Introduction to X-ray Spectroscopy

Olivia Johnson

February 28, 2002

1 Introduction

These notes are meant to introduce the basics of X-ray spectroscopy, outlining how spectral information is obtained and interpreted. I also aim to describe the quality of spectral data available from the current generation of X-ray observatories. While the concepts I discuss are general, I will focus mainly on Chandra data. To illustrate the scientific yield of modern-day X-ray spectroscopy, I'll mention applications related to the well-studied supernova remnant, Cassiopeia A.

I have broken the subject into three sections, each dealing with increasingly precise ways of studying the energy distribution of X-ray data. The first two - broad-band colors (§2.2) and pulse height spectra (§2.3) - fall under the heading of imaging spectroscopy as they rely on the intrinsic energy resolution of the X-ray detector. I'll lead into this topic by discribing X-ray CCDs and X-ray datasets in §2.1, and I'll cover the necessary data reduction techniques in §2.2.1 and §2.3.1. In §3, I'll briefly discuss high resolution dispersive spectroscopy, in which transmission gratings are inserted into the light path between the mirrors and detector.

2 Imaging spectroscopy

2.1 X-ray CCDs and X-ray data

In most ways, X-ray CCDs like those on Chandra and XMM are analogous to familiar optical CCDs. Incident photons release charge in a semiconducting detector through photoelectric absorption. A matrix of potential wells is maintained by a gate structure, which keeps the charge localized throughout the integration. A total charge which is proportional to the photon flux is then moved to a read-out node by clocking the voltages on the gates in sequence, as illustrated in Figure 1. The current produced within in each gated region, or pixel, during the integration is then amplified and recorded. For optical CCDs, this information is used to create a digital image which indicates incident intensity at each location.

X-ray CCDs, however, take advantage of two major differences between X-ray and optical light so as to collect and preserve more information about the incident radiation. The first is that each X-ray photon is energetic enough to liberate a large amount of charge, so there is no need for lengthy integrations to obtain a measurable signal as there is in optical imaging. Secondly, the X-ray countrate from most astrophysical sources is

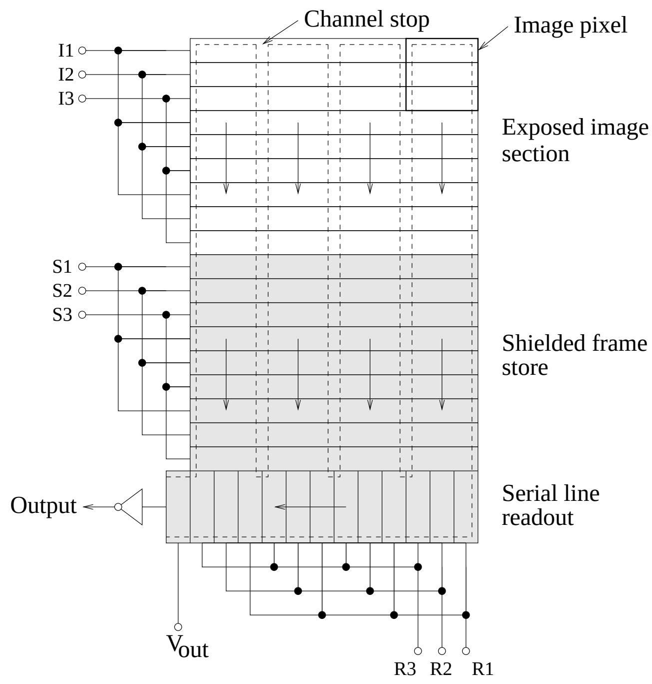

Figure 1: Schematic of a frame store CCD, such as those comprising the ACIS instrument on Chandra. Charge in pixels in the exposed region is moved rapidly (for ACIS, 41 μs) down the columns into the frame store by varying the voltages I. Charge in the shielded region is then moved vertically and horizontally (by clocking voltages S and R, respectively)to the read-out node (V_{out}). For ACIS, the nominal time to empty the full frame store, and therefore the nominal X-ray integration, is 3.2 s. (Figure credit: James Manners.)

very low; it is unlikely that more than one X-ray photon will hit a pixel during a period of order seconds. At X-ray energies, therefore, rapid read-out CCDs can be used as photon counting detectors. Charge resulting from individual photon interactions with the photosensitive detector are moved to the output, where the total signal measured is now indicative of the photon energy.

In contrast to optical imaging data, X-ray data consists of an event list in which various attributes of individual photons are tabulated. A truncated example is shown in Figure 2, in which six types of information are shown for seven photons. The time column gives the arrival time of the photon in unit of seconds (measured since 1 January 1998), and can be extremely precise (3 ms for ACIS; 0.03 ms for XMM pn). The 'sky' column contains the x and y pixel position of the event, and the 'grade' is a measure of how an event affects surrounding pixels. The latter can be used to reject cosmic rays and other false events. The 'pha' column reports the signal recorded at the readout, which is nominally proportional to the incident photon energy. However, due to the finite energy resolution of the CCD ($E/\Delta E \sim 40$), a range of energies will produce the same pulse height. The energies listed in Figure 2 are therefore only approximate.

Much of X-ray analysis is comprised of filtering and binning the event list in different ways. For example, binning on sky position will produce X-ray images, binning on time creates lightcurves, and binning on energy yields spectra. In the next two sections I'll

	time	ccd_id	sky	pha	energy	grade
Units	s		pixel	adu	eV	
Types	double	short	double	long	float	short
1	81629044.87727128	3	(double,double)	938	3917.56	0
2	81632328.00838332	3	(double,double)	3379	13904.6	2
3	81638427.56623022	3	(double,double)	68	386.842	0
4	81640479.11777654	3	(double,double)	252	1138.14	0
5	81650347.95572355	3	(double,double)	51	309.659	0
6	81650947.54026583	3	(double,double)	254	1136.53	0
7	81668711.44813648	3	(double,double)	1594	6586.68	6

Figure 2: Example of a Chandra event list, displaying information about seven detected photons.

discuss two methods of investigating the spectral properties of an X-ray source, and the data reduction required to account for instrumental effects in each.

2.2 Broad-band colors

Coarse spectral analysis can be accomplished by simply comparing the relative numbers of X-ray counts in broad spectral bands. This is commonly done through construction of broad-band 'true-color' images and computation of X-ray hardness ratios. Broad band images are made by first filtering the input event list by energy, and then binning the resulting list on sky position. Hardness ratios are then measured for individual sources by performing aperture photometry on these images. Defined as $HR = \frac{H-S}{H+S}$, where H is the number of counts in a hard (high energy) band and S the number of counts in the soft (low energy) band (e.g. 0.5 - 2 keV), hardness ratios range from -1 for very soft sources to +1 for extremely hard sources. It is important to note that, as explained above, the energy recorded in the event list is not necessarily equal to that of the incident photon. Therefore, the spectral information contained in broad band images is necessarily approximate. In addition, several instrumental effects must be taken into account when creating images from eventlists.

2.2.1 Exposure maps

Mirror effective area, filter transmission, and detector quantum efficiency are all position and energy dependent, and must be corrected to produce a physically meaningful image. This is commonly done by constructing an exposure map which determines the relative sensitivity of different parts of the detector. Dividing by the raw data by this map then corrects for instrumental variation across the chip, analogous to the 'flat-fielding' of optical data. Specifically, the raw image counts C during exposure time $\tau_{\rm eff}$ at position $\hat{\mathbf{p}}$ within pulse height range Δh are related to source photon flux S at a given energy λ by

$$C(\Delta h, \hat{\mathbf{p}}) = \tau_{\rm eff} \int_{\lambda \in \Delta\lambda} d\lambda E(\Delta h, \lambda, \hat{\mathbf{p}}) \mathbf{S}(\lambda, \hat{\mathbf{p}}) \quad (1)$$

where the integral in wavelength extends over all wavelenths which contribute to pulse heights in Δh. The quantity $E(\Delta h, \lambda)$ is the exposure map, in units of cm^2 counts

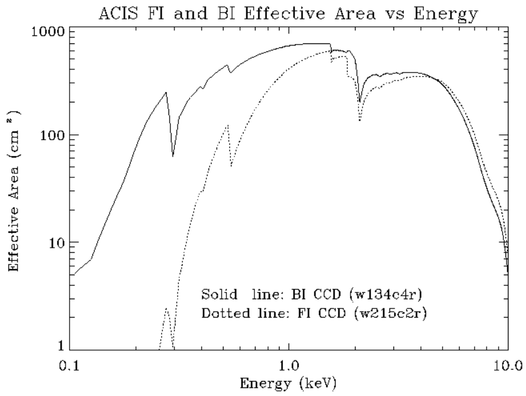

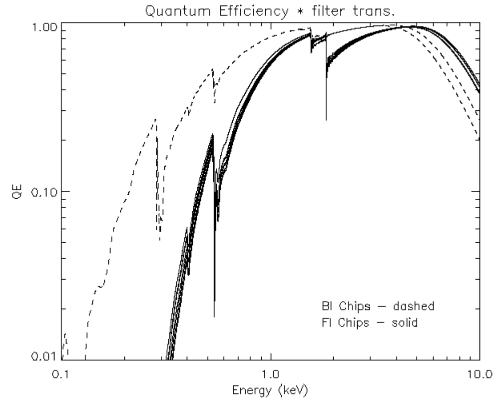

Figure 3: Energy versus a) mirror effective area and b) detector quantum efficiency for Chandra ACIS.

photon^{-1}. Choosing a wavelength range over which this quantity is reasonably constant allows its extraction from the integral, and creation of a 'fluxed' image by division of the raw image by the exposure map and the exposure time:

$$\int_{\lambda \in \Delta\lambda} d\lambda S(\lambda, \hat{\mathbf{p}}) = \frac{C(\Delta h, \hat{\mathbf{p}})}{\tau_{\text{eff}} E(\Delta h, \lambda, \hat{\mathbf{p}})} \tag{2}$$

In practice, construction of an exposure map is a three step process. First, the motion of the source on the detector due to spacecraft dither must be specified. Then, an instrument map, which is essentially the product of the mirror effective area and the filter/detector quantum efficiency is produced. Finally, the aspect information is combined with the instrument map to determine the accumulated exposure at any point on the sky. As shown in Figure 3, effective area and quantum efficiency both vary significantly with energy. Therefore the calculation of the instrument map requires assumption of a source spectrum. Since Equation 2 is valid only so far as $E(\Delta h, \lambda)$ is constant over $\Delta\lambda$, creation of an exposure map suitable for broad bands is done piecewise over locally constant bins. These bins are then weighted according to the assumed source spectrum. For many observations, a single spectrum will not be suitable for all objects within a field, or for all areas of an extended object. An indication of the magnitude of the resulting errors in the inferred flux can be seen in Figure **??**. For multi-color images, made by applying appropriate exposure maps to images made over several broad energy bands, this is a second source of imprecision.

2.2.2 Examples

Multicolor image of Cas A
Elemental images of Cas A

2.3 Pulse height spectra

More detailed and accurate study of X-ray spectral distribution is allowed by the moderate resolution ($E/\Delta E \simeq 50$) pulse height spectra. Pulse height spectra are created by filtering events spatially to isolate source and background regions, then binning the resulting lists on PHA. These spectra are then interpretted through comparison with models of source emission which are convolved with models of instrument response.

2.3.1 Response matrices

Two calibration files must be constructed to facilitate spectral fitting. Spectral 'blurring' due to the finite resolution of the detector is expressed in the X-ray spectral redistribution matrix (RMF), while instrumental effects due to mirror effective area, filter transmission, and detector quantum efficiency are specified in the ancillary response file (ARF).

As discussed above, due to the limited intrinsic energy resolution of X-ray CCDs, photons of a range of energies ΔE are detected at pulse height h. Mapping from energy to pulse height then requires a spreading of the observed counts, which is conveniently expressed as a matrix multiplication. The RMF contains the ensemble of probabilities, $R(E, h)$, of detecting photons of energy E at pulse heights h. For high-resolution systems (i.e. dispersive gratings), the response is very nearly diagonal, while for low-resolution detectors such as proportional counters it is much less so. CCDs such as those discussed in §2.1 are intermediate cases. While the response is largely diagonal there is some off-axis structure, indicating significant non-linearities due to escape peaks and low-energy tails.

RMFs are relatively static and are distributed as standard instrumental calibration products. Their construction for a given dataset requires only correct specification of the detector conditions (i.e. temperature, degree of chip degredation, etc.) and source position to select the appropriate RMF from a pre-defined library. In contrast, the ARF, which contains the mirror, filter, and detector efficiencies as a function of energy, must be calculated for each observation. Like exposure maps, ARFs are based on an aspect histogram which details the position of the source on the chip throughout the exposure. This information is then used to average instrumental effects over exposure time to deduce the combined system response versus energy.

The RMF and ARF are combined to produce a total response with which to convolve candidate source models for comparison with the real data. The source model multiplied by the ARF alone would produce the spectrum observed by a 'perfect' detector with infinite energy resolution; folding this spectrum through the RMF then mimics the resolution of the actual data. Since the RMF is probabilistic, it is possible only to say how likely the observed counts distribution is to have arisen from a given source spectrum. This is done by minimizing a fit statistic, such χ^2 or Cash, through variation of the source model parameters. For spectra with very few counts, which are quite common in X-ray observations, a statistically acceptable fit is often obtained for a variety of models.

2.3.2 Examples

Cas A regions with PHA spectra

3 Dispersive Spectroscopy

Much higher resolution is afforded by dispersive spectrography, which is available on both Chandra and XMM. Functioning simultaneously with the imaging CCDs by sharing the incident light beam, XMM's Reflection Grating Spectrometer provides spectra with $E/\Delta E \simeq 200 - 800$ between 0.35 and 2.5 keV. Chandra has two sets of gratings - High Energy Transmission Gratings (HETG) and Low Energy Transmission Gratings (LETG) - which are inserted behind the mirrors during gratings mode observations and produce spectra with $E/\Delta E \simeq 60 - 1000$.

3.1 Operating principles

The two sets of Chandra gratings are contained in two circular aluminum plates which can be swung into the light-path behind the X-ray mirrors, as shown in Figure 5. Each plate contains hundreds of circular grating facets, each of which is comprised of finely spaced gold-plated grating bars off which the X-rays are dispersed. The support plates are machined so that the centers of the grating facets lie along one surface of the Rowland torus (see Figure 6. The grating spacing, height, and width are tuned to enhance efficiency in odd orders, and minimize even orders. When inserted behind the HRMA, the LETG diffracts X-rays in a dispersed spectrum according the the gratings diffraction relation $m\lambda = p \sin\theta$, where m is the integer order number, λ the photon wavelength, p the spatial period of the grating lines, and θ the dispersion angle. The dispersed spectrum will be focused along a second surface of Rowland torus, which determines the positioning of the spectroscopy mode detectors. The Rowland geometry maintains the telescope focal properties in the dispersion direction for a large range of diffraction angles.

For the LETG, four rings of facets with the same period create a single dispersed image which is usually imaged with the High Resolution Camera spectroscopic array. This camera is a multi-channel plate detector, which provides higher spatial resolution ($\sim$ 6 micron pixels vs. $\sim$ 24 for ACIS), but has little intrinsic energy resolution. Higher-order overlaps can therefore be a significant problem. Its functional wavelength coverage is $\sim$ 2 - 170 $\mathring{A}$. The HETG contains two types of gratings, optimized for medium (0.4 - 5.0 keV) and high (0.8 - 10.0 keV) energy bands, whose grating angles are offset so that the resulting spectra form an 'X' on the detector. As the usual detector for the HETG is a CCD array, its energy resolution can be used to seperate higher order lines, as shown in Figure 7 which plots the dispersion determined wavelength versus the pulse height determined energy.

Figure 8 shows how raw counts are pared by order by spectral filtering, extracted via spatial filtering, and binned into a 1D spectrum.

3.2 Examples

E0102 in light of an oxygen transition, with Doppler motions indicated

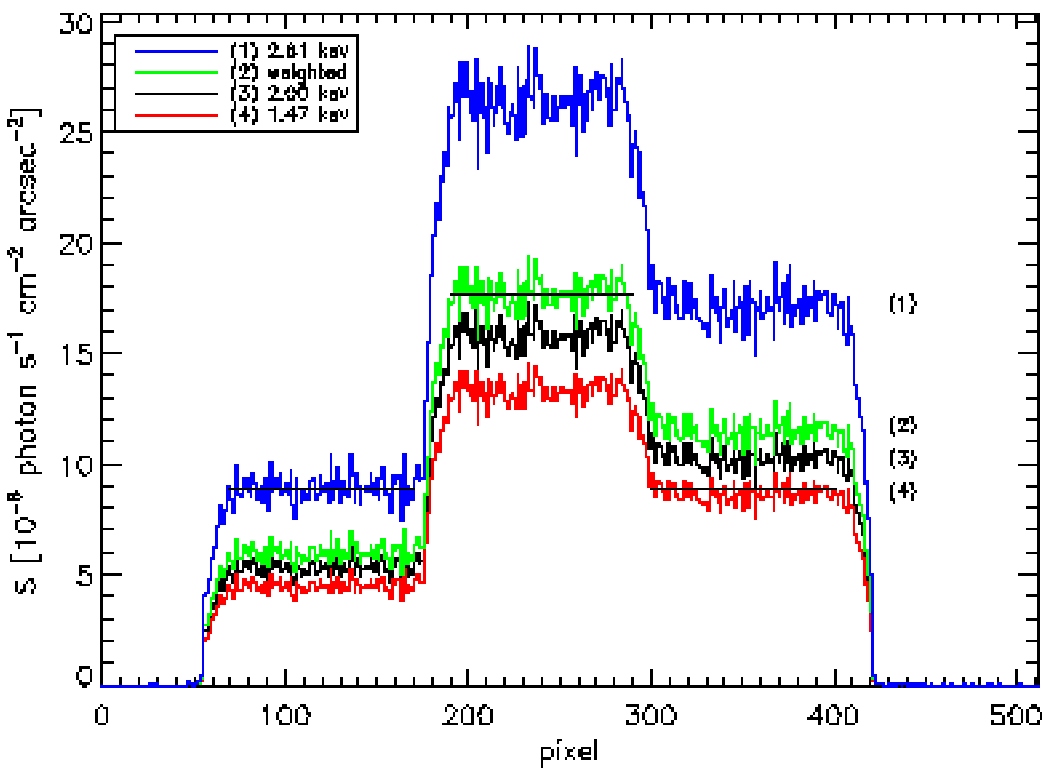

Figure 4: This figure, taken from John Houck's memo *An Introduction to Exposure Maps*, illustrates the flux errors which can result from assumption of a single spectrum in creating an exposure map for a field with various spectral regions. The simulated data shown consists of two partially overlapping disks of uniform and equal surface brightness, one emitting at 1.47 keV and the other at 2.61 keV. The overlapping region therefore has twice the flux comprised of photons of both energies. The black lines indicate the input flux, while the grayscale lines indicate the recovered flux using four different exposure maps - three monochromatic maps at 1.46 keV, 2.61 keV, and 2.0 keV, and one map which weights equally 1.46 keV and 2.61 keV. Even for a known input spectrum, it is impossible to correct the entire image with a single map, and flux errors of up to 40% are seen.

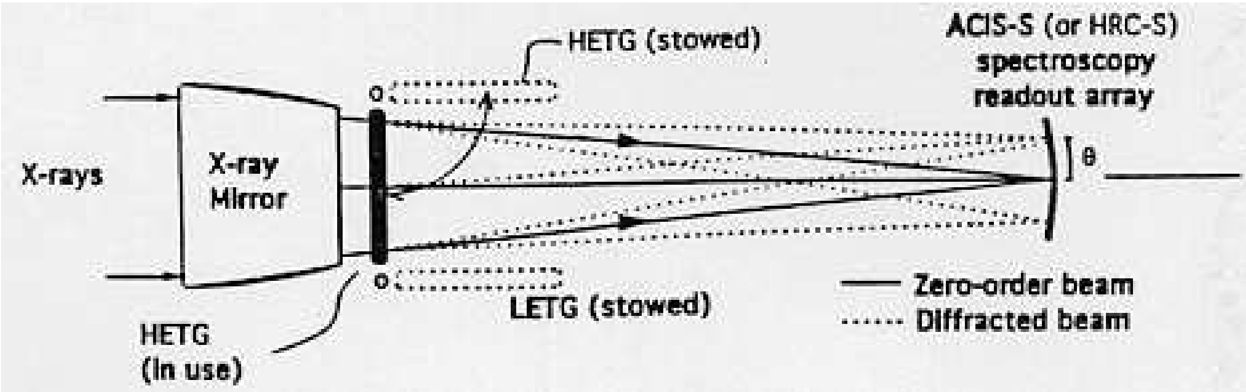

Figure 5: Schematic drawing of Chandra's gratings. In dispersive spectroscopy mode the grating support frames are swung into the lightpath just after the mirrors. The dispersed spectrum is focused on the spectroscopic detector arrays, which are curved to follow the Rowland torus (see next Figure). Figure from M. Zombeck.

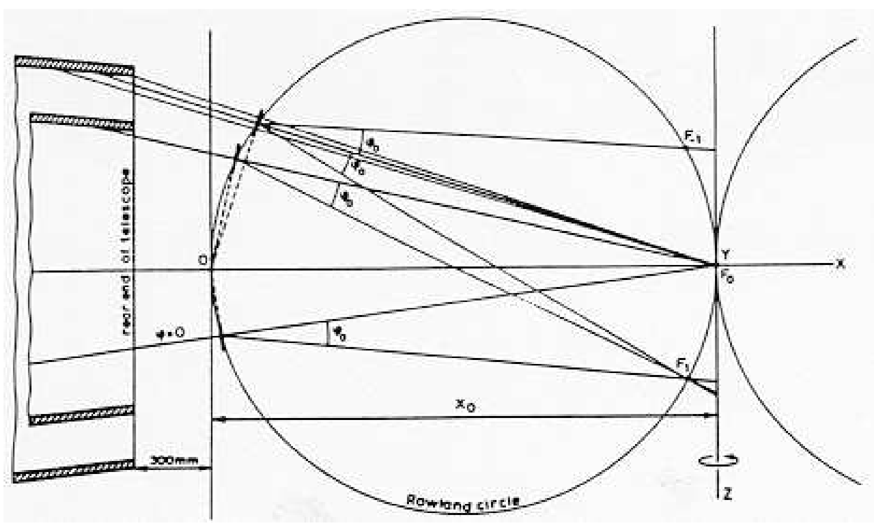

Figure 6: Diagram of Rowland geometry. The circle whose diameter is the distance from the grating plane to the focal plane along the optical axis is the Rowland circle. Rotating this circle around the z axis describes the Rowland torus. To minimize grating induced optical abberation, both grating facets and spectroscopic detectors follow the Rowland geometry. Figure from M. Zombeck.

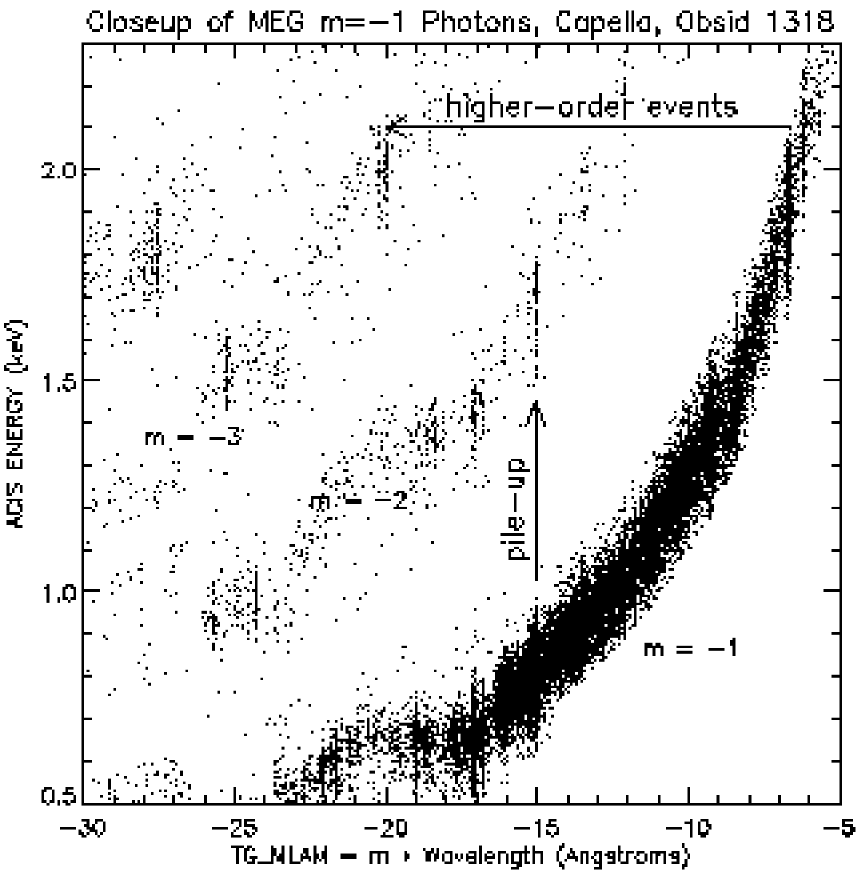

Figure 7: A plot of dispersed wavelength vs. PHA energy, illustrating separation of higher order photons.

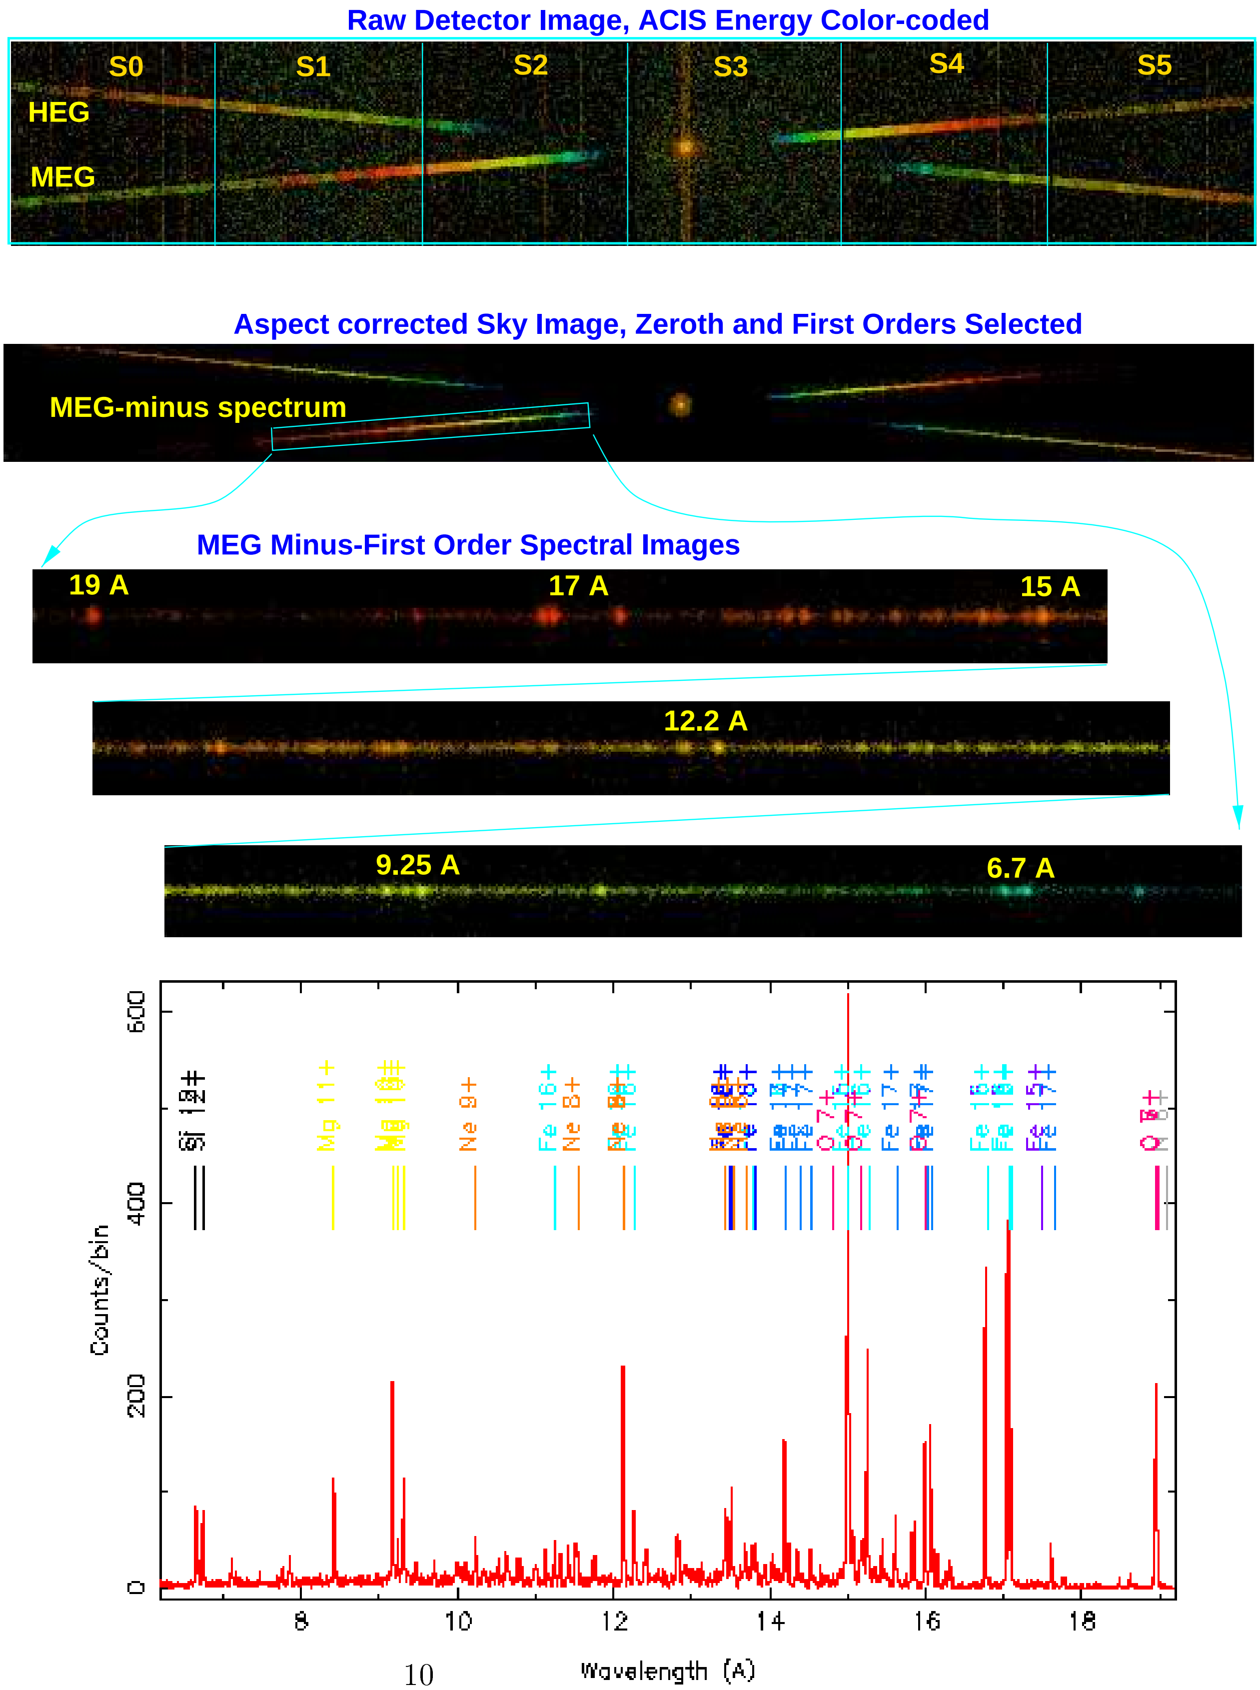

Figure 8: An illustration of the pipeline processing of an HETG observation, beginning with raw detected counts, and ending with a 1D high resolution spectrum.