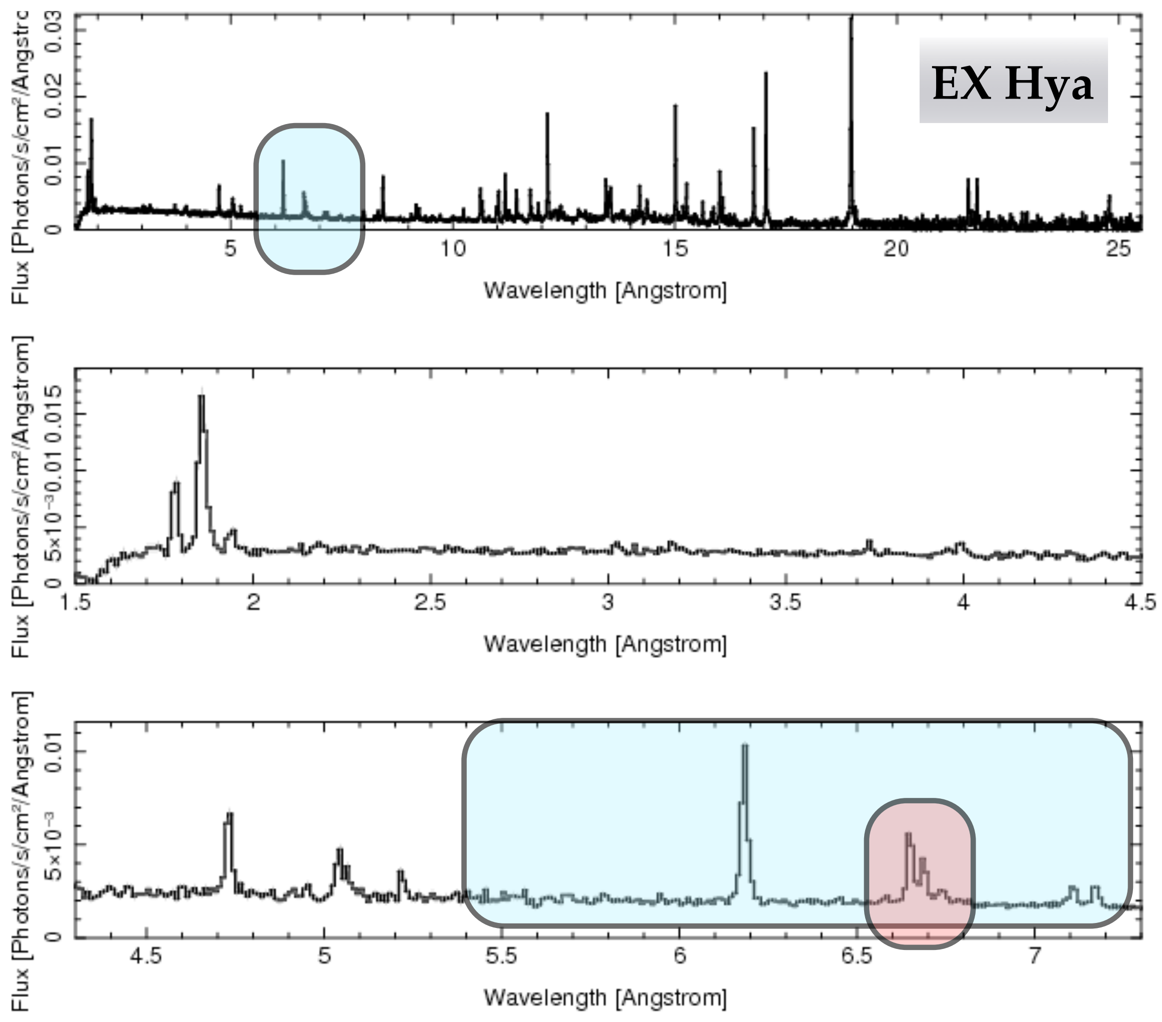

Absorption line spectra...

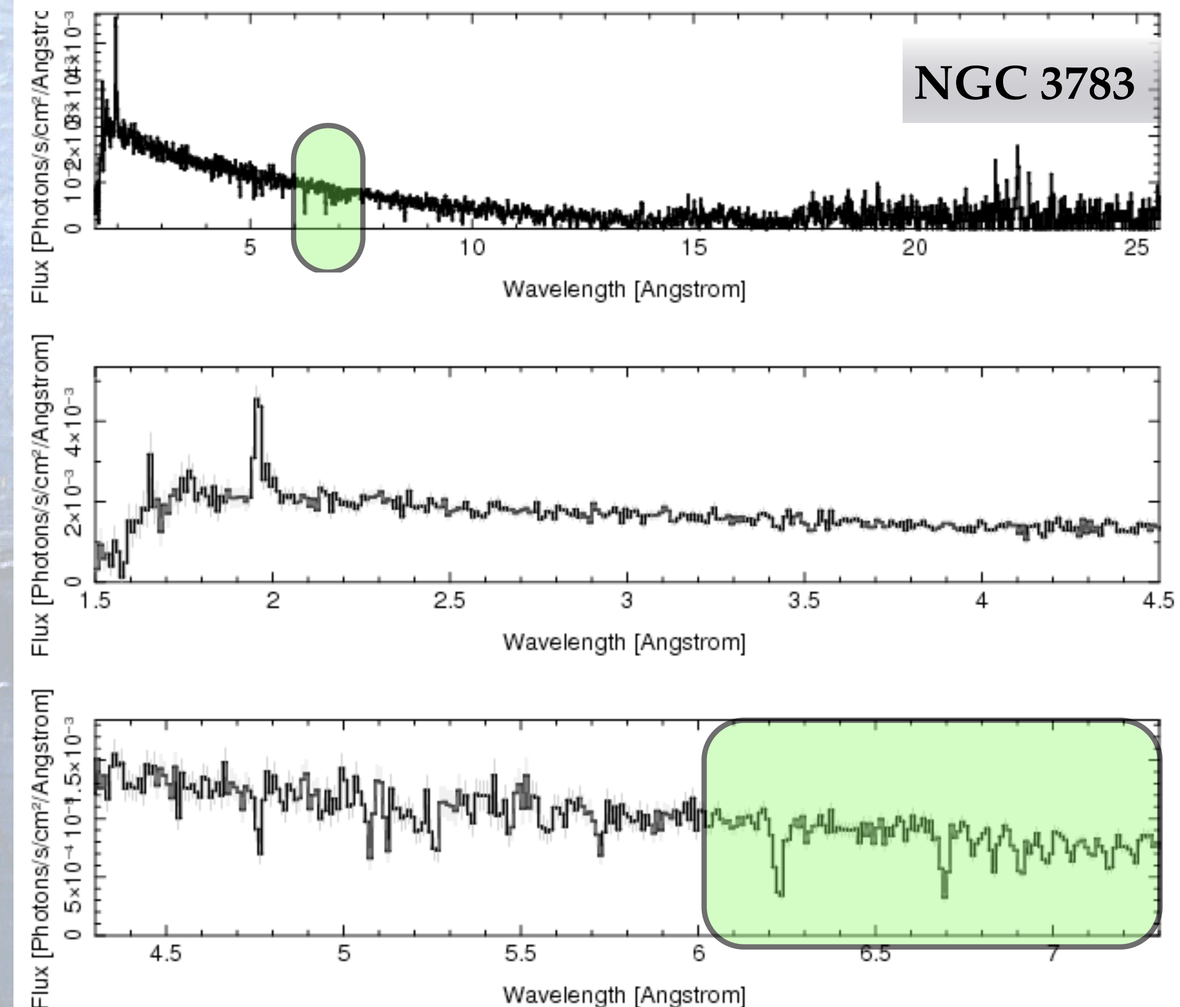

Spectral diagnostics...

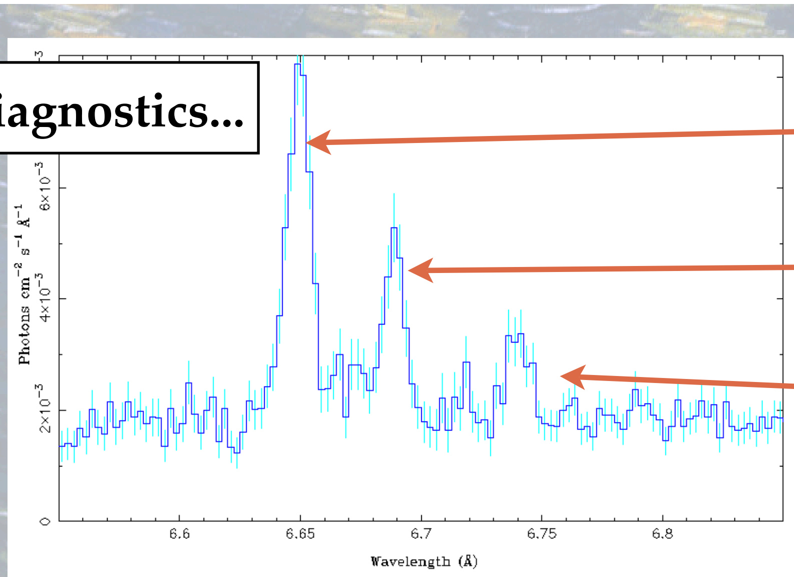

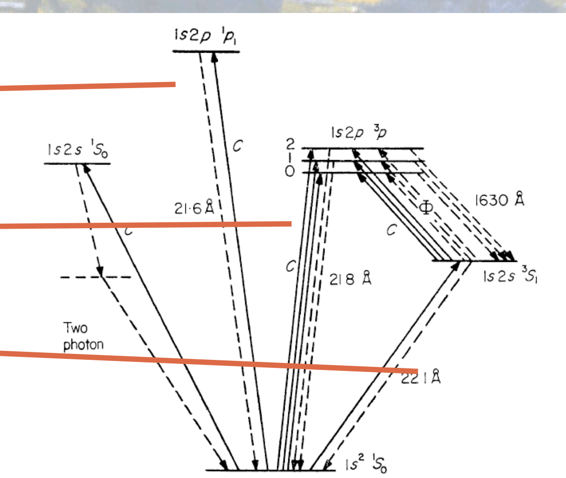

He-like ion triplets

... can provide temperatures, densities (electron, photon, or column), dynamics, composition, ...

High-Resolution Spectroscopy: details

Know your instrument! Dispersive spectrometers (HETG, LETG, RGS) each have their own peculiarities.

LETG grating supports distort the PSF (left: zeroth order) and they cross-disperse, which creates a faint pattern next to the spectrum (right).

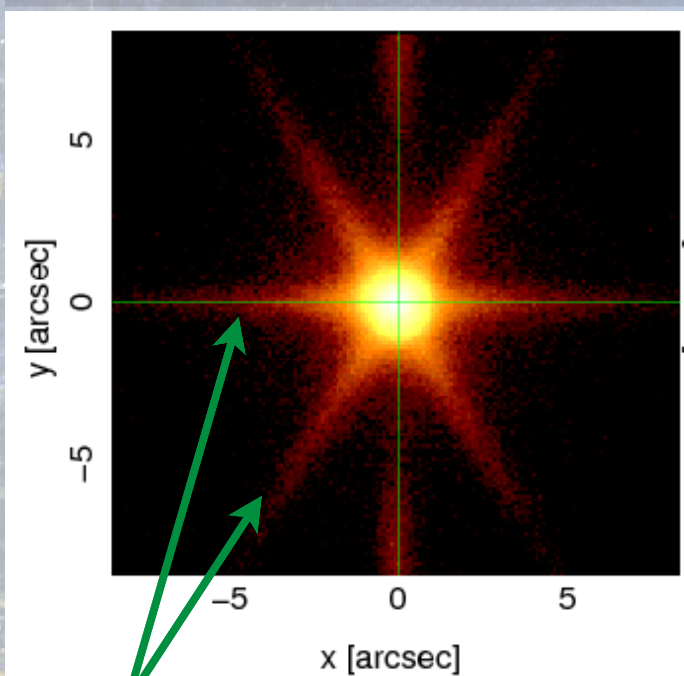

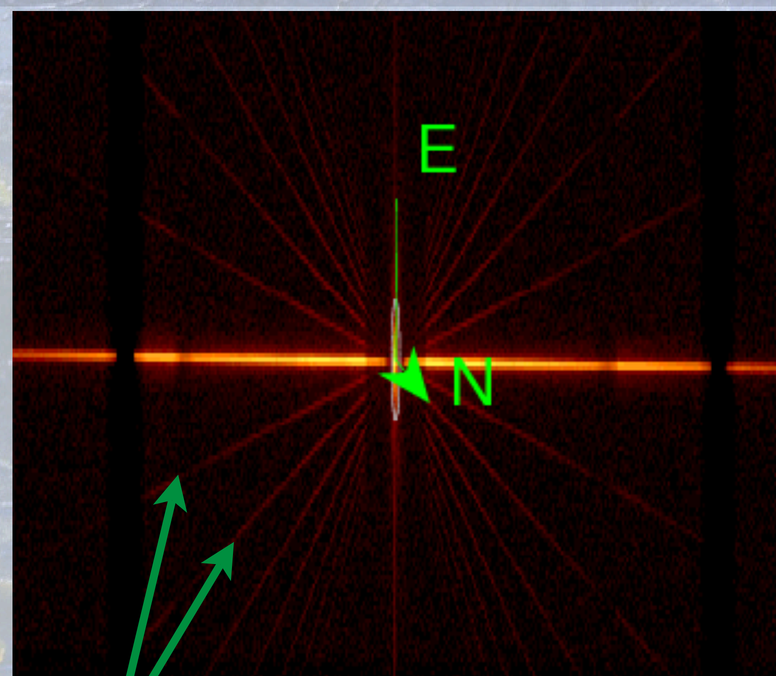

Coarse support diffraction

Fine support diffraction

LETG/HRC-S and RGS have significant background. XMM doesn't dither, and RGS has many bad columns, so beware of sharp "absorption" features.

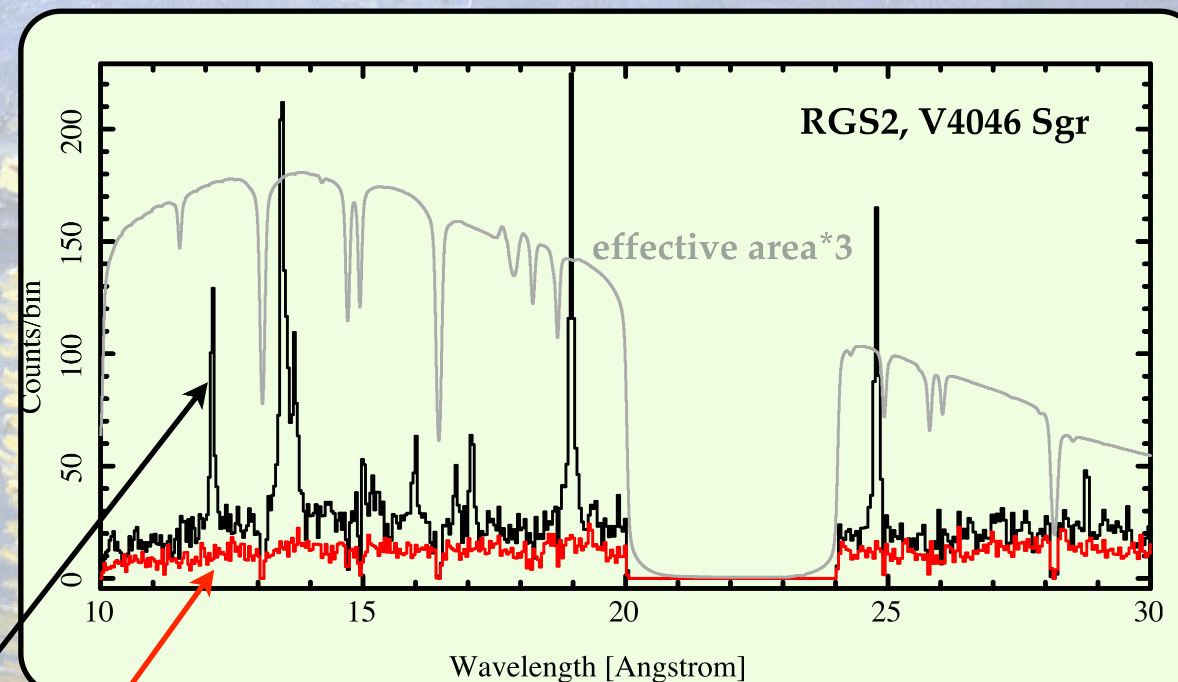

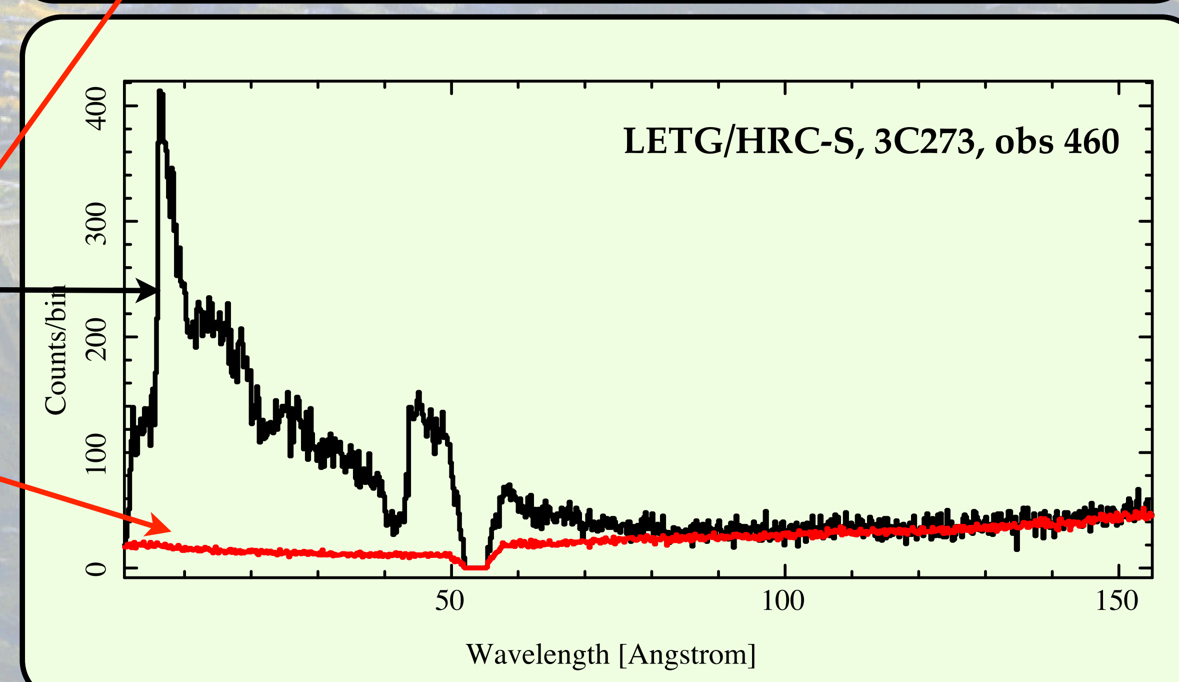

total counts

background counts

54

High-Resolution Spectroscopy: order, order

Dispersive spectrometers create spatially overlapping orders. Using intrinsic detector energy resolution (as with a CCD), the orders can be sorted. At the right is an HETG/ACIS order-sorting image (as found on `tgcat.mit.edu`) and below is a schematic. The *y*-axis plots the grating *m*λ (uniquely determined from the dispersion distance) divided by the CCD "wavelength" (scaled PHA).

These order traces should be horizontal; if not, then the zeroth order position is inaccurate.

The HRC-S has little energy resolution, so the overlapping LETGS orders must be modeled, typically including 8 orders.

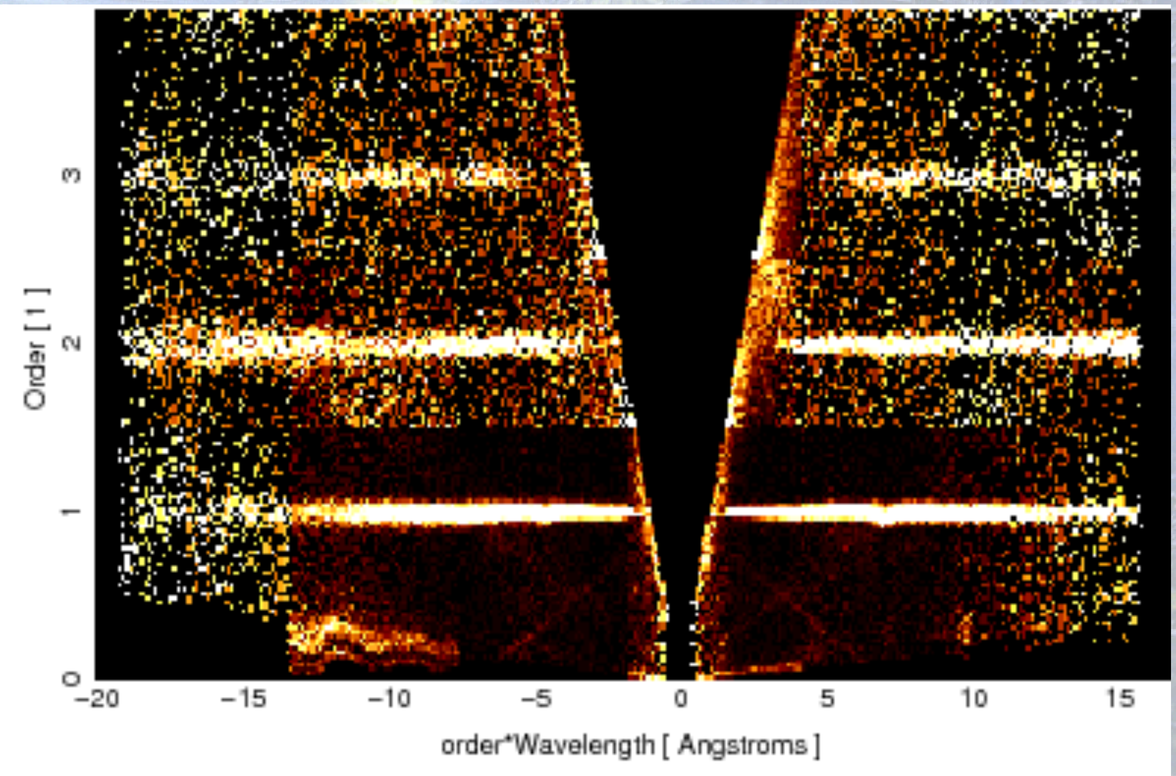

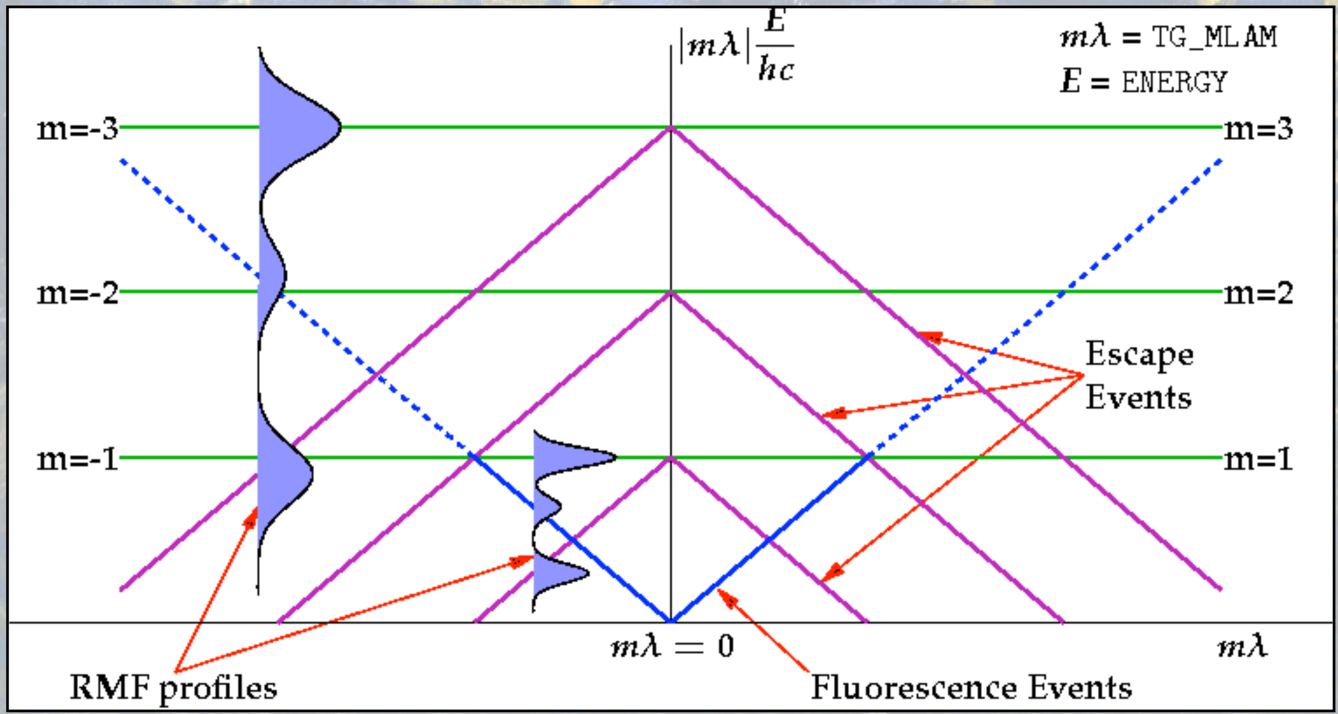

High-Resolution Spectroscopy: unresolved orders

An example of LETG/HRC-S overlapping order contributions.

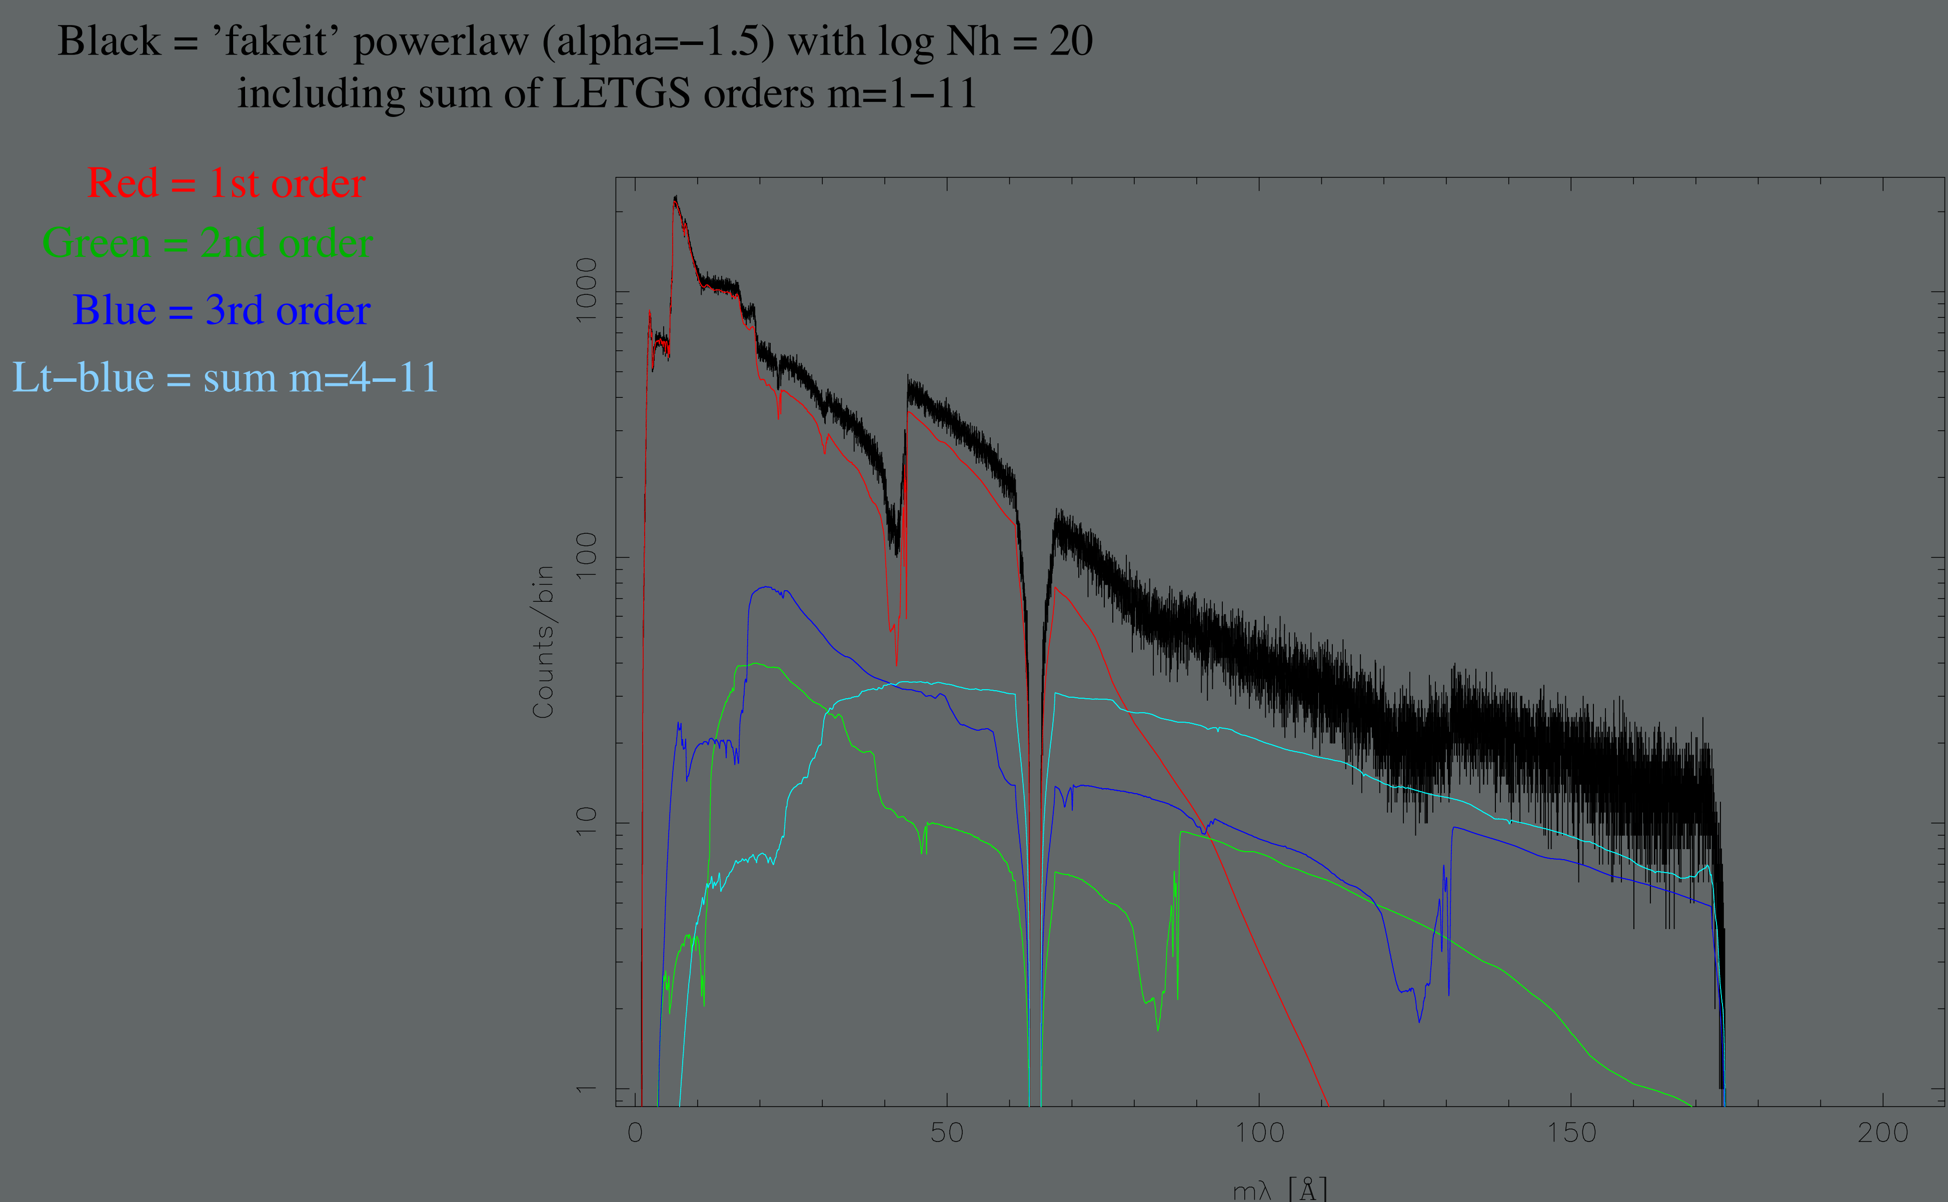

Spectroscopy: CCD Pileup

Pileup is the coincidence of photons in a single charge cloud processed as a “good grade”; in other words, multiple photons per detection cell per frame-time. It distorts the spectrum, it is highly non-linear, and must either be mitigated or modeled.

Very piled zeroth orders

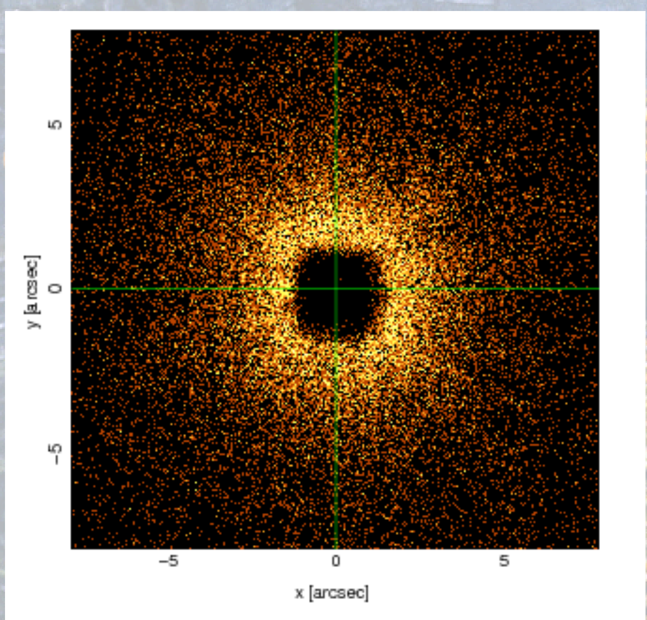

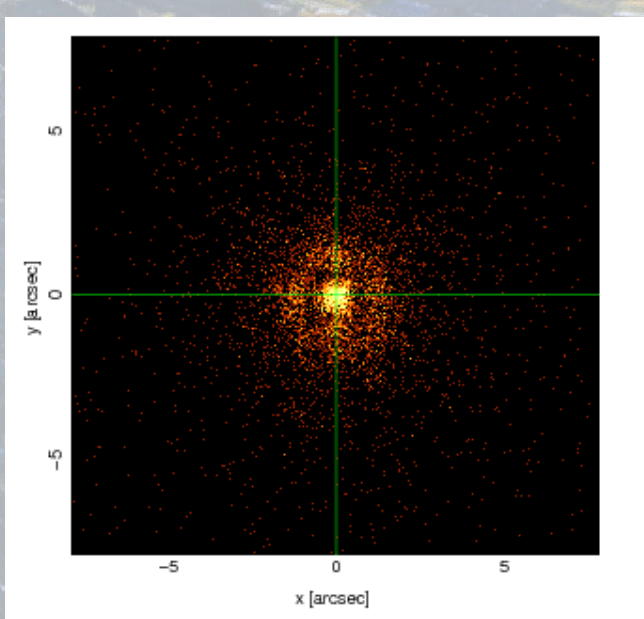

A Suzaku image of pileup fraction (exclude the red region from analysis)

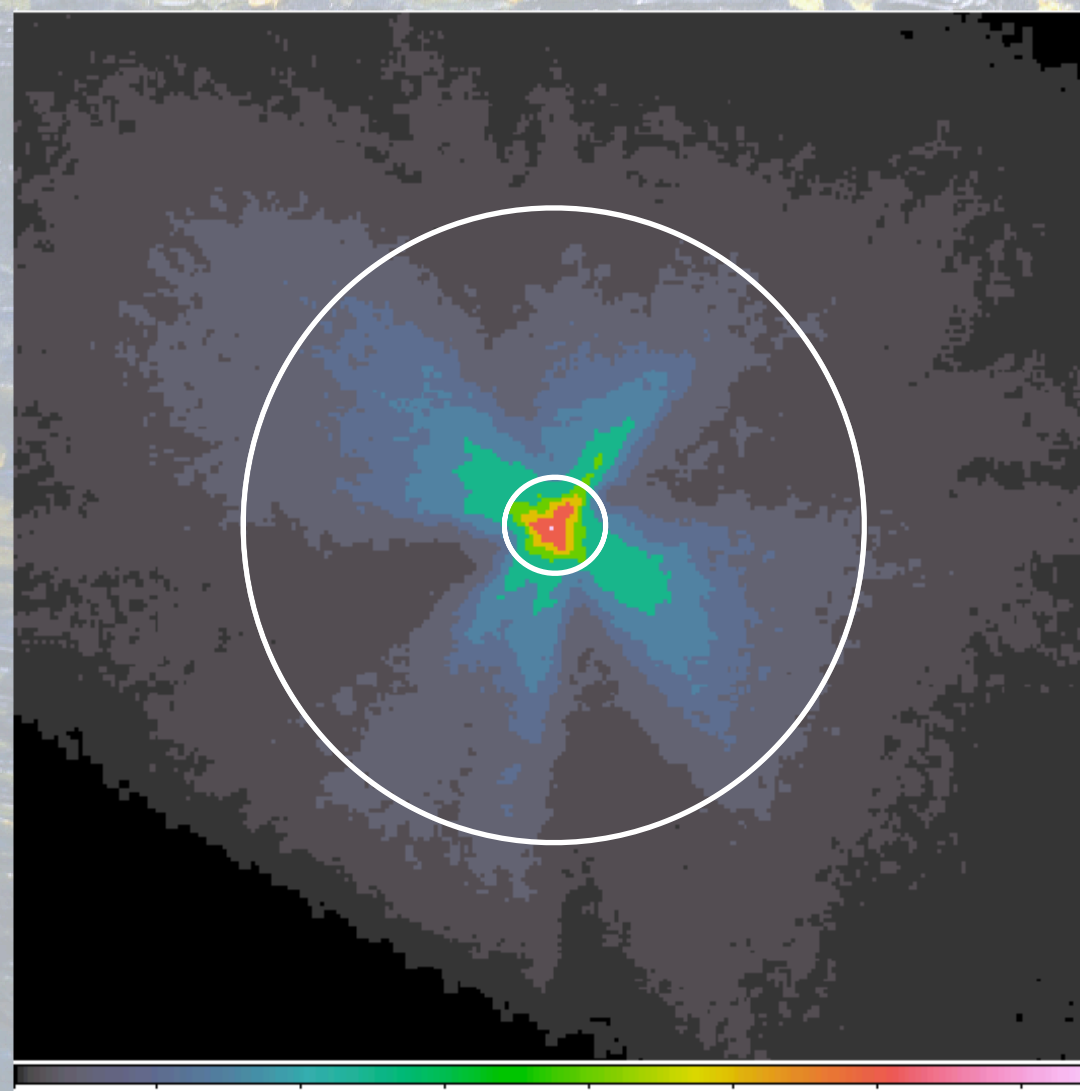

You can’t always tell from the image or spectrum alone whether pileup matters. You need to know the count rate per frame per source.

See the pileup information at:

`cxc.harvard.edu/ciao/why/pileup_intro.html`

`space.mit.edu/cxc/software/suzaku/`


xas
X-ray Astronomy School