

EVALUATION OF ALGORITHMS FOR RECONSTRUCTING ELECTRON SPECTRA FROM THEIR BREMSSTRAHLUNG HARD X-RAY SPECTRA

JOHN C. BROWN,¹ A. GORDON EMSLIE,^{2,3} GORDON D. HOLMAN,⁴ CHRISTOPHER M. JOHNS-KRULL,⁵
 EDUARD P. KONTAR,¹ ROBERT P. LIN,⁶ ANNA MARIA MASSONE,⁷ AND MICHELE PIANA^{8,9}

Received 2005 October 14; accepted 2006 January 15

ABSTRACT

The *Ramaty High Energy Solar Spectroscopic Imager (RHESSI)* has yielded solar flare hard X-ray spectra with unprecedented resolution, enabling reconstruction of mean source electron energy spectra $\bar{F}(E)$ by deconvolution of photon energy spectra $I(\epsilon)$. While various algorithms have been proposed, the strengths and weaknesses of each have yet to be explored in a systematic fashion. For real data $\bar{F}(E)$ is unknown, so these various algorithms must instead be tested on simulated data for which the “true” $\bar{F}(E)$ is known. Accordingly, we devised several forms of $\bar{F}(E)$ with “interesting” features, generated the corresponding (noise-added) $I(\epsilon)$, and recovered $\bar{F}(E)$ using a variety of algorithms, including zero- and first-order Tikhonov regularizations, triangular matrix row elimination, and forward fitting using a parametric form consisting of a double power law with low/high cutoffs plus an isothermal component. All inversion methods reconstructed the general magnitude and form of $\bar{F}(E)$ well, suffering only from (1) blurring of sharp features and (2) poor recovery at low electron energies E in cases in which $\bar{F}(E)$ was positive and large. Addition of a steep thermal component at low E did not prevent recovery of features at higher values of E . Forward fitting did recover large-scale forms and features well but, inevitably, failed to recover local features not expressible within the parametric used. This confirms that inversions are the most dependable way to discover such features. However, examination of the pattern of $I(\epsilon)$ residuals can suggest feature locations and so help refine the parametric form used. Since quite smooth $\bar{F}(E)$ forms do reproduce the observed $I(\epsilon)$ form with relatively small residuals, it appears that sharp features may be uncommon in actual flares.

Subject headings: methods: data analysis — Sun: flares — Sun: X-rays, gamma rays

1. INTRODUCTION

Much of the information on electrons accelerated in solar flares must come from indirect (remote sensing) analysis of the bremsstrahlung radiation signature produced by the accelerated electrons. The quantity of central interest is the electron phase space distribution function $f_e(\mathbf{v}, \mathbf{r})$ [electrons $\text{cm}^{-3} (\text{cm s}^{-1})^{-3}$]. However, for an isotropic cross section (a reasonable approximation for bremsstrahlung emission at nonrelativistic energies; see, however, Massone et al. 2004), the bremsstrahlung production rate is independent of the direction of the electron velocity vector, and hence we need consider only the *speed* distribution $f_e(v, \mathbf{r})$, which has the units of electrons $\text{cm}^{-3} (\text{cm s}^{-1})^{-1}$. We can therefore define the “energy flux spectrum” $F(E)$ (electrons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$) through the relation

$$F(E, \mathbf{r}) dE = v f(v, \mathbf{r}) dv. \quad (1)$$

Since $dE = m_e v dv$, where m_e is the electron mass, we see that $F(E, \mathbf{r})$ and $f(v, \mathbf{r})$ differ only by a multiplicative constant: $F(E, \mathbf{r}) = f(v, \mathbf{r})/m_e$.

In determining the solar accelerated electron spectrum through its bremsstrahlung signature at the Earth, we observe in practice only the volume-integrated hard X-ray photon spectrum $I(\epsilon)$ (photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$) at the Earth. For the optically thin conditions pertinent to the problem at hand, $I(\epsilon)$ is a convolution of the electron flux spectrum $F(E, \mathbf{r})$ with the proton target density $n(\mathbf{r})$ and the (solid-angle-averaged, per target particle) bremsstrahlung cross section $Q(\epsilon, E)$ for production of a photon of energy ϵ by an electron of energy E . In such cases, the only quantity that can be inferred unambiguously from the observations is the (density-weighted volumetric) “mean source electron spectrum” $\bar{F}(E) = \int F(E, \mathbf{r}) n(\mathbf{r}) dV / \int n(\mathbf{r}) dV$ (electrons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$), where dV is an elementary element of volume. Formally, the observed photon spectrum $I(\epsilon)$ is related to $\bar{F}(E)$ and $Q(\epsilon, E)$ by the relation

$$I(\epsilon) = \frac{\bar{n}V}{4\pi R^2} \int_{\epsilon}^{\infty} \bar{F}(E) Q(\epsilon, E) dE, \quad (2)$$

where the mean density $\bar{n} = (1/V) \int n(\mathbf{r}) dV$ and R is the Earth-Sun distance (Brown 1971; Brown et al. 2003). Extracting information on $\bar{F}(E)$ from observations of $I(\epsilon)$ involves either (1) inversion of integral equation (2)—i.e., finding the best model-free nonparametric fit to the data subject to constraints such as smoothness, or (2) forward fitting of a parametric model form of $\bar{F}(E)$ to the observed $I(\epsilon)$ to find the best-fit parameter values for that model.

Both forward-fit and inversion methods are limited in the extent to which they can recover features in $\bar{F}(E)$. Regularized inversion techniques (e.g., Piana 1994; Piana & Brown 1998; Kontar et al.

¹ Department of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, UK; john@astro.gla.ac.uk, eduard@astro.gla.ac.uk.

² Department of Physics, University of Alabama in Huntsville, Huntsville, AL 35899.

³ Department of Physics, Oklahoma State University, Stillwater, OK 74078; gordon.emslie@okstate.edu.

⁴ Code 682, NASA Goddard Space Flight Center, Greenbelt, MD 20771; holman@stars.gsfc.nasa.gov.

⁵ Department of Physics and Astronomy, Rice University, Houston, TX 77005; cmj@rice.edu.

⁶ Space Science Laboratory, University of California at Berkeley, Berkeley, CA 94720; boblin@ssl.berkeley.edu.

⁷ CNR-INFN LAMIA, via Dodecaneso 33, I-16146 Genova, Italy; massone@ge.infn.it.

⁸ Dipartimento di Matematica, Università di Genova, via Dodecaneso 35, I-16146 Genova, Italy; piana@dima.unige.it.

⁹ Dipartimento di Informatica, Università di Verona, strada le Grazie 15, 37134 Verona, Italy.

2004) reduce the level of spurious noise in $\bar{F}(E)$ by requiring overall smoothness in the solution; this inevitably smears sharp features in $\bar{F}(E)$ and reduces their detectability. In the matrix row elimination procedure, followed by Johns & Lin (1992) in their solution of equation (2), regularization is replaced by judicious energy binning (i.e., local smoothing) of the data before inversion to avoid unacceptably noisy $\bar{F}(E)$; the necessary pre-smoothing of the data, and widening of the solution bins, means again that sharp features escape detection because of the need to suppress noise. In the forward fitting of $\bar{F}(E)$ (e.g., Holman et al. 2003), the (generally) small number of parameters utilized, and the corresponding severe constraints on the allowable form of $\bar{F}(E)$, imply that spectral features outwith the scope of the parametric forms used cannot be detected.

With the high-resolution ($\Delta\epsilon \lesssim 1$ keV) hard X-ray spectra now available from the *Reuven Ramaty High Energy Solar Spectroscopic Imager (RHESSI)*; Lin et al. 2002), it is important to assess the relative merits of the various approaches by which $\bar{F}(E)$ can be determined. Since we do not know ab initio the true form of $\bar{F}(E)$, it follows that analysis of actual solar flare data is not a viable way of testing the appropriateness, strengths, or weaknesses of any given method. Rather, what is required is testing on simulated data for which $\bar{F}(E)$ is known. Thus, here we invent forms of $\bar{F}(E)$ with interesting features in them, generate the corresponding $I(\epsilon)$ forms through forward application of equation (2) and addition of a realistic level of data noise, then recover $\bar{F}(E)$ by the various available methods and compare them with the originally assumed $\bar{F}(E)$ to provide a measure of the fidelity and accuracy of each method.

For such an exercise to have objective validity, it is important that the researchers seeking to recover the form of $\bar{F}(E)$ from the synthesized spectra $I(\epsilon)$ have no knowledge of the input form of $\bar{F}(E)$. To this end, the forms of $\bar{F}(E)$ were generated by one of the authors (J. C. B.) and provided to another (A. G. E.), who generated the corresponding (noisy) data $I(\epsilon)$ and then distributed these alone to the other authors. These latter authors, working in the blind, used their reconstruction algorithms to generate forms of $\bar{F}(E)$, which were then sent to A. G. E. for comparison with the original forms. Only after $\bar{F}(E)$ spectra had been recovered by all methods in all cases were the results shared for discussion.

In § 2 we provide an overview of the process by which the synthetic spectra (both electron and photon) are generated. In § 3 we discuss the general principles of the various methodologies used and what one might anticipate as their likely strengths and weaknesses. In § 4 we present the forms of $\bar{F}(E)$ used and the rationale behind them. In § 5 we show the forms of $\bar{F}(E)$ reconstructed by the various recovery techniques and compare them critically with the originally supplied forms. In § 6 we present our conclusions.

2. GENERATION OF THE SYNTHETIC DATA

For the proposed exercise to be credible and of practical value, it is important that (1) the various participants use exactly the same cross section $Q(\epsilon, E)$ in their analysis, and (2) that the synthetic spectra generated have a realistic level of noise comparable with typical large solar flare events as detected by current instrumentation such as *RHESSI*. After much discussion, we agreed to (1) use an isotropic form of the bremsstrahlung cross section, viz., the relativistic 3BN cross section of Koch & Motz (1959) and (2) ignore modifications to the hard X-ray spectrum resulting from photospheric albedo backscatter (e.g., Langer & Petrosian 1977; Bai & Ramaty 1978), the addition of which effectively changes the form of $Q(\epsilon, E)$ to a form $Q_{\text{eff}}(\epsilon, E, \bar{F}[E])$. Although the resulting nonlinearity of the basic problem can be circumvented by a Green's function method (Kontar et al. 2006),

we do not include such considerations here, since our concern is with comparing the reconstruction algorithms in an ideal situation where the true primary photon spectrum is known.

The top left panel of Figure 1 shows photon spectra, characterized by relatively flat spectral forms at all energies, generated from three of the assumed mean source electron spectral forms (*bottom left panel*). These have all been scaled to produce ~ 30 photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$ at the Earth at 30 keV. While the photon spectra look smooth and broadly similar in form, the electron spectra used to generate them are enormously different. This is equally true of the additional three cases shown in the right panels of Figure 1, which are all characterized by fairly steep components at the lower end of the observed energy range; these components are intended to represent thermal sources with temperatures significantly lower than the photon energy range under observation. It is apparent that the various features exhibited by the electron spectra are heavily filtered by the bremsstrahlung cross section so that all the photon spectra are, at first sight, rather featureless. Nevertheless, information regarding the electron spectra is contained in the details of these photon spectral forms, and it is the primary purpose of this paper to evaluate the ability of the various reconstruction algorithms to extract this information from these details.

3. RECOVERY METHODOLOGIES

3.1. Basics

Since the aim of the present paper is to detect *features in* (i.e., the shape of) mean source electron spectra $\bar{F}(E)$, the absolute magnitudes of $\bar{F}(E)$ or of $I(\epsilon)$ are not critical, except inasmuch as they determine the level of noise in the (simulated) data. The observed photon fluxes of Figure 1 are approximately 10 times larger than a major solar flare (see, e.g., Holman et al. 2003; Piana et al. 2003). Hence, in order to generate count spectra with a reasonable level of uncertainty, we used an integration time of 1 s; the count spectra (and their associated uncertainties) therefore correspond to an actual large flare observed over a more reasonable (~ 10 s) time interval. The synthetic hard X-ray spectra $I(\epsilon)$ were generated over the range $\epsilon_{\text{min}} \rightarrow \epsilon_{\text{max}} = (10-150)$ keV. To this photon spectrum we then added random (Poisson) photon noise.

We do recognize that, in reality, additional systematic noise (background, detector response, etc.) should be added to reproduce observed *count* spectra. However, since the goal of the present paper is an isolated comparison of various methods for recovering electron spectra $\bar{F}(E)$ from *photon* spectra $I(\epsilon)$, no attempt was made to include such additional noise terms in the data function. Another way of looking at this is to assume that the observed count spectra have been converted to photon spectra through a preprocessing algorithm prior to their use as data functions in the recovery of the mean source electron spectra.¹⁰

In this section we briefly summarize the various methodologies used to recover the form of $\bar{F}(E)$ from the “observed” photon spectrum $I(\epsilon)$. More details of each method can be found in the references cited. All “inversion” methods (Craig & Brown 1986) involve first discretizing both the mean source electron

¹⁰ Indeed, such a conversion from count spectra to photon spectra is routinely done by the instrument calibration teams, utilizing equations similar in form to eq. (2), but interrelating the count and photon (rather than photon and electron) energy spectra. The conclusions of this paper regarding the ability to discern features in the source electron spectrum from the observed photon spectrum therefore also apply to the ability to discern features in the emitted photon spectrum from analysis of the observed count spectrum. However, it should be noted that count \rightarrow photon reconstructions through instrument response usually involve matrices that are much nearer diagonal in form than those associated with “source-intrinsic” problems (Craig & Brown 1986) such as the photon \rightarrow electron reconstruction problem addressed here.




FIG. 1.—Synthetic mean source electron spectra (*bottom*) and their resulting photon spectra (*top*).

spectrum $\bar{F}(E)$ and the hard X-ray spectrum $I(\epsilon)$ into vectors, so that the integral relation (eq. [2]) may be written in the matrix form

$$\mathbf{I} = \mathbf{Q} \cdot \bar{\mathbf{F}}, \quad (3)$$

where \mathbf{Q} is a matrix with elements corresponding to the values of the cross section integrated over the appropriate photon and electron energy ranges. Formally, inversion of equation (3) provides the vector of \bar{F} values. In practice, for large matrices \mathbf{Q} , the determinant is small and the condition number large, so that direct inversion leads to large amplification of noise in \mathbf{I} and spurious high-frequency structure in $\bar{\mathbf{F}}$. Thus, we have instead to solve equation (3) subject to some constraints that suppress such spurious structure (e.g., Craig & Brown 1986).

3.2. Matrix Inversion with Data-adaptive Binning

The essence of this approach (Johns & Lin 1992) is to take advantage of the upper triangular form of \mathbf{Q} to solve equation (3) by standard square matrix row elimination, starting at high energies. Physically, this comes down to using the fact that

$$\begin{aligned} I(\epsilon + d\epsilon) - I(\epsilon) &= \int_{\epsilon}^{\epsilon+d\epsilon} \bar{F}(E) Q(\epsilon, E) dE \\ &= \bar{F}(\epsilon) Q(\epsilon, \epsilon) d\epsilon, \end{aligned}$$

so that essentially $\bar{F}(E)$ is derived locally from $I(\epsilon)$. In practice, the \mathbf{I} and $\bar{\mathbf{F}}$ vectors must be judiciously adjusted by combining individual fluxes into aggregated bins larger than the data resolution (1 keV in this example) to prevent unacceptably large oscillations in the recovered vector $\bar{\mathbf{F}}$. The method is thus what Craig & Brown

(1986) termed “data-adaptive” and is tantamount to stabilizing the inversion of equation (3) by adjusting \mathbf{Q} through judicious energy binning until its condition number is sufficiently small to yield a meaningful solution for a prescribed data precision.

To avoid numerical problems, the electron bins are shifted upward slightly relative to the photon bins, so that the lower limits of the electron energy bins correspond to the (logarithmic) midpoints of the photon energy bins. Further, because photon data above a certain energy ϵ_{\max} (150 keV in this exercise) are truncated, and since an appreciable fraction of the photon spectrum in the observed range $\epsilon \leq \epsilon_{\max}$ is due to electrons with energies $E > \epsilon_{\max}$, it is necessary in this method (see Johns & Lin 1992) to include an extrapolated form $I_{\text{extrap}}(\epsilon)$ for $\epsilon_{\max} \leq \epsilon \leq \epsilon_{\text{upper}}$, where ϵ_{upper} is an arbitrarily assigned value, here taken to be 1 MeV, or ~ 7 times the maximum photon energy observed.¹¹ The method inverts the spectrum over the augmented range $\epsilon_{\min} \leq \epsilon \leq \epsilon_{\text{upper}}$, but, recognizing that the recovered $\bar{F}(E)$ in the range $E > \epsilon_{\max}$ is based on an assumed form, rather than on data, only $\bar{F}(E)$ points in the energy range $\epsilon_{\min} \leq E \leq \epsilon_{\max}$ are retained.

¹¹ Extending this extrapolated power law to even higher energies has a negligible effect on the results, due to the very low number of photons in the region of observation that are produced by the electrons with energies corresponding to the high-energy end of the extrapolated photon spectrum. The actual choice of upper cutoff ϵ_{upper} (and hence the maximum electron energy for which $\bar{F}[E]$ is calculated) will depend to some extent on the steepness of the observed photon spectrum—steeper spectra may have upper energy cutoffs safely applied at lower values of ϵ_{upper} ; a 5 times maximum photon energy rule of thumb was found empirically to be adequate. To ensure that the results are not sensitive to the value of the upper energy cutoff used, one could recalculate $\bar{F}(E)$ for an increasing range of ϵ_{upper} values until no significant change in the recovered $\bar{F}(E)$ in the range $E < \epsilon_{\max}$ results.

3.3. GSVD Regularized Matrix Inversion

Instead of reducing the resolution locally by binning as in § 3.2, the approach here (Piana et al. 2003; Kontar et al. 2004) is to solve equation (3) with an added smoothness (minimized roughness) constraint. Rather than stabilizing the solution of $\mathbf{Q}\bar{\mathbf{F}} - \mathbf{I} = 0$ by binning adjustment of \mathbf{Q} , we instead solve (Tikhonov 1963)

$$|\mathbf{Q}\bar{\mathbf{F}} - \mathbf{I}|^2 + \lambda|\mathbf{L}\bar{\mathbf{F}}|^2 = \text{minimum}, \quad (4)$$

with the “roughness” penalty being determined by the form of the (user-specified) operator \mathbf{L} and the value of λ chosen. If $\mathbf{L} = \mathbf{I}$ (the identity matrix), then the method is referred to as “zero-order regularization.” Higher-order regularization methods (which force even greater smoothness—e.g., higher order differentiability—on the solution) can be invoked through a different choice of \mathbf{L} . Kontar et al. (2004) discuss how some of these different orders of constraint can be interpreted physically in terms of properties of $\bar{F}(E)$. In practice, the solution to equation (4) is commonly generated (Golub & van Loan 1933) by utilizing the generalized singular value decomposition (GSVD) of the matrix generators \mathbf{Q} and \mathbf{L} (although other approaches to regularization are possible; Thompson et al. 1992; Craig & Brown 1986). The choice of the regularization parameter λ (with small λ corresponding to solutions that are faithful to the data, but noisy, and large λ corresponding to solutions that are smooth, but possibly unfaithful to the data) is determined according to a statistical trade-off between the sizes of the two terms in equation (4).

In these methods, there is no requirement for the matrix \mathbf{Q} to be square, so that the solution for $\bar{F}(E)$ can be found over a range of E that extends beyond the maximum observed photon energy ϵ_{max} . The method (and its SVD/GSVD implementation) depends on a number of factors, such as the order of the method (or, equivalently, the form of the matrix \mathbf{L}), the form of the matrix \mathbf{Q} (square or rectangular), the determination of what array of photon energy bins to associate with the supplied photon data and what array of electron energy bins are used for the recovered \mathbf{F} , and the method by which the regularization parameter λ is determined. In addition, Kontar et al. (2004) have shown how pretransformation of the input and output vectors for \mathbf{I} and \mathbf{F} (either by scaling or subtraction of reference spectral forms) can lead to a better solution.

Two different regularization methods were used in the present study. The first algorithm used a zero-order regularization ($\mathbf{L} = \mathbf{I}$), nonsquare SVD method. The photon data were binned into 1 keV bins, and the midpoint of the photon energy bin was used for the photon energy array; 1 keV bins were also used for the recovered mean source electron spectrum. To transform the problem into one more easily solved, we first prescaled (Kontar et al. 2004) the photon and electron vectors by ϵ^p and E^p , respectively, using an optimum value $p(\gamma)$ obtained from a lookup table. Here γ is the average of the best-fit local power-law spectral indices at low and high energies to the observed photon spectrum; in practice, $p \simeq \frac{1}{3}(\gamma - 2)$. The solution was then rescaled back to the original variables. The regularization parameter λ was determined by a two-step procedure. First, a trial value was obtained using the “discrepancy principle” (Morozov 1966). Then, recognizing that this principle generally results in over-smoothed solutions with large photon spectrum residuals, the value of λ was reduced until the cumulative normalized photon spectrum residual (see Piana et al. 2003; Kontar et al. 2004) lay within $\pm 3 \sigma$ at all points.

The second method was also a nonsquare algorithm but using a first-order (GSVD) algorithm with $\mathbf{L} = \mathbf{D}$, the derivative op-

erator. The photon spectrum was preconditioned by first subtracting the best-fit thermal plus double power-law fit to the observed spectrum (see below); this subtracted spectrum was then added back to the result to get the final overall solution. The method used the same photon and electron energy arrays as the zero-order method, and the same two-stage analysis to determine λ . As explained by Kontar et al. (2004), the use of first-order regularization may result in recovered forms of $\bar{F}(E)$ that are more suited to subsequent differentiation, such as when determining the corresponding injected thick target electron flux spectra $F_0(E_0)$ (Brown 1971; Brown & Emslie 1988). Discussion of the $F_0(E_0)$ problem is deferred to a later paper, but we note here that our preliminary results for $F_0(E_0)$ are also good [albeit noisier than those for $\bar{F}(E)$ due to the second deconvolution involved]. They still reveal the presence of marked features in $F_0(E_0)$ and also yield unphysical results for $F_0(E_0)$ in cases in which the model $\bar{F}(E)$ used is not in fact compatible with a thick target model (Brown & Emslie 1988).

3.4. Forward Fitting

In this approach (Holman et al. 2003), a parametric form of $\bar{F}(E)$ is assumed, and the corresponding form of $I(\epsilon)$ produced. The residuals between the $I(\epsilon)$ generated by the forward-fit form and the data provided) are normalized to the standard error on each photon data point, and the χ^2 minimized with respect to each parameter to determine the “best” form for $\bar{F}(E)$. In this study we assumed parametric forms corresponding to a single-temperature thermal distribution (characterized by a total emission measure EM and a temperature T) plus a broken power law with adjustable low- and high-energy cutoffs ($AE^{-\delta_1}$, $E_{\text{low}} \leq E \leq E_{\text{brk}}$; $AE_{\text{brk}}^{\delta_2 - \delta_1} E^{-\delta_2}$, $E_{\text{brk}} < E \leq E_{\text{max}}$). There are thus (up to) eight independent parameters (EM, T , A , E_{low} , E_{brk} , E_{max} , δ_1 , δ_2) in the fit. This method does not restrict the energy range of $\bar{F}(E)$ to $E < \epsilon_{\text{max}}$. However, since a specific form is assumed and since the number of parameters that one can specify is necessarily limited, it follows that some types of feature in $\bar{F}(E)$ can *never* be inferred by this method, although their presence might be suspected where the residuals are large.

4. INPUT FORMS OF $\bar{F}(E)$

After some experimentation, we considered six different forms for the input mean source electron spectrum $\bar{F}(E)$ (electrons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$). The first three are empirical basic forms chosen with challenging properties, while the second three are similar but with the addition of a large low-energy “thermal” components to test whether its presence (as in real flares) masks recovery of high-energy features. As discussed in § 3.1, all are scaled to produce photon spectra with ~ 30 photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$ at the Earth at a photon energy of 30 keV.

4.1. Basic Forms

These $\bar{F}(E)$ forms are shown (on logarithmic scales) in the bottom panels of Figure 1.

Case A:

$$\bar{F}_A(E) = \frac{8.0 \times 10^{56}}{\bar{n}V} \begin{cases} \left(\frac{E}{30}\right) \left\{ \frac{2}{27} \left[1 - \left(\frac{30}{35}\right)^3 \right] + \frac{1000}{E^3} \right\}, & E \leq 30, \\ \left(\frac{100}{E^2}\right) \left[1 - \left(\frac{2}{3}\right) \left(\frac{E}{35}\right)^3 \right], & 30 < E \leq 35, \\ \frac{100}{3E^2}, & E > 35. \end{cases} \quad (5)$$

This spectral form is the smoothest of those studied; it is continuous and decreasing everywhere, but has a “knee” in the 30–35 keV range, within which its gradient sharply increases. It is the form closest to those used in forward-fit parametric forms.

Case B:

$$\bar{F}_B(E) = \frac{1.8 \times 10^{54}}{\bar{n}V} \begin{cases} E \frac{21}{64} + \frac{1}{9} \left[1 - \left(\frac{E}{50} \right)^9 \right], & E \leq 50, \\ \frac{E}{192} \left(\frac{200}{E} \right)^3 - 1, & 50 < E \leq 200, \\ 0, & E > 200. \end{cases} \quad (6)$$

This form is slowly increasing below 50 keV and also has an upper cutoff at 200 keV, which is above the maximum photon energy in the supplied hard X-ray “data.” This form was chosen for two reasons. First, we wish to see if the upper energy cutoff at 200 keV can be detected through its influence on the photon spectrum in the “observed” (10–150) keV range. Second, we wish to determine the ability of the various methods to recover an increasing form of $\bar{F}(E)$.

Case C:

$$\bar{F}_C(E) = \frac{5.5 \times 10^{55}}{\bar{n}V} \times \begin{cases} \left(\frac{E}{30} \right)^8, & E \leq 30, \\ \left(\frac{E}{30} \right)^{-3}, & 30 < E \leq 250, \\ 0, & E > 250. \end{cases} \quad (7)$$

This case is an exaggerated version of case B. It is very steeply increasing below a “cusp” at 30 keV and has an upper cutoff at 250 keV, which is again above the maximum photon energy in the supplied hard X-ray data. The steeply increasing form of $\bar{F}(E)$ below the cusp implies that, contrary to more conventional cases, the photon emission at low energies is *not* dominated by the electron spectrum near these energies but rather comes from $E \gg \epsilon$. Hence, even small errors in determining the large $\bar{F}(E)$ at, say, 30 keV have a profound effect on the photon emission at, say, 15 keV and hence (see below) on the determination of the electron spectrum at 15 keV, which makes only a small contribution to $I(\epsilon = 15 \text{ keV})$.

4.2. Forms with a Thermal Component Added

These forms are shown in the right panels of Figure 1.

Case D:

$$\bar{F}_D(E) = \frac{1.0 \times 10^{57}}{\bar{n}V} \times \begin{cases} 500e^{-E/2.7} + \left(\frac{10}{E} \right)^{2.5} e^{-E/150}, & E \leq 42, \\ 0.2 \left(\frac{10}{E} \right)^{2.5} e^{-E/250}, & 42 < E \leq 50, \\ \left(\frac{10}{E} \right)^{2.5} e^{-E/150}, & E > 50. \end{cases} \quad (8)$$

This case is the first of those characterized by a large low-energy thermal component. This part of the spectrum corresponds to a differential emission measure $\xi(T) \sim T^{-1.5}$ above $kT_0 = 2.7 \text{ keV}$. We deliberately chose to model the low-energy part of $\bar{F}(E)$ with a nonisothermal form in order to test how well such a form is dealt with by forward-fit methods that constrain the low-energy part of $\bar{F}(E)$ to be isothermal.

Case E:

$$\bar{F}_E(E) = \frac{2.0 \times 10^{57}}{\bar{n}V} \times \begin{cases} 300e^{-E/3} + \left(\frac{10}{E} \right)^{3.7}, & E \leq 55, \\ \left(\frac{1}{7} \right)^4, & 55 < E \leq 70, \\ \left(\frac{10}{E} \right)^4, & 70 < E \leq 250, \\ 0, & E > 250. \end{cases} \quad (9)$$

This form is characterized by a sudden decrease (by almost an order of magnitude) at 55 keV and by a plateau region from 55 to 70 keV. It also has the steepest spectrum at high energies. Apart from the thermal component, which has the same form as case D but with $kT_0 = 3 \text{ keV}$, it has a power-law index of 3.7 at low and 4.0 at high energies.

Case F:

$$\bar{F}_F(E) = \frac{1.5 \times 10^{57}}{\bar{n}V} \times \begin{cases} 100e^{-E/2.7} + 0.02 \left(\frac{E}{40} \right)^7, & E \leq 40, \\ 0.02 \left(\frac{40}{E} \right)^3, & 40 < E \leq 350, \\ 0, & E > 350. \end{cases} \quad (10)$$

This form is characterized by a relatively smooth “dip” extending from ~ 20 to $\sim 40 \text{ keV}$ and is reminiscent of the spectrum obtained by Piana et al. (2003) in their analysis of the 2002 July 23 flare. It is, in fact, similar to the cusp form of case C plus a thermal form as in cases D and E, with $kT_0 = 2.7 \text{ keV}$.

5. RESULTS

5.1. Presentation

Figure 2 shows the reconstructed mean source electron spectra, using the various methodologies outlined in § 3, for all six cases. The top set of lines in each figure shows the input spectrum (*dashed line*), shifted upward by a factor of 10 (for ease in visualization of the results), together with the similarly scaled reconstruction using the zero-order regularization algorithm. The second set of lines shows the unshifted spectrum, superposed on the reconstruction, obtained using the first-order regularization algorithm. The third set of lines shows the reconstruction obtained by the matrix binning method (shifted downward by a factor of 10), and the bottom set shows the results of the forward-fitting method, (shifted downward by a factor of 100).

In the case of the adaptive binning method the bounds on the solution are shown as error boxes whose width and height indicate, respectively, the local bin E range (“resolution”) and the uncertainty in \bar{F} . (We refer the interested reader to the Appendix for a brief discussion of how these are related and should be interpreted). The error boxes were obtained successively, starting at high energy, by increasing the bin width ($\Delta \log E$) until, for each energy band, the fraction of the dynamic range of $\log F$ represented by each 1σ error box height ($\Delta \log F$) was approximately the same as the fraction of the dynamic range of $\log E$ represented by that energy band ($\Delta \log E$). (In this way, 1σ error boxes would look approximately square on a log-log plot with axis scales just covering the full range of each variable.) This equality is a subjective choice akin to the choice of smoothing parameter in the Tikhonov regularization methods and could be




FIG. 2.—Recovered forms of $\bar{F}(E)$ obtained by the various methods. Top set of curves (*green*): zero-order regularization, shifted upward by a factor of 10; second set of curves (*red*): first-order regularization, unshifted; third set of curves (*brown*): binned-matrix-inversion method, shifted downward by a factor of 10; bottom set of curves (*blue*): forward fit, shifted downward by a factor of 100. The dashed curves show the true forms of $\bar{F}(E)$. The forward-fit solution is plotted without error bars. For the top two curves (regularized solutions), the results are plotted as a 3σ river, i.e., the band within which $\sim 99\%$ of all (sufficiently smooth) solutions lie; the horizontal lines illustrate the energy resolution as a function of electron energy that is achieved. Results for the binned-matrix-inversion method are shown as error boxes; the vertical and horizontal extents of each box correspond to the 3σ uncertainty level and energy resolution, respectively.

increased or decreased if we wanted to give more emphasis to bin width (resolution) or to accuracy, respectively. That is, energy resolution and uncertainty are coupled by an “uncertainty principle”: increasing one reduces the other (Craig & Brown 1986).

In the zero- and first-order regularization approaches, the regularized solution is shown as a confidence strip (CS) or solution “river.” This is created by generating solutions $\bar{F}(E)$ from random Gaussian realizations of $I(\epsilon)$ out to $\pm 3\sigma$ (i.e., rejecting any random realization of a data point that exceeds these limits) and plotting the extremes of these solutions. This CS river then defines the bounds within which almost all ($\sim 99\%$) *suitably smooth* solutions compatible with the data will lie. Note, however, that nonsmoothed solutions lying inside the CS river

are *not* acceptable [e.g., $\bar{F}(E)$ lying along the top bank for $E < E_1$ and along the bottom bank for $E \geq E_1$], since they violate the smoothness constraint and may, in some cases, give unacceptable residuals. Likewise, the resolution of a regularized solution is a somewhat technical issue and is closely linked to the solution uncertainty. These issues, to be discussed in more detail in a future paper, are briefly touched on in the Appendix. Here we simply note that, in the zero-order regularization method, the resolution ΔE was estimated from the half-width between the zeros of the highest order singular function used in the recovered solution (Bertero et al. 1988), while for first-order regularization the FWHM of the rows of the *resolution matrix* (Moore 1920; Penrose 1955; Press et al. 1988; Mankat 1989) were used. Although these give somewhat different answers, they are of

comparable order. We have superimposed selected illustrative horizontal lines of width ΔE on the CSs in Figure 2 to indicate the resolution as a function of electron energy E .

5.2. Fidelity of Recovered Spectra

All methods do a reasonable job of determining the rather simple spectral form of case A, the forward-fit method averaging out the knee in $\bar{F}(E)$, replacing it with a spectral index break. In case B, the forward-fitting method gives reasonable agreement with the upper energy end of the spectrum, but it fails to capture the increasing form of $\bar{F}(E)$ at low energies, the best approximation achievable by the parametric form being to impose a low-energy cutoff at ~ 20 keV to bring the hard X-ray intensity at low energies toward agreement with the simulated data. In the data-adaptive matrix binning method, the $\bar{F}(E)$ points are determined by row elimination in order of decreasing energy, starting at a high-energy point (here 1 MeV) where the photon spectrum above $\epsilon_{\max} = 150$ keV is a power-law extrapolation of the observed data. While some inference on the electron spectra above ϵ_{\max} is possible, it is generally assumed to be unreliable in detail so $\bar{F}(E)$ is retained only for $E \leq \epsilon_{\max}$, making it impossible for the method to detect the upper energy cutoff in $\bar{F}(E)$. On the other hand, the two regularization algorithms determine $\bar{F}(E)$ over the entire range. The zero-order method does a remarkably good job of finding the upper energy cutoff at 200 keV (the recovered value is 205 keV), while the first-order regularization method overestimates this cutoff somewhat, due to the high penalty it imposes on rapid variations in slope.

Case C presented the most difficult challenge. Because of the steeply rising electron spectrum below 30 keV, the bulk of the photon emission at, say, 15 keV is dominated not by electrons with an energy slightly greater than 15 keV (as it would if the electron spectrum were steeply decreasing), but rather by electrons near 30 keV. Thus, points in the electron spectrum below 30 keV contribute little to the observed photon spectrum and so are progressively more poorly determined as one goes below 30 keV. The forward-fit method simply truncates the spectrum at ~ 27 keV to avoid an overproduction of low-energy photons; the inversion methods fare quite well until just above 20 keV, at which point the errors become unacceptably large and the solution becomes unphysically negative.

Case D is characterized by two main features—a steep thermal (but not isothermal) low-energy component, and a “notch” from 42 to 50 keV. The regularized inversion methods, with their requirement of spectral smoothness, smear the shape of the notch feature from rectangular to Gaussian-like, while the matrix binning method recovers the location of the notch feature only at the expense of requiring rather coarse binning, so that the feature is characterized by only a couple of energy points. Not surprisingly, the limited set of functions allowed by the parametric form of the forward-fit method cause this feature to be completely ignored.

From the results, it is apparent that the presence of a strong low-energy thermal component does not detract from the ability of the inversion methods to determine features in the high-energy part of the spectrum (as long as these are appreciably beyond the “bulk” of the thermal component). All three inversion methods also determine well the (nonisothermal) form for the low-energy component of $\bar{F}(E)$. [This does not imply that the emission measure $\xi(T)$ differential in T is well determined.]

The single-temperature-plus-broken-power-law forward-fit model does a respectable job of finding a single representative temperature ($T = 2.3$ keV) for the thermal component. Interestingly, this is just *below* the low cutoff temperature $T_0 = 2.7$ keV in the actual $\xi(T)$ distribution (similar results were obtained in

cases E and F, with the recovered T being 81%–84% of T_0 in each case). However, the residuals between the observed photon spectrum in case D and that calculated using the recovered $\bar{F}(E)$ are unacceptably large (up to almost 10σ in the range from 30 to 40 keV), so that the poor fit to the low-energy end of the spectrum is well established. Whether the form of these photon spectrum residuals contains sufficient information to hint at the actual (nonisothermal) form of $\bar{F}(E)$ remains to be determined. While inversion algorithms recover the low-energy (nonisothermal thermal) part of $\bar{F}(E)$ rather well, the problem of converting the form $\bar{F}(E)$ into the corresponding form of $\xi(T)$ is a formidable challenge (Prato et al. 2005).

Case E is similar to case D, but with a sharp drop plus plateau in lieu of the notch. The quality of the recoveries is generally less good than that in case D. The first-order regularized method (red “river”) smooths through the plateau region from 55 to 70 keV, showing only some inflection and systematically overestimating the electron flux there and then, consequently (to compensate for the excess photons created by this overestimate), *underestimates* the electron flux from ~ 70 to 150 keV. Both the zero-order regularized method and the matrix binning method show more sign of a steep decrease in $\bar{F}(E)$ at ~ 55 keV and some sign of flattening of $\bar{F}(E)$ from 55 to 70 keV, but within the errors, no method ensures confident inference of the true form of $\bar{F}(E)$ in this range. The differences between methods are again simply a reflection of the fact that there is no unique answer to an inverse problem with noisy data, and the solution one gets varies according to the prior information/assumptions added.

Case F is characterized by a medium-energy “dip” in $\bar{F}(E)$ reminiscent of that found by Piana et al. (2003) in their analysis of the 2002 July 23 solar flare. All the inversion methods do faithfully recover this dip, while the parametric limitations of the forward-fitting model again necessarily forces a bad fit there. However, it should be noted that since the photon residuals for this forward fit are unacceptably large, they warn the modeler of the inadequacy of the parametric form chosen.

Overall, then, the “inversion” methods far outperform forward-fitting methods in their ability to recover sharp spectral features—a limitation intrinsic to forward-fitting methods, which by their nature can recover only the overall shape of the underlying spectrum, within the limits of the form used.¹² However, in practice, forward-fit methods can still reveal the *presence* of features outside their parametric form through examination of the pattern of photon spectrum residuals. Figure 3 illustrates this for case F. While the residuals from the inversion methods are acceptably small everywhere, with little evidence of strong clustering, the residuals from the forward-fit method are vastly greater and indeed cluster at negative values (of order -5σ) in the vicinity of 30 keV, the location of the dip in $\bar{F}(E)$. The form of the cumulative residual function (Fig. 3, *bottom*) is only marginally acceptable at the 3σ level. On examination, this behavior is seen to be a consequence of the steep isothermal form in the forward fit extending to quite high energies in the forward fit (before it jumps up to the high-energy broken power-law form), which renders the recovered spectrum in the 30–40 keV range much lower than the actual form. Recognizing this, an aptly modified parametric form of the forward-fit function might thus be deduced and applied to the data.

It must be noted that acceptable patterns of residuals (in addition to acceptable values of χ^2) are generally obtained in forward fits to solar flare data analyzed so far (e.g., Holman et al.

¹² Similar arguments apply to the inference of spatial features in image reconstruction algorithms (Bertero & Boccacci 1998).




FIG. 3.—Photon residuals for case F. The top panel shows the residuals (data-reconstructed photon spectrum) as a function of photon energy, normalized to the Poisson standard deviation at each photon energy, for two recovery methods—zero-order regularization (*light*) and forward fitting (*dark*). The middle panel shows the distribution of residual sizes for each method; each is Gaussian in form with a standard deviation of approximately unity. The bottom panel shows the behaviors of the “cumulative residual” (the average of the residuals from the low energy end of the spectrum to the energy in question) as a function of ϵ , as discussed by Piana et al. (2003); an acceptable behavior should lie within the 3σ bounds shown by the dashed lines. While the residuals for the regularized inversion show acceptable behavior, those for the forward fit do not. Specifically, both the pattern of residuals and the cumulative residual plot reveals a high degree of clustering of residuals around 30 keV—the vicinity of the dip in the actual $\bar{F}(E)$ spectrum. This shows that the forward fit $\bar{F}(E)$ is not an acceptable recovery of the true mean electron source spectrum and even suggests where adjustments should be made in order to obtain a better fit.

2003), including the few cases found in which the $\bar{F}(E)$ determined through regularized inversion methods seems to involve a significant dip (e.g., Piana et al. 2003). The absence of obvious contrary features in the pattern of residuals argues against the presence in these real solar events of sharp features, such as those being investigated here, of amplitude large enough to be detected above current levels of data uncertainty, although it should be noted that the assumed presence (Holman et al. 2003) of systematic uncertainties over and above the Poisson uncertainty level will increase the reference noise value and thus lower the magnitude of the normalized residuals.

A feature we found in the results using inversion methods is that the form of the recovered mean source electron spectrum depends slightly on the photon energy to which one ascribes the photon flux in a given bin; methods that use the low-energy end of the bin (on the grounds that for steeply decreasing spectra the average photon energy is near the lower bound) obtain results that are slightly different from methods that use the midpoint

energy of the photon bin. Of course, neither of these methods is correct except in special cases—the average energy in a given bin is a weighted average of the (a priori unknown) spectrum across that bin. This problem of “where to plot the points in a hard X-ray spectrum” has been known for years in the context of low-resolution scintillator spectra (Dennis 1978), and it is interesting that it is still notable in the era of keV-resolution hard X-ray spectra.

6. CONCLUSIONS

We have demonstrated that various methods in use for recovering the model-independent mean bremsstrahlung source electron spectrum $\bar{F}(E)$ are dependable in recovering various large- and small-scale spectral features, albeit to varying degrees. Which of these methods should be applied to observed spectra depends to a large extent on the type of information desired.

The Tikhonov regularization methods generally recognize and recover most of the interesting spectral features, including high-energy cutoffs, dips, and bumps. The adaptive-binning matrix inversion method faithfully recovers features within the energy range defined by the input photon spectrum but, by its nature, cannot determine the form (or even existence) of features outside this range. The adaptive-binning matrix inversion method recovers certain features of the input data (e.g., the low-energy end of case B, the plateau of case E, and the notch in case F) particularly well; this would appear to be a consequence of the local smoothing involved as compared with the global smoothing in the regularization algorithms. Very sharp features (such as cutoffs) are generally less well fit by first-order regularization methods because of the requirement for a smooth first derivative.

Parametric forward-fit methods, by their very design, cannot recover such features unless they are built into the parametric form. However, analysis of the photon spectrum residuals can provide clear indication of the inadequacy of the assumed functional form and may even suggest locations and/or likely amplitudes of additional features. Forward-fit methods should be adequate when either a gross indication of the spectral form, or a quantity related to its integral (e.g., the electrical current, or the injected energy flux), are desired. However, to obtain the detailed shape of the spectrum (of interest, e.g., in searching for sudden transitions in form associated with a candidate acceleration or transport mechanism), the use of overall regularized inversion methods or data-adaptive binning is essential. Codes to perform all of these recovery techniques, written in IDL, will very soon be available on the Solar Software Tree.

This work was supported by NASA’s Office of Space Science through Grant NAG5-207745, by a PPARC Rolling Grant, and by the International Space Science Institute.

APPENDIX

RESOLUTION AND UNCERTAINTY IN REGULARIZED SOLUTIONS

For practical purposes of data interpretation, the confidence strips (CSs) in Figure 2 should be seen as delineating that part of the E , $\bar{F}(E)$ plane in which $\sim 99\%$ of solutions (with the imposed degree of smoothness) lie. Caution is needed in reading any more into the CS river than that and especially in considering how it relates to the resolution ΔE and the uncertainty $\Delta \bar{F}$ of the solution. These issues will be discussed fully in a future paper; here we simply give some cautionary pointers.

1. ΔE and $\Delta \bar{F}$ are *not* independent. If we reduce ΔE (by binning or by reducing smoothing parameter λ) to improve resolution, then the uncertainties $\Delta \bar{F}$ worsen. On the other hand, if we sacrifice energy resolution, we can get a much smaller uncertainty on a mean value of $\Delta \bar{F}$ over a wide energy range. The tradeoff between these two is subjective.

2. In the case of the adaptive binning method, the resolution ΔE is *local* and fixed by the criterion chosen for the acceptable level of error $\Delta \bar{F}$. Thus, as can be seen in Figure 2, ΔE and $\Delta \bar{F}$ can vary considerably from bin to bin, depending on data noise and on the true shape of $\bar{F}(E)$. On the other hand, for the regularization methods, the smoothness constraint is applied *globally*, i.e., to the whole solution, giving the CS river of globally smoothed solutions a relatively small value of $\Delta \bar{F}$ compared to the error in the (locally smoothed) adaptive binning solution.

3. The resolution of, and uncertainty in, a regularized solution can be defined in various ways. Both are affected by the level of noise in the data and by the degree of filtering action imposed by the kernel (Q here). Setting a value for the acceptable uncertainty $\Delta \bar{F}$ basically fixes the resolution ΔE in terms of the properties of Q and the value of the smoothing parameter λ (eq. [4]). For the zero-order regularization method, the value of λ determines the truncation of the singular function expansion of the solution and a natural measure of resolution is the separation of the zeros in highest order singular function used (Bertero et al. 1988). For the first-order regularization method, the FWHM of the rows of the “resolution matrix” (Moore 1920; Penrose 1955; Press et al. 1988; Mankat 1989) were used. This matrix represents the combined properties of the unconstrained matrix problem and the action of whatever smoothing algorithm is used.

4. It is tempting to interpret the “height” of the CS river as the uncertainty in the solution at that point, and the “width” of the river (in E space) as the resolution of the method. This is not the case, however, since the CS height and width depend on the true shape of $\bar{F}(E)$. For example, if $\bar{F}(E)$ were constant, then the CS river height would indeed indicate the uncertainty in $\bar{F}(E)$, but its width would be the entire solution energy range, a quantity far greater than the achievable resolution ΔE , i.e., the width of the narrowest feature whose presence could be inferred. Similarly, if $\bar{F}(E)$ had a large local discontinuity, then the height of the CS river at that point would be far larger than the noise error in $\bar{F}(E)$.

5. On the other hand, the uncertainties ($\Delta \bar{F}$) in the adaptive binning method *do* have the traditional meaning with the caveat that what is really determined in the method is the product $\bar{F}(E)\Delta E$ in each energy bin. Thus, for example, if the $\Delta \bar{F}$ correspond to 1 σ error bars, the true value of $\int \bar{F}(E)dE$ in each energy bin is within the range $(\bar{F} \pm \Delta \bar{F})\Delta E$ of the recovered spectrum 68% of the time.

6. In forward fitting, the concepts of resolution and error do not have the same meaning since a prescribed functional form is assumed. However, it is possible to estimate confidence intervals for the model parameters and hence a strip in the E, \bar{F} plane by the modified χ^2 method given for example, by Simmons et al. (1982). Basically, recognizing that the model does not in fact come close to fitting the data within data errors, one fits it within “model error” bars estimated from the variance in the fit itself. Such a method has been applied by Kasparová et al. (2005) in the solar flare case.

REFERENCES

- Bai, T., & Ramaty, R. 1978, ApJ, 219, 705
 Bertero, M., & Bocacci, P. 1998, Introduction to Inverse Problems in Imaging (Bristol: IOPP)
 Bertero, M., De Mol, C., & Pike, E. R. 1988, Inverse Problems, 4, 573
 Brown, J. C. 1971, Sol. Phys., 18, 489
 Brown, J. C., & Emslie, A. G. 1988, ApJ, 331, 554
 Brown, J. C., Emslie, A. G., & Kontar, E. P. 2003, ApJ, 595, L115
 Craig, I. J. D., & Brown, J. C. 1986, Inverse Problems in Astronomy (London: Adam Hilger)
 Dennis, B. R. 1978, OSO Int. Tech. Memo. OSO-8
 Golub, G. H., & van Loan, C. F. 1993, Matrix Computation (Baltimore: Johns Hopkins Univ. Press)
 Holman, G. D., Sui, L., Schwartz, R. A., & Emslie, A. G. 2003, ApJ, 595, L97
 Johns, C., & Lin, R. P. 1992, Sol. Phys., 137, 121
 Kasparová, J., Karlický, M., Kontar, E. P., Schwartz, R. A., & Dennis, B. R. 2005, Sol. Phys., 232, 63
 Koch, H. W., & Motz, J. W. 1959, Phys. Rev., 31, 920
 Kontar, E. P., MacKinnon, A. L., Schwartz, R. A., & Brown, J. C. 2006, A&A, 446, 1157
 Kontar, E. P., Piana, M., Massone, A. M., Emslie, A. G., & Brown, J. C. 2004, Sol. Phys., 225, 293
 Langer, S. H., & Petrosian, V. 1977, ApJ, 215, 666
 Lin, R. P., et al. 2002, Sol. Phys., 210, 3
 Mankat, W. 1989, Geophysical Data Analysis: Discrete Inverse Theory (New York: Academic Press)
 Massone, A. M., Emslie, A. G., Kontar, E. P., Piana, M., Prato, M., & Brown, J. C. 2004, ApJ, 613, 1233
 Moore, E. H. 1920, Bull. Am. Math. Soc., 26, 394
 Morozov, V. A. 1966, Soviet Math. Dokl., 7, 414
 Penrose, R. 1955, Proc. Cambridge Philos. Soc., 52, 17
 Piana, M. 1994, A&A, 288, 949
 Piana, M., & Brown, J. C. 1998, A&AS, 132, 291
 Piana, M., Massone, A. M., Kontar, E. P., Emslie, A. G., Brown, J. C., & Schwartz, R. A. 2003, ApJ, 595, L127
 Prato, M., Piana, M., Emslie, A. G., Kontar, E. P., Massone, A. M., & Brown, J. C. 2005, Sol. Phys., submitted
 Press, W. H., Saul, A., Teukolsky, S. A., Wetterling, W. T., & Flannery, B. P. 1988, Numerical Recipes in C (1st ed.; New York: Cambridge Univ. Press)
 Simmons, J. F. L., Aspin, C., & Brown, J. C. 1982, MNRAS, 198, 45
 Thompson, A. M., Brown, J. C., Craig, I. J. D., & Fulber, C. 1992, A&A, 265, 278
 Tikhonov, A. N. 1963, Soviet Math. Dokl., 4, 835