


Incorporating Systematic Uncertainties into Spectral Fitting

H.Lee*, V.Kashyap, J.Drake, A.Connors, R.Izem, T.Park, P.Ratzlaff, A.Siemiginowska, D.van Dyk, A.Zezas

H. Lee*
Harvard-Smithsonian
Center for Astrophysics
hlee@cfa.harvard.edu




Introduction

X-ray spectral fitting is a process of solving the inverse problem (eq.1) to infer θ , the parameter(s) of a source model S . The observed Poisson photon counts $O(E_i)$ are the main source of the uncertainty in θ estimates, so called statistical error, whereas the systematic uncertainties in A (effective area) and R (response matrix) have been ignored in spectral fitting. This ignorance generally underestimates the error bars of θ . This presentation focuses on handling the A uncertainty in the spectral fitting process and illustrates an efficient way to obtain calibration uncertainty incorporated error bars.

$$O(E_i) = \int S(E; \theta) R(E_i; E) A(E) dE \quad (1)$$

Effective Area (arf)

The plot below shows the coverage of a sample of 1000 ACIS-S arfs generated by Drake et al. (2006) and the default arf (\mathbf{a}_0) is in a black line.




In order to incorporate the arf uncertainty into spectral fitting that affects error calibration results, we propose Bayesian hierarchical modeling for spectral fitting by devising the MCMC algorithms by van Dyk et al. (2001).

Summarizing the black box of arfs (\mathcal{A})

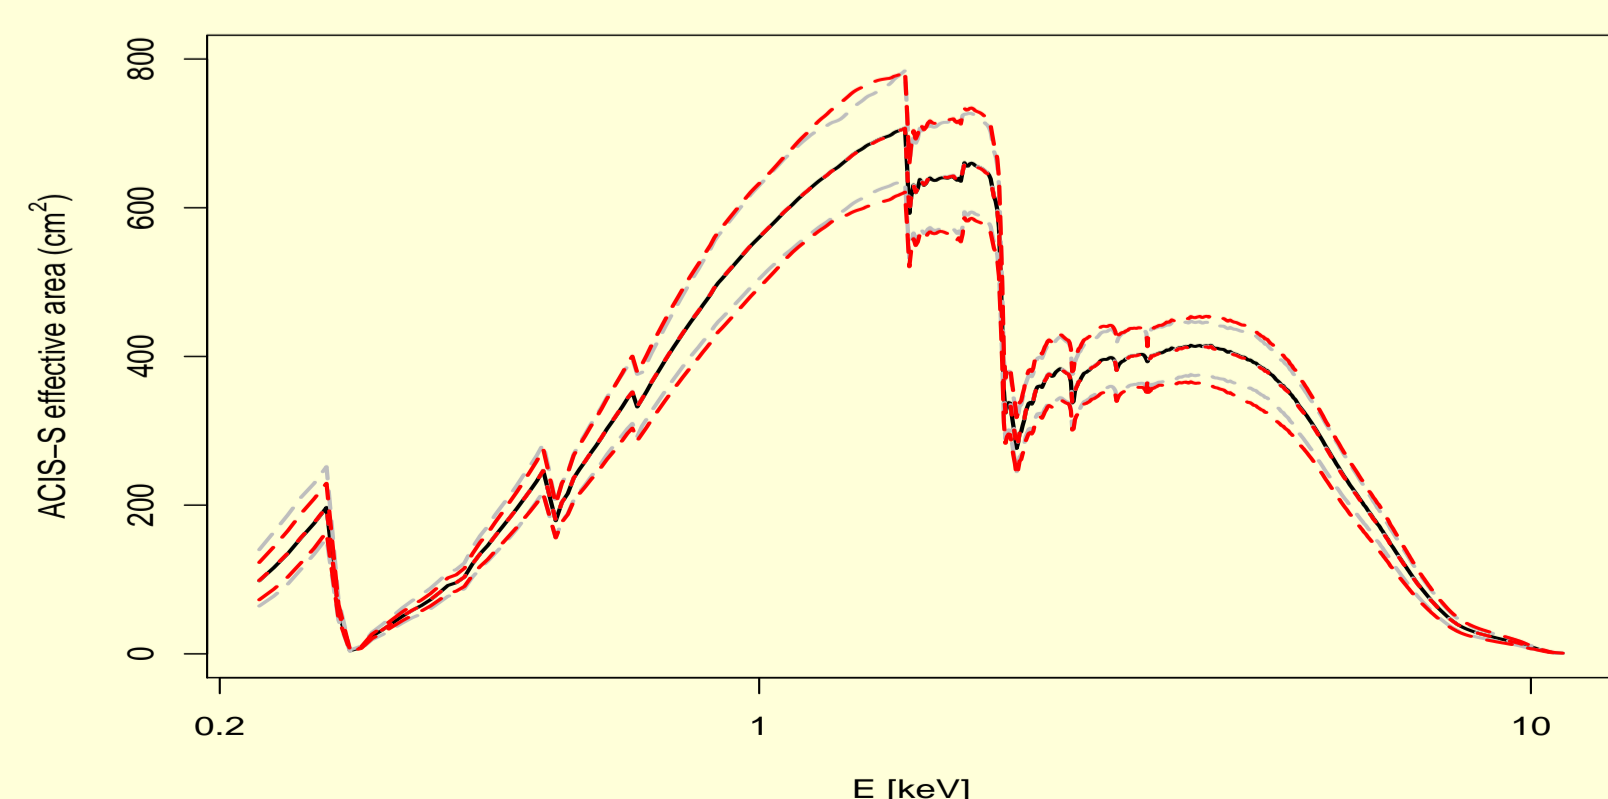
Principal Component Analysis (PCA) reduces dimensionality and summarizes the arf set \mathcal{A} with a small number of principal components (PCs), ready to be utilized into spectral fitting instead of the entire 1000 arf sample by calibration scientists. Let $\mathbf{a}_{(j^*)} = \{\mathbf{a}_{(j^*)}\}_{j=1, \dots, M} \in \mathcal{A}$ be given arfs by calibration scientists on which we perform PCA.

Scree plot: 8 PCs explain 96% of total variation; 12 PCs explain 99%. We will use the first 8 PCs (\mathbf{v}_n) and 8 coefficients (r_n) to simulate arfs. Then, an arf $\mathbf{a}_{(j^*)}$ is generated via:



$$\mathbf{a}_{(j^*)} = \mathbf{a}_0^* + \delta \mathbf{a} + \sum_{n=1}^8 e_n r_n \mathbf{v}_n, \quad e_n \sim \mathcal{N}(0, 1) \quad (2)$$

where \mathbf{a}_0^* is the supplied default arf, \mathbf{a}_0 is the default arf, $\bar{\mathbf{a}}$ the mean of $\mathbf{a}_{(j)}$ s, and $\delta \mathbf{a} = \bar{\mathbf{a}} - \mathbf{a}_0$. The figure below shows the coverage of simulated 1000 arfs in red lines. The 8 PCs are sufficient to match the arf uncertainty represented by gray lines.



Marginalizing over arfs

$$p(\theta|\mathbf{y}) = \int_{\mathcal{A}} p(\theta|\mathbf{y}, \mathbf{a}) p(\mathbf{a}) d\mathbf{a} = \frac{1}{M} \sum_{j=1}^M p(\theta|\mathbf{y}, \mathbf{a}_j)$$

How to marginalize over arfs?

Drake et al. (2006) proposed a strategy [B.0] using standard packages (e.g., XSPEC). We propose three algorithms [B.1-B.3] with BLoCXS (van Dyk et al., 2001). [B.0] tends to be tedious and time consuming depending on the size of the arf library, whose Bayesian counterpart is [B.1]. To speed up [B.1], we introduce [B.2] by selecting arfs randomly from the arf library. To improve computational efficiency, we introduce [B.3]. Given the observed spectrum, [B.0-B.3] work as follows:

[B.0] Fit with XSPEC

Require: M arfs and spectral fitting engines;
for $j = 1, \dots, M$ **do**
 Set a new arf \mathbf{a}_j and fit the spectrum yielding a best fit $\hat{\theta}_j$.
end for
 Compute mean and variance of $\{\hat{\theta}_j\}_{j=1, \dots, M}$.

Repeat fitting procedures as many times as the size of the arf library instead of the supplied default arf. Very tedious!!!

[B.1] Fit with Gibbs sampler

Require: M arfs and Bayesian spectral fitting engines;
Set initial values including priors
for $j = 1, \dots, M$ **do**
 repeat
 Augment data $y_{k|j}$ given $\theta_{k-1|j}$ and \mathbf{a}_j
 Draw $\theta_{k|j}$ from $p(\theta|\mathbf{y}_{k|j}, \mathbf{a}_j)$
 until the chain $\{\theta_{k|j}\}$ is stable, $k = 1, \dots, n_j$.
 Drop n_b draws of a burn-in period.
end for
 Compute mean and variance of $\{\theta_{k|j}\}_{j=1, \dots, M}$.

Extra-tedious! The individual gibbs sequence $\{\theta_{k|j}\}$ offers a statistical error that varies depending on an arf. See the plot at the lower right.

[B.2] Fit with randomized arfs

Require: M arfs and Bayesian spectral fitting engines;
Set initial values including priors
repeat
 Choose $\mathbf{a}_{(j)}$ randomly among M arfs.
 Augment data $y_{k(j)}$ given $\theta_{k-1(j-1)}$ and $\mathbf{a}_{(j)}$
 Draw $\theta_{k(j)}$ from $p(\theta|\mathbf{y}_{k(j)}, \mathbf{a}_{(j)})$
 until the chain $\{\theta_{k(j)}\}$ is stable, $k = 1, \dots, n$.
 Drop n_b draws of a burn-in period.
 Compute mean/variance or mode/HPD from $\{\theta_{k(j)}\}$.

Randomizing arfs saves the *for* loop in [B.1].


[B.3] Fit with PC simulated arfs

Require: PCs (\mathbf{v}_n), coefficients (r_n), and spectral fitting engines;
Set initial values including priors
repeat
 Simulate $\mathbf{a}_{(j^*)}$ based on PCs. (see eq.(2).)
 Augment data $y_{k(j^*)}$ given $\theta_{k-1(j^*-1)}$ and $\mathbf{a}_{(j^*)}$
 Draw $\theta_{k(j^*)}$ from $p(\theta|\mathbf{y}_{k(j^*)}, \mathbf{a}_{(j^*)})$
 until the chain $\{\theta_{k(j^*)}\}$ is stable, $k = 1, \dots, n$.
 Drop n_b draws of a burn-in period.
 Compute mean/variance or mode/HPD from $\{\theta_{k(j^*)}\}$.


We distinguish (j^*), PC simulation from (j), randomization.

Comparison across algorithms

Results from these algorithms work very similarly as shown below but [B.3] is most efficient. One histogram of best fits [B.0] and three posterior density profiles [B.1-B.3] from fitting an *absorbed power-law* spectrum of photon index $\alpha = 2$, column density $N_{\text{H}} = 10^{23} \text{cm}^{-2}$, and total counts $\sim 10^5$ are shown. The black bar indicates a best fit $\pm \sigma$ only with the default arf. The widths of posterior densities represent errors including calibration uncertainty.



Another *absorbed power-law* spectrum ($\alpha = 1$, $N_{\text{H}} = 10^{21} \text{cm}^{-2}$, $\sim 10^5$ cts).



How many arfs?

PCs and coefficients depend on the arf library provided by calibration scientists but our results from PCA indicate that a relatively small number of arfs is sufficient to incorporate calibration uncertainty instead of thousands.


Law of Total Variance (LTV)

LTV explains the complexity of the error decomposition. A best fit depends on arfs and its uncertainty has two components, statistical error and calibration error which are not independent. LTV indicates that the calibration error is dominant with high count data where the statistical error becomes minuscule. This law also explains that [B.3] of 8 PCs (96% calibration error) tends to result in slightly narrower profiles than other algorithms.

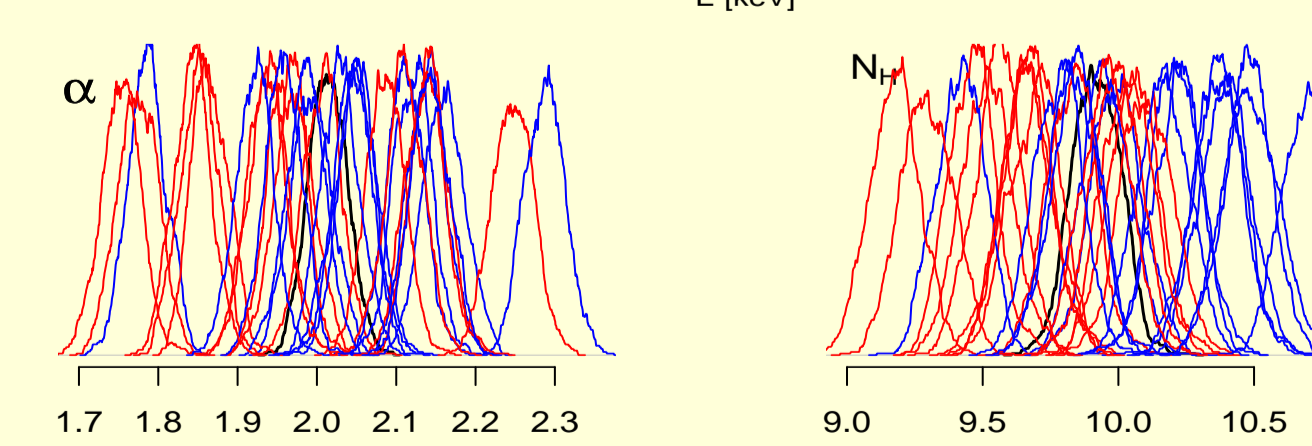
$$V[\theta] = V[E[\theta|\mathbf{a}]] + E[V[\theta|\mathbf{a}]]$$

Behaviors of calibration and statistical errors


Depending on the model used, these two errors may not be separable. In the plot below, two groups of 15 similar arfs are colored and the histograms of gibbs sequences are colored according to the arf colors (default arf in black). The shifting patterns of posteriors do not match between these two spectra. This figure clearly shows that best fit values change with arfs and that calibration uncertainty must be incorporated into spectral fitting.



$\alpha = 2$,
 $N_{\text{H}} = 10^{23} \text{cm}^{-2}$,
 $\sim 10^5$ counts




$\alpha = 1$,
 $N_{\text{H}} = 10^{21} \text{cm}^{-2}$,
 $\sim 10^5$ counts




Asymptotics of calibration error

In the figure below, the horizontal solid line represents the average uncertainty derived from [B.0] and the dashed lines represent the range in this uncertainty obtained from 20 simulations ($\alpha = 2$, $N_{\text{H}} = 10^{23} \text{cm}^{-2}$, $\sim 10^5$ counts). Also shown are the results obtained from combining posterior pdfs by using different numbers of arfs. Dots represent the mean uncertainty and vertical bars denote errors on the means; in other words, N arfs from 1000 are randomly chosen to get the uncertainty of $\frac{1}{N} \sum_{(j)} p(\theta|\mathbf{y}, \mathbf{a}_{(j)})$ for 200 times, and the means and rms errors of these uncertainties are the dots and bars. This figure shows that after $N \approx 25$, the estimated uncertainty is stabilized and therefore, ~ 25 fits with different arfs are sufficient to account for calibration uncertainty provided that the full posterior pdf on the parameters is obtained.



Analyzing Quasar Spectra

Sixteen radio loud quasar spectra from Chandra Data Archive (CDA) are analyzed based on the *PowerLaw*abs* model. Both panels display the error characteristics of estimated power law index α . Calibration error with/without the arf uncertainty is denoted by $\sigma_{\text{tot}}/\sigma_{\text{stat}}$. Three or 4 digit numbers indicate ObsID in CDA. The left panel shows non zero limits in errors due to calibration uncertainty. The right panel displays that systematic errors become more significant in high count spectra than low count ones.



Summary

We have developed a fast, robust, and general method to incorporate effective area calibration uncertainties in model fitting of low-resolution spectra. Because such uncertainties are ignored during spectral fits, the error bars derived for model parameters are generally underestimated. Incorporating them directly into spectral analysis with existing analysis packages is not possible without extensive case-specific simulations, but it is possible to do so in a generalized manner in a Markov chain Monte Carlo framework. We describe our implementation of this method here, in the context of recently codified Chandra effective area uncertainties. We develop our method and apply it to both simulated as well as actual Chandra ACIS-S data. We estimate the posterior probability densities of absorbed power-law model parameters that include the effects of such uncertainties. Overall, a single run of the Bayesian spectral fitting algorithm incorporates calibration uncertainty effectively.

References

Drake, J. et al. (2006). Proc. SPIE, 6270, p.49
van Dyk, D. et al. (2001). ApJ, 548(1), p.224

Acknowledgment

This work was supported by NASA/AISRP grant NNG06GF17G.