4 Data Analysis

4.1 Techniques of DEM Recovery

New data demand new analysis techniques. The MSSTA is exciting because it produces a dataset fundamentally unlike any other currently available: a set of 5 calibrated narrowband EUV spectroheliograms. However, the novelty of the data makes interpreting these results challenging. The DEM extraction techniques we employ on the MSSTA data cannot be tested on a well-studied dataset because no dataset suitable for such an investigation exists. There have been efforts to extract spatially-resolved DEMs from CDS observations (Schmelz, Scopes et al. 2001); however, the rastered-slit spectrograms produced by CDS are fundamentally quite different from the narrowband full-disk images obtained by the MSSTA. The inability of TRACE and EIT observations alone to constrain a DEM is discussed at length in Section 1.3.2.

4.1.1 SIMULATING OBSERVATIONS

Therefore, in order to refine the DEM extraction technique and establish realistic expectations for its performance, we must begin by producing artificial datasets consisting of simulated observations based on a realistic instrument model and a reasonable assumption about the state of the corona. Running our analysis algorithm on these synthetic observations and comparing the extracted DEM to the input DEM will accomplish some important goals:

- First, working with simulated observations will establish whether DEM extraction is possible at all with the kind of data we have obtained.

- Assuming the basic concept is viable, we can test a number of approaches to DEM recovery and determine which works best. The detailed parameters of that approach can be fine-tuned in order to speed up its execution, enhance its reliability or improve its flexibility.

- We will be able to test the sensitivity of the analysis to unexpected problems in the data set – calibration errors, gaps in the coverage of one or more bandpasses, etc. Conversely, we can explore how improving the data by more accurate calibration or additional bandpasses might improve the breadth and accuracy of the results.

- Most important of all, we will be able to see whether the results of our analysis are trustworthy. Simply putting error bars on the recovered DEM is not enough; we need to get an intuitive sense for whether DEMs obtained in this way are stable, consistent, and realistic descriptions of the conditions underlying our dataset.

Ideally, an artificial dataset would consist of synthetic images in each of our multilayer bandpasses of a region of the corona, based on a three-dimensional model of the temperature structure of that region (Weber, DeLuca et al. 2004). However, the complexity involved in generating such a model is substantial, and in fact such an artificial dataset would not give us any more information about our ability to recover DEMs from a set of pixel values than we can gain from a much simpler approach. Instead, we start by assuming a single pixel. Without fully enumerating the distribution of plasma within the column-of-sight defined by that pixel, we characterize that pixel with an ideal DEM. (As noted in section 1.3.4, a fundamental conceit of analyzing multilayer images is that the DEM contains all the information about the state of the plasma along a line of sight that is needed to calculate the intensity of EUV emission from that plasma.) Using the assumed DEM, we calculate the pixel value found at that point in each of the images in the MSSTA dataset. This artificial dataset can then be fed into a variety of candidate DEM extraction algorithms to produce a recovered DEM. The recovered DEM is compared to the input DEM to test the accuracy of the algorithm.

A number of different input DEMs were chosen in order to reproduce the range of possible observed pixel values expected from the flight dataset. Most tests were run using one of three

standard DEMs (active region, bright quiet-sun network, and coronal hole) from the CHIANTI package, derived from the observations of (Vernazza and Reeves 1978). These are all broad, multithermal temperature distributions based on spectrograms averaged over large areas on the solar disk; all have been used extensively in the solar literature and are widely accepted as being characteristic of solar plasma in a general class of feature. In addition to these three broad DEMs, a number of artificial datasets were generated using sharply-peaked DEMs, in order to evaluate the algorithms' response to an image of an isothermal plasma. Since the question of whether broad or narrow DEMs are the norm for coronal loops is one of the central problems that we hope to solve using high-resolution multi-spectral observations, it is important to be sure that the DEM extraction technique is capable of reproducing an input distribution with any shape.

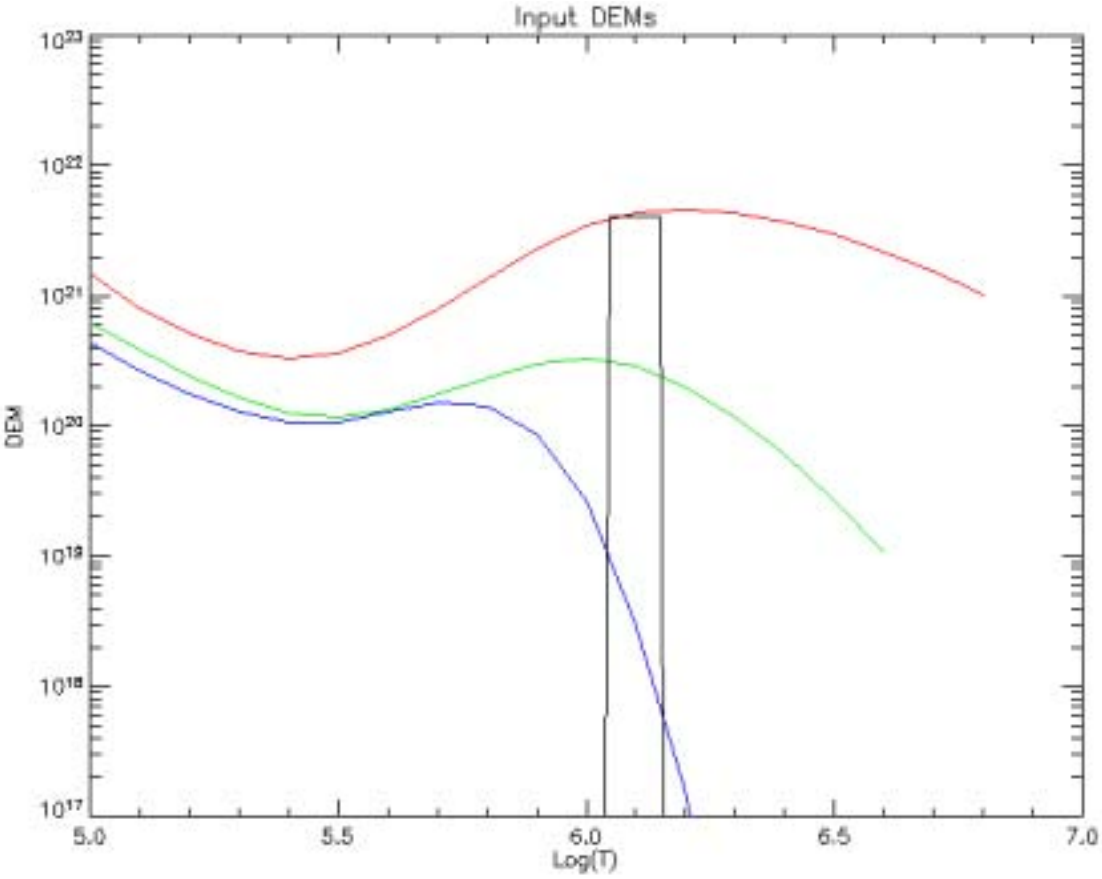

Figure 60. The DEMs used to simulate observations in order to test the data analysis techniques. The colored curves represent standard averaged active region (red), bright quiet sun (green) or coronal hole (blue) type emission. The black curve represents an isothermal coronal loop based on loop 1 from (Aschwanden, Newmark et al. 1999). Its DEM is centered at 1.25 MK; the width of the peak is 0.1 in Log(T), corresponding with the temperature precision of the line contribution functions that go into the DEM. Similar isothermal DEMs at other temperatures were occasionally used to ensure that the algorithm didn't have substantial "blind spots" in temperature space.

Of course, the exact shape of the isothermal DEM is somewhat arbitrary; a truly isothermal plasma would have a delta-function DEM, but asking a general DEM-fitting routine to recover such a function is unrealistic. The temperature precision of the DEM is only 0.1 in Log(T); that is the step size at which the atomic data in CHIANTI is tabulated. Interpolating to a finer temperature grid does not increase the accuracy of the extracted function. Thus, when we generated isothermal input DEMs, we used box functions with a width of 0.1 in Log(T), normalized to a reasonable isothermal emission measure. The "isothermal" DEM shown in Figure 60, chosen to reproduce the observations of (Aschwanden, Newmark et al. 1999), is centered at Log(T) = 6.1 and normalized to an emission measure of $\sim 5 \times 10^{27}$ cm^{-5}.

The observations expected from a pixel with each of these input DEMs are listed in Table 17. For the 9 MSSTA EUV telescopes, the predicted observations are in units of ergs cm^{-2} s^{-1}, the flux units of MSSTA pixels after film calibration has been applied to the image. For the EIT observations, the units are in DN (counts) per second. Comparing these predictions with those based on the pre-flight calibration of the EIT made in Table XV of (Dere, Moses et al. 2000) reveals discrepancies of up to a factor of 2. These differences are due to different assumptions made in computing the instrument temperature kernels; I recalculated the EIT kernels using a lower assumed pressure, and updated abundances, ionization equilibria, and line strengths thanks to enhancements to the CHIANTI database (Young, Del Zanna et al. 2003). As a result, the EIT observations predicted in Table 17 offer somewhat better agreement with actual EIT data[3] (including the 1996 and 1998 observations mentioned in Dere et al.'s Table XV) than the predictions of (Dere, Moses et al. 2000). This should serve as a reminder of the sensitivity to detailed atomic physics that underlies any analysis of EUV emission by coronal plasmas; such a reminder is needed, since the original temperature kernels presented in (Dere, Moses et al. 2000) are still widely used in the literature.

[3] The exception is the 304 Å bandpass, which is not entirely susceptible to temperature kernel analysis. It is included here for completeness; Dere et al. used an *ad hoc* modification to the 304 Å temperature kernel for their predictions.

Bandpass Central Wavelength [Ångstroms]	Active Region	Quiet Sun	Coronal Hole	Isothermal Loop
MSSTA Telescopes	Predicted Pixel Value (10^{-2} ergs cm^{-2} s^{-1})			
150	6.0	0.3	0.1	0.6
171	254.0	28.1	8.1	47.0
180	165.1	12.3	1.3	61.5
195	180.5	9.7	0.4	55.6
211	65.8	2.8	0.1	9.9
131	238.3	18.8	9.5	13.1
98	27.1	2.0	0.5	4.2
58	3.6	0.1	0.0	0.7
256	1.7	0.1	0.0	0.3
EIT Telescopes	Predicted Pixel Value (DN s^{-1})			
171	464.9	47.0	11.3	106.9
195	358.4	21.6	1.7	119.1
284	10.1	0.4	0.1	0.5
304	13.3	0.9	0.3	1.7

Table 17. Predicted pixel values based on the input DEMs shown in Figure 60. These simulated observations were used to test DEM reconstruction techniques.

Extracting a DEM from a perfectly clean set of synthetic observations like those in Table 17 is not difficult; the real test is handling noisy observations. I simulated a noisy dataset using a Monte Carlo technique: a random number was added to each of the predicted pixel values before they were passed to the DEM extraction procedure. The random number was chosen from a Gaussian distribution centered at zero, with a FWHM set by the overall noise level times the predicted pixel value. (Thus, when simulating observations with 25% error, most of the pixel values in the synthetic dataset would be within about ±25% of their true predicted value.) In most cases, several hundred different noisy datasets were created for each true dataset, and the DEMs extracted from all of these Monte Carlo realizations were compared.

While straightforward, a Gaussian description of noise and error is not entirely accurate. As described in Section 2.6.2, the overall error on the MSSTA observations is 25-35%, but much of this error comes from systematics. Only a small part is due to sources that might be accurately modeled by a Gaussian distribution centered at 0 (although film grain and some aspects of development chemistry can add what looks like Gaussian noise to a small, pixel-sized sample of a

negative). The same is true for EIT. Systematic errors in calibration and in underlying atomic physics generally predominate over the statistical noise found in EUV spectroheliograms. Unfortunately, modeling systematic error is tricky. One approach is to add an offset to each pixel value before generating the randomized noisy datasets. However, the number of Monte Carlo realizations needed to fully explore the effect of systematic errors in this way is unreasonably large. For the most part, the simulated datasets included only Gaussian noise at a realistic overall level, while the effect of treating that noise as systematic error instead was briefly studied after the fact.

4.1.2 MATRIX INVERSION

With a goal (the input DEM) and a few hundred sets of simulated, noisy observations in hand, we can now test our ability to reconstruct a DEM from MSSTA data. We begin by re-examining Equation 7 from Section 1.3.3:

(Equation 7) $$p_i(\mathbf{x}) = \int_0^\infty K_i(T) DEM(T, \mathbf{x}) \mathrm{d}T \qquad [\text{counts s}^{-1}]$$

This is really a set of N equations, where N is the number of telescopes contributing to the dataset. For each telescope i (i = 1, …, N), the pixel value is found by integrating the temperature kernel of the ith telescope times the DEM over temperature. In order to solve this set of equations for the DEM, the most obvious approach is to discretize the problem in temperature. We define a set of temperatures T_j (for simplicity, we choose j = 1, …, N), changing the DEM from a continuous function of temperature to a set of tabulated values. The set of temperature kernels now becomes an N x N temperature response *matrix* describing the sensitivity of the entire instrument package, and the pixel values are determined by summing over temperature, rather than integrating:

(Equation 19) $$p_i(\mathbf{x}) = \sum_{j=1}^{N} K_{ij} DEM_j(\mathbf{x}) \qquad [\text{counts s}^{-1}]$$

Inverting such a simple system of linear equations is tantalizingly easy. We invert the square temperature response matrix by Gaussian elimination (using IDL's INVERT routine) to obtain the inverse matrix K_{ij}^{-1}, such that

(Equation 20) $$DEM_j(\mathbf{x}) = \sum_{i=1}^{N} K_{ij}^{-1} p_i(\mathbf{x}) \quad [\mathrm{cm}^{-5}\ \mathrm{K}^{-1}]$$

Thus, in principle, any set of observations can be turned into a DEM almost instantaneously, with a simple matrix multiplication.

A quick test of this method using noise-free simulated data gave very promising results. Since the MSSTA dataset consists of images in 5 EUV bandpasses, we defined a discrete temperature grid with 5 points T_j such that $\mathrm{Log}(T_j) = 5.35 + 0.2 \times j$. The instrument response matrix K_{ij} was generated by sampling the temperature kernel of each instrument i at all temperatures j. This matrix was then inverted, and a set of simulated MSSTA observations generated using the CHIANTI active region DEM was multiplied by the inverse matrix to yield DEM values at each of the temperature grid points. This output DEM was compared to the input active region DEM; the results, shown in Figure 61, are accurate to about 10%. Similar agreement was found with quiet region and coronal hole DEMs; even isothermal DEMs were recovered fairly well, although those results were quite sensitive to the spacing of the temperature grid.

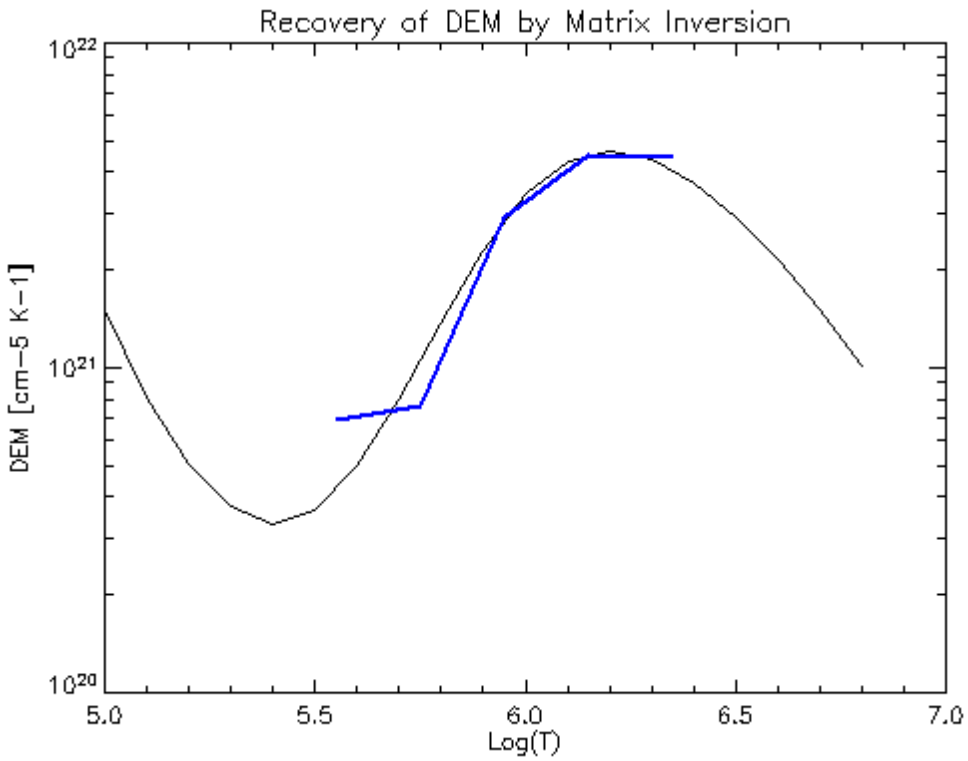

Figure 61. An active region DEM (black curve) was used to generate noise-free simulated observations in the 5 MSSTA EUV channels; these simulated observations were then multiplied by the inverse temperature kernel matrix to give a recovered DEM. The result agrees with the input DEM to ±10% over the temperature range sampled by the MSSTA.

However, the agreement quickly deteriorates when we add noise to the observations. As shown in Figure 62, even a very small amount of noise (5%) results in a recovered DEM that is generally only close to the input DEM at one or two temperature points; with more realistic noise levels (25%), the solution becomes wildly oscillatory, and utterly unreliable. Typically, the recovered DEM takes on unphysical negative values at one or more points. Including additional bandpasses (like the EIT instruments, or the temperature kernels of the MSSTA telescopes whose filters failed on flight) in the kernel matrix does not help; indeed, a large number of bandpasses forces us to use a finer temperature grid (or extend it over a broader temperature space), which tends to increase the instability of the solution.

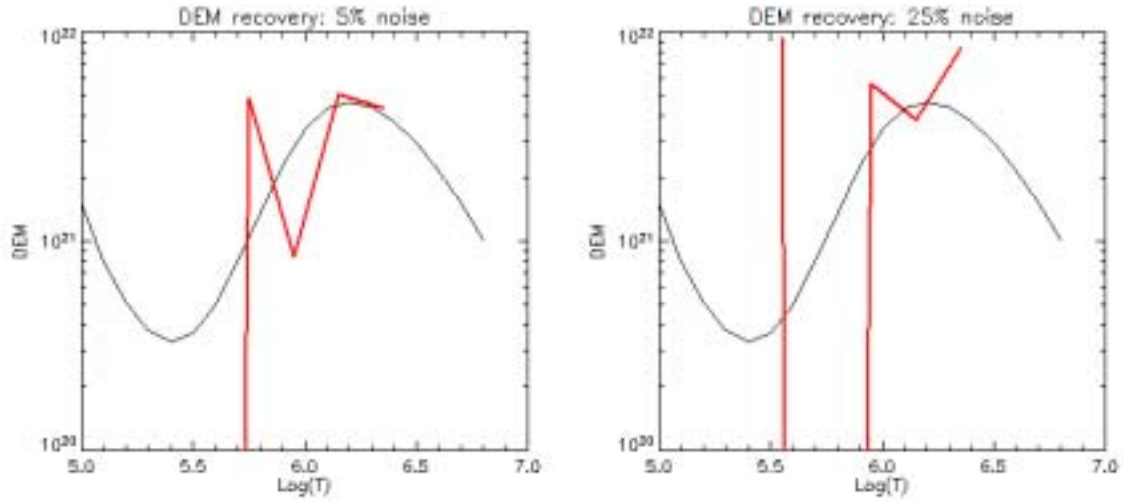

Figure 62. Adding noise to the simulated observations before inverting them quickly destroys the accuracy of the recovered DEM. The result (shown in red) is unstable under even a small amount of Gaussian noise (left), frequently going negative as the noise level increases (right).

The problem is that the temperature response matrix K is very badly-conditioned. We can quantify this ill-conditioning, and gain some insight into what it means and how to deal with it, by looking at the singular value decomposition (SVD) of the matrix. Following the derivation by (Hansen 1998), we define the SVD as

$$K = U\Sigma V^T = \sum_{i=1}^{N} \mathbf{u}_i \sigma_i \mathbf{v}_i^T$$

where $U = (\mathbf{u}_1, \ldots, \mathbf{u}_N)$ and $V = (\mathbf{v}_i, \ldots, \mathbf{v}_N)$ are matrices of orthonormal columns and $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_N)$. The SVD is similar to an eigenvalue decomposition of the matrix KK^T; it has the following general properties that are useful for our case:

- The σ_i are the singular values of the matrix; they decrease gradually and monotonically to 0 with increasing i. The condition number of the matrix σ_i / σ_N is generally quite large; increasing the dimension of the matrix tends to increase the number of small singular values.

- The $\mathbf{u}_i$ and $\mathbf{v}_i$ are the singular vectors of the matrix; they can be thought of as a set of basis observations and basis DEMs, onto which the actual or simulated observations and measured DEM are projected. They tend to have more sign changes as i increases, becoming highly oscillatory at large i.

If we now consider Equation 19 as a mapping of a vector of observations **p** onto a DEM vector **d** (leaving out the indices from Equation 19, which serve only to make the matrix operations explicit; note that the index *i* now refers to a particular singular value, rather than a particular instrument), we can write

$$\mathbf{p} = K\mathbf{d} = \sum_{i=1}^{N} \sigma_i (\mathbf{v}_i^T \mathbf{d}) \mathbf{u}_i$$

It is clear that the high-frequency components of the DEM vector **d** are damped out by multiplication with the small singular values; thus the observations are relatively insensitive to small fluctuations in DEM. However, when the problem is inverted, any high-frequency oscillations in the data **p**, such as those introduced by small-scale noise, are hugely amplified, more so as the condition number of the kernel matrix (the smallness of the high-order singular values) increases.

The singular values of the MSSTA temperature kernel matrix are plotted in Figure 63. The effect of choosing different combinations of instruments, and different temperature grids, is illustrated by the four different curves. The blue curves, based on the 5 successful MSSTA EUV telescopes, show that substantially better results are obtained when the discretization is defined over a fairly narrow range of temperatures; spreading the 5 bandpasses over a broad temperature range leaves only one large singular value, indicating that the information content of the resulting response matrix is quite small. The black curves include three extra bandpasses, but do not generally produce many more large singular values; most of the additional singular values are small, and thus only contribute to noise in the inversion. However, it is clear that, with more bandpasses, it is advantageous to consider a broader temperature range.

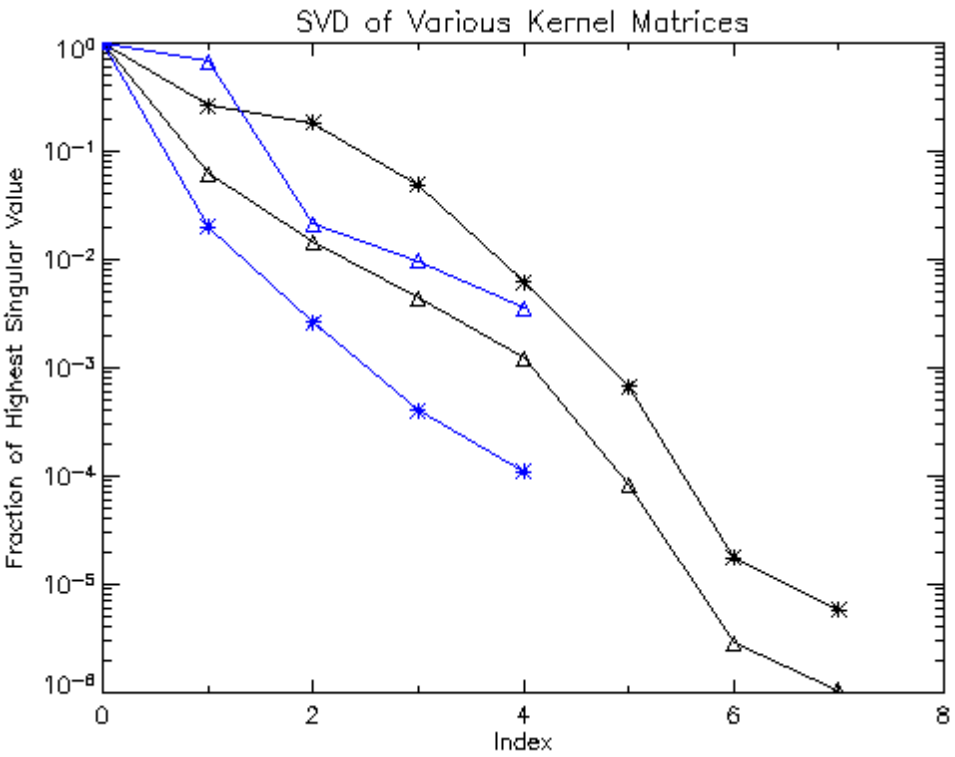

Figure 63. The SVD of the kernel matrix is sensitive to the temperature range and number of bandpasses used. Blue lines are for matrices that include the MSSTA 150 Å, 171 Å, 180 Å, 195 Å and 211 Å telescopes; black lines include these bandpasses as well as EIT's 171 Å, 195 Å and 284 Å telescopes. Triangles indicate kernels that use a temperature grid from Log(T) = 5.55 – 6.75; stars indicate kernels that were gridded over Log(T) = 5.35 – 7.15.

The SVD serves several important functions. First, simply examining plots like Figure 63 provides invaluable insight into the information content of the instrument response matrix. It enables us to fine-tune the temperature-discretization process in order to maximize the usefulness of our results, and it indicates the relevance of including various bandpasses. Studying the singular vectors $\mathbf{u}_i$ and $\mathbf{v}_i$ corresponding to large singular values can help us understand what sorts of DEMs an instrument is actually capable of discriminating. Such information could be used to optimize the design of future EUV instruments (indeed, it would have been helpful to look at the SVD during the selection of bandpasses for the MSSTA), and to constrain any analysis programs that use their data.

Furthermore, the SVD can serve as the stepping stone to a regularized solution to Equation 20, one that is still a unique inverse solution but that reduces some of the extreme noise-sensitivity illustrated by Figure 62 (Larsen, Kosovichev et al. 2004). A simple truncated SVD approach,

where the observations are projected onto the singular vectors $\mathbf{u}_i$ with large singular values (i =1, …, M, $M < N$) and a solution is constructed from the corresponding superposition of singular vectors $\mathbf{v}_i$, is certainly possible; slightly more complex schemes like Tikhonov regularization (Tikhonov and Arsenin 1977) can provide better approximations to the exact solution while minimizing unphysical oscillations in the DEM. Regularization of ill-conditioned matrices is a vast and richly-studied field (see (Hansen 1998) for an overview); however, while a regularized linear solution has been discussed by (Judge, Hubeny et al. 1997), among others, as a possible method for analyzing spectral data from SERTS or SUMER, most of the literature of coronal EUV observations relies on forward fitting techniques similar to the approach described in the following section.

My own efforts to construct a regularized solution to the problem of obtaining a DEM by inverting MSSTA observations are preliminary at best. The approach is certainly interesting, but enforcing the requirement that the singular vectors, or at least the resulting DEM, be non-negative has proven to be a considerable obstacle (though perhaps not an insurmountable one). The extracted DEMs are generally either trivial or completely unphysical, and thus uninteresting from a solar physics standpoint. I will continue to study the linear inversion approach, but for now will move on to discuss a more typical forward-modeling procedure that consistently generates believable solutions. This technique lacks some of the rigor of the linear inversion method, and does not guarantee a unique solution, but it will enable us to determine whether any solutions to the DEM problem can be found, and provide a glimpse of what those solutions might look like.

4.1.3 FORWARD FITTING

The forward approach to solving Equation 7 does not require that we discretize the problem in temperature and invert the equation to solve for DEM(T). Instead, we begin by examining our observations $p_i(\mathbf{x})$ and making a guess about the shape of the DEM that produced them. The guessed DEM is plugged into Equation 7 to predict a set of observations $p_i'(\mathbf{x})$, and the predicted observations are compared to the actual observations. Then the guessed DEM is iteratively modified to obtain the best possible agreement between the predicted and the observed pixel values.

In order to implement a forward-fitting solution, it is necessary to define a metric describing the agreement between predicted and observed pixel values. I used a standard χ^2:

$$\chi^2 = \sum_i \frac{(p_i' - p_i)^2}{p_i'}$$

An additional penalty function was applied to the χ^2 when the DEM became excessively oscillatory; however, the penalty function was not generally a factor in the analysis. It is also necessary to parameterize the DEM so that it can be modified in a systematic way. That is, we must model the DEM as a function of some set of parameters π_i (i = 1, …, N). Any particular combination of parameters then specifies a particular DEM, and thus a particular set of predicted observations, which lead to a value of χ^2. The fitting algorithm (I used the downhill simplex method as implemented in IDL's built-in AMOEBA function) then searches the N-dimensional space for the set of parameters that leads to the minimum value of χ^2. Defining the functional dependence of the DEM on the π_i is a subtle procedure; the choice of the number of parameters, and the functional form that turns those parameters into a DEM, has a strong effect on the ability of the fitting algorithm to find stable and unique solutions.

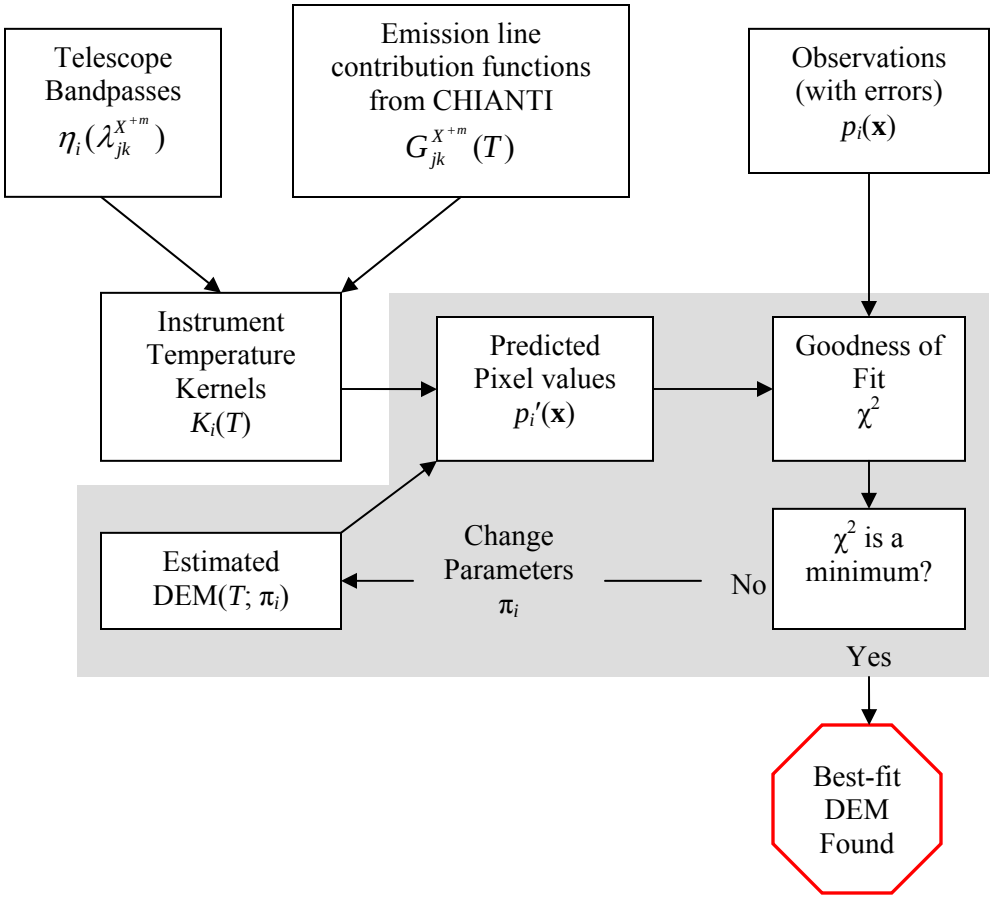

Figure 64. The forward-fitting procedure. Diagram modeled after one in (Golub, DeLuca et al. 2004).

After experimenting with polynomial and Fourier-series functions, I settled on defining the DEM with a cubic spline, where the parameters specified the DEM at temperature points equally spaced over the range of sensitivity of the MSSTA instruments. The spline is intuitive and flexible enough to reproduce almost any shape of curve, given enough grid points. The number of grid points specifies the dimension of the solution space. While it is tempting to use a number equal to or just less than the number of observations (see, *e.g.*, (Weber, DeLuca et al. 2004), who used N-1 spline points), it is important to remember the lessons about the poor conditioning of the instrument's temperature response that we learned by looking at the SVD of the kernel matrix. As noted in (Craig and Brown 1976), the contribution functions of EUV lines (and thus the temperature kernels of narrowband EUV instruments) are fundamentally broad in temperature, and thus the line intensity is relatively insensitive to small fluctuations in DEM. (Craig and Brown 1976) found that using more than 3-4 spline points was not justified by even the best

spectral observations. My testing of multi-point spline DEMs supports that conclusion; while a DEM based on 5 or 6 temperature grid points was better able to match a variety of input DEMs in the absence of noise, its performance quickly deteriorated under even small perturbations to the simulated observations.

Therefore, three-point splines were used to construct the model DEMs during the fitting procedure. I fine-tuned the DEM parameterization by generating simulated observations with a range of input DEMs and testing the ability of the model DEMs to reproduce these observations. At this stage, the location of the spline points and the domain over which the model DEM was defined were varied by hand until a good compromise between breadth and stability was found. The fitting algorithm was disturbingly sensitive to the location of the spline points; however, a good compromise was found by defining the model DEM over the interval Log(T) = 5.3 – 6.8, with spline points at Log(T) = 5.6, 6.0 and 6.4. The endpoints (*i.e.* the DEM values at Log(T) = 5.3 and 6.8) were forced to be lower than the adjacent spline points in order to prevent the model from diverging to a high value outside of the range over which the MSSTA instruments were sensitive. The results are shown in Figure 65. Clearly, this model is quite capable of matching any reasonable input DEM, including isothermal DEMs, over the range of Log(T) = 5.7 – 6.4. That such good agreement can be achieved with only three free parameters is quite encouraging; the fitting algorithm's sensitivity to noise increases greatly with the number of free parameters, so adding an extra spline point just to improve the fit at the low-temperature end of the range would defeat the purpose.

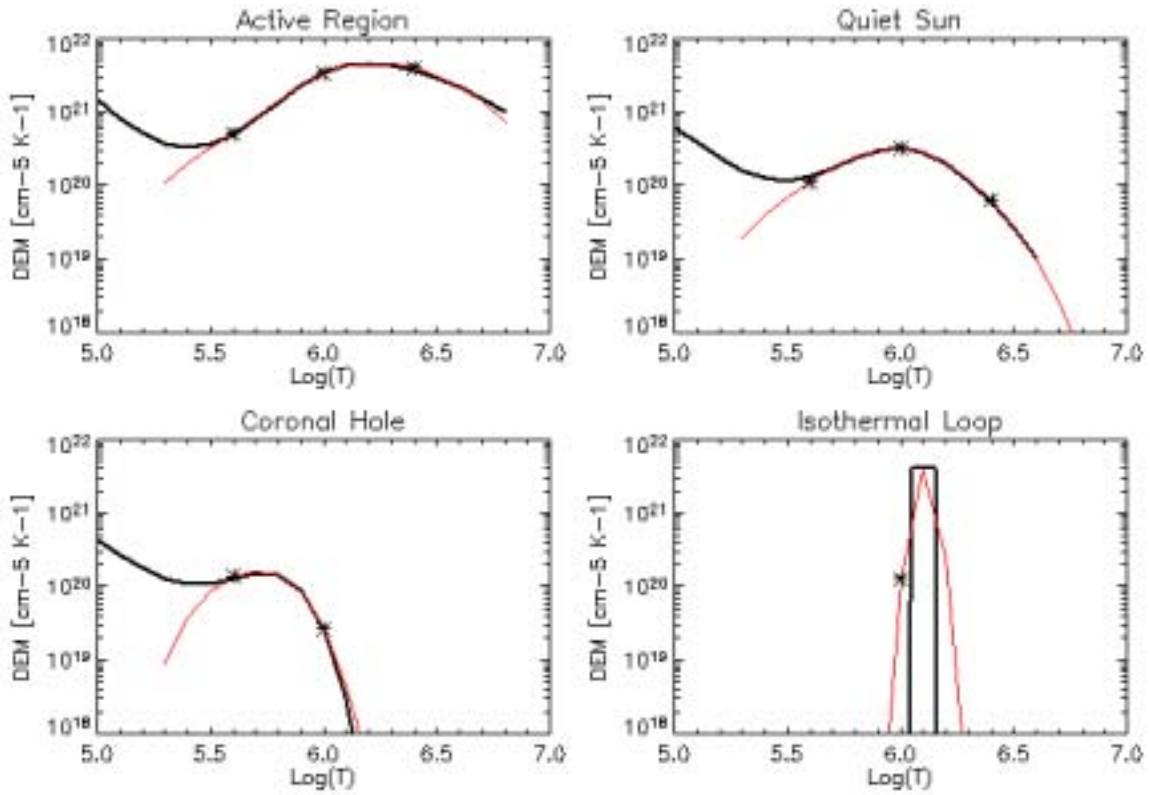

Figure 65. The red curves show the best-fit DEMs found by forward fitting to noise-free simulated observations. The input DEMs used to generate the simulated observations are shown by the heavy black lines. The stars mark the location of the spline points. Three-point splines provided adequate flexibility to match a wide range of possible input DEMs, including very sharp isothermal DEMs like the one shown in the lower right panel. Plots such as these were used to fine-tune the temperature range, temperature grid, and boundary conditions of the DEM splines.

There remains the danger that the location of the spline points will bias our results. This concern is somewhat alleviated by the chosen grid's success in recovering many different DEMs, including a number of isothermal DEMs at various temperatures not plotted in Figure 65. However, while simplicity demands that we chose a single temperature grid and apply it to all observations, it is important to double-check the resulting DEM by periodically attempting to fit a set of observations with DEMs defined on different temperature grids and making sure that they converge to the same function. I made some effort to introduce a fourth free parameter dictating the location of the spline points, but this did not produce reliable results.

Another potential source of bias is the initial guess at the DEM. The starting point for each set of observations was a flat DEM, whose value was determined by normalizing to the MSSTA 195 Å pixel value. Different initial guesses appeared to have little effect on the final result; the algorithm

spent a little more time hunting for a minimum if the starting point was chosen without reference to the data, but tended to converge to the same solution.

A variety of simulated datasets were fed into the forward fitting algorithm to test its response to noise and its dependence on the number of different bandpasses. A summary of the results is shown in Figure 66. For each panel, the CHIANTI active region DEM was used to generate 100 randomized datasets, as described in Section 4.1.1. The upper left panel shows the result of simulations using a realistic low-end estimate of 25% noise on the temperature kernels of the 5 MSSTA EUV telescopes that succeeded in flight: the 150 Å, 171 Å, 180Å, 195 Å and 211 Å instruments. The 1-σ accuracy of the recovered DEM is about ±50% over the interval from Log(T) = 6.0 – 6.3; the results are substantially worse at the edges of the temperature interval.

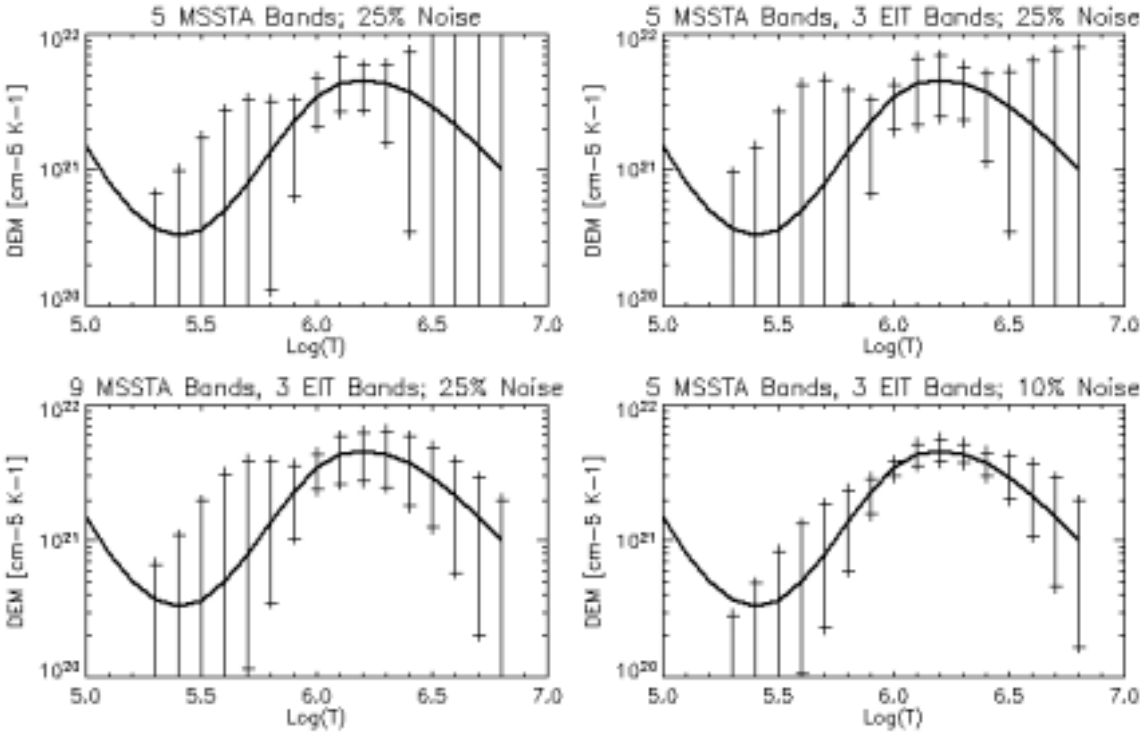

Figure 66. Reconstruction of DEM from simulated datasets. The simulated data sets consisted of 100 randomized observations of CHIANTI's active region DEM (dark black line). Clockwise from upper left: observations using the MSSTA 150 Å, 171 Å, 180Å, 195 Å and 211 Å bandpasses, with 25% noise; the previous bandpasses, plus EIT's 171 Å, 195 Å and 284 Å bandpasses, also with 25% noise; the same 8 bands, with only 10% noise; and those 8 bands, plus the MSSTA 58 Å, 98 Å, 131 Å and 256 Å bands, all with 25% noise. Error bars show 10%-90% solutions (>1σ).

Increasing the number of bandpasses does not substantially improve the accuracy in the center of the curve, but it does extend the range over which the reconstruction is valid. The upper right panel shows the effect of adding the 3 EIT EUV bandpasses to the analysis, still with a 25% error on all observations. (These conditions most closely resemble the expected state of the data obtained from the MSSTA flight and supporting observations). The 1-σ errors are now about ±35% over the 1-2 MK region (Log(T) = 6.0 – 6.3). In the bottom left panel, we examine how our capabilities could have improved had the other four MSSTA EUV telescopes (at 58 Å, 98 Å, 131 Å and 256 Å) recorded usable images. In this instance, the performance of the reconstruction improves notably at the high temperature end of the range; but, somewhat surprisingly, it does not look much better at temperatures below 10^6 K, and the accuracy near the peak of the DEM curve is not noticeably better.

The accuracy of the reconstructed DEM depends strongly on the amount of error in the observations. In the bottom right panel of Figure 66, we see the effect of assigning only 10% error to the temperature kernels obtained on the day of flight. These results are quite encouraging: the 1-σ errors are down to ±20% over the range Log(T) = 5.9 – 6.4. The persistent poor performance of the fit at temperatures below Log(T) = 5.8 can be attributed in part to the fact that the active-region DEM plotted here is dominated by a peak around Log(T) = 6.1; similar analyses using the quiet region and coronal hole DEMs show better agreement at lower temperatures. However, in no case is the algorithm successful at temperatures below about Log(T) = 5.7; the instruments are simply not tuned to transition region material.

Clearly, both the level of error in the observations and the exact nature of that error are critical to interpreting the results of a DEM recovery. Considerable effort went into pinning down the sources of error in the MSSTA observations; however, it is important to reiterate that this error is not necessarily well-modeled by the Monte Carlo process used in the above simulations. While the simulations introduce noise by adding uncorrelated random numbers to the true value to each observation, the actual error associated with each pixel value consists of several components:

- There is some statistical noise (photon counting or film grain statistics) in every pixel. This component, which does behave in the same way as the noise in the simulations, is probably only at the level of a few percent;

- Each instrument has some intrinsic systematic error (for example, the difference between the true reflectivity of its multilayer mirrors and the efficiency curves used to generate its temperature kernel). Generally at the level of ~10%, this error will not be centered on 0 in a set of 100 Monte Carlo realizations, and thus will distort the shape of the recovered DEM (or render DEM recovery impossible).

- Uncertainties and inaccuracies associated with the calculation of the EUV emission line contribution functions were not measured, but they certainly exist. Abundance, ionization equilibria and atomic physics of the emitting ions are probably uncertain at about the 10% level. This adds to the systematic error of each bandpass.

- Some systematic errors are common to one or more instruments. For example, most of the bandpasses used in the analysis are dominated by emission from ionized iron, so errors in the assumed iron abundance affect all instruments similarly. The uncertain calibration of the MSSTA photographic film also has similar (though not identical) effects on all the MSSTA images.

Sadly, the net result of all this uncertainty, and uncertainty about the uncertainty, is that there is no obvious answer to the question of how reliable our results might be. That the largest component of the error in the MSSTA observations – the film calibration – is shared by all the MSSTA bandpasses suggests that the relative error on the MSSTA observations may be only about 10%, while the absolute error is up around 25%-35%. In that case, we might hope to recover the *shape* of the DEM curve quite accurately, though the overall normalization of the DEM might be off. However, a more pessimistic reading of the above caveats would emphasize the fact that all the bandpasses contain substantial uncorrelated systematic errors, and that therefore even the [somewhat underwhelming] results shown in the "realistic" simulation in the upper right panel of Figure 66 are not achievable with the data in hand.

The data themselves will provide some insight into these questions. Keep in mind that one of the central advantages of using a DEM-recovery approach to analyzing narrowband EUV images,

instead of filter ratios or direct physical modeling, is the possibility that no DEM will be found that reproduces the observations to within their error bars. Obviously, the simulations assume that a true DEM exists; nevertheless, they can provide a glimpse of the diagnostic power of unsuccessful DEM recovery. Figure 67 shows one instance of the DEM reconstruction using 12-band simulated data with 25% noise (similar to the ones plotted in the lower left panel of Figure 66). This instance was selected for its rather poor fit. The left panel shows a recovered DEM that bears little resemblance to the true DEM used to generate the observations. At right, we see that the recovered DEM reproduced the noisy observations to within their 25% error bars in 8 of the 12 bandpasses. To the eye, the agreement seen here is fairly typical of simulations with 25% noise; however, the χ^2 of this iteration was about 50% higher than the median for the run of 100 iterations. The high error statistic could serve as a clue that the recovered DEM was not as reliable as most of the others found in the run. Still, it is a fairly subtle clue, and would be unlikely to set off any alarms if these results were obtained on flight data. The fact that the observations were reproduced so well by a DEM so different from the true target simply reinforces the point that line fluxes are not the sensitive indicators of small fluctuations in temperature structure that we might like them to be.

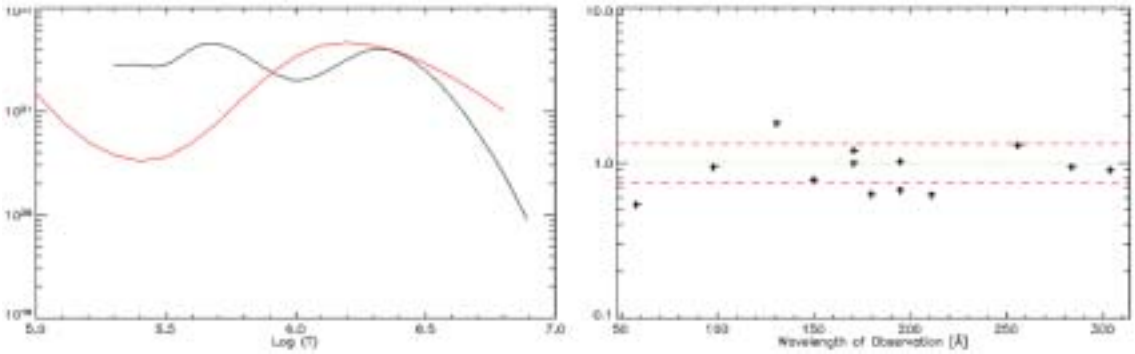

Figure 67. A rather poor DEM recovery in the presence of 25% noise. On the left is the recovered DEM in black with the input active region DEM in red. On the right is the ratio of noisy observations to observations predicted by the recovered DEM. Note that the observations agree fairly well with the predictions.

Thus the final word on the simulations is a cautionary one. Realistic error estimates limit the accuracy of any DEM recovery attempted with MSSTA and EIT data to about ±35% over the range Log(T) = 5.9 – 6.3. While there is some hope that correlations in the systematic errors might allow us to do better, at least as far as the shape of the DEM is concerned, it is also possible that systematic errors not accounted for may make our answers even worse. Furthermore, the χ^2

statistic will only alert us to fairly serious failings of our recovered DEM. This picture of what we might hope to accomplish is notably bleaker than the one painted by (Golub, DeLuca et al. 2004) using simulated data from the AIA; however, their investigation assumed primarily random noise at or below 10%, and did not consider uncertainties in the atomic physics parameters.

Even so, as long as our eyes are open to its shortcomings, the forward-fitting DEM recovery technique is well worth applying to real data. Indeed, it is an essential step if we are ever to extract meaningful quantitative information about the thermal structure of the corona from the beautiful and plentiful EUV spectroheliograms that make up such an important segment of the data currently available to solar physicists. As noted in Section 1.3.4. the limitations that apply to DEM extraction apply to any analysis of narrowband images; the DEM extraction technique just serves to highlight those limitations. With that in mind, we will now turn that technique on the images from the MSSTA III.

4.2 Extracting DEMs from the MSSTA data

The calibration of the MSSTA data is described at length in Chapters 2 and 3. As noted there, calibrated images consist of pixel values giving flux density (ergs cm^{-2} s^{-1}) at the film plane averaged over the area of the pixel. Thus, averaging pixel values over any region in the image and feeding those pixel values to the reconstruction algorithm will produce a line-of-sight DEM averaged over a region of the corona. For the purpose of the analysis, I also included the EIT 171 Å, 195 Å and 284 Å images; the MSSTA images were co-aligned and cross-calibrated with EIT as described in Section 3.3.

The images were then passed to an interactive sub-region extraction program (PB_INTERACT_DEM; the IDL code is in the Appendix) based on Dominic Zarro's image mapping software. There are some concerns that this software does not perfectly conserve energy when used to regrid images on different pixel scales (De Forest 2004), but perfunctory testing with MSSTA and EIT images suggested that any such distortions were below the 5% level. The program prompts the user to crop a region from the MSSTA 195 A image and select the MSSTA and EIT images to be included in the analysis; it then identifies the pixels in each selected image that correspond to the selected sub-region, averages them, and passes the resulting list of pixel values to one of the DEM-extraction routines tested in Section 4.1.3.

4.2.1 FULL-DISK

I began by examining the full-disk average pixel values; that is, the "sub-region" was a square 2000 arc-seconds on a side, centered on the solar disc (including the full solar disc, out to at least 10 Megameters above the limb on all sides). Gaps in the MSSTA (due to vignetting or misalignment of the telescopes) and EIT (due to dropped blocks within the images) data were ignored. The resulting average pixel values allow comparison with calibrated full-disk spectra such as those observed by [(Malinovsky and Heroux 1973), (Woods, Rodgers et al. 1999)], and are a valid test of the principle of analyzing flight data, although any results obtained from these

full-disk averages should probably not be considered for any in-depth analysis due to the data gaps and the non-rigorous square shape of the sub region.

Central λ [Å]	Exposure Time [s]	Full-Disk Average Pixel Value
MSSTA		
150	15	9.95
	20	10.33
	40	11.57
	100	11.26
171	15	74.51
	40	82.59
	100	66.03
180	3	61.76
	30	60.12
	40	61.99
	100	45.10
EIT		
171	8	133.54
195	13	114.24
284	123	7.98

Central λ [Å]	Exposure Time [s]	Full-Disk Average Pixel Value
MSSTA		
195	20	46.39
	30	47.31
	40	32.09
	100	29.29
211	65	30.26
	15	20.27
	30	24.07
	40	18.83
	100	38.93

MSSTA Units: 10^{-2} ergs cm^{-2} s^{-1}

EIT Units: DN s^{-1}

Table 18. Full-disk observations used to generate disk-averaged DEMs. All corrections described in Section 3.3 have been applied to these data.

The results, shown in Table 18, were then used to generate DEMs with the three-point temperature spline method. The output was extremely sensitive to the width of the temperature region and the placement of the pivot points – much more so than all but the noisiest simulations (done with 50% random error). This sensitivity immediately suggests that the data contain errors larger than the nominal 25%. Nevertheless, it was possible to find reasonable-looking DEM solutions that provided good agreement with the observations. The best results came from extending the temperature range from Log(T) = 4.5 – 7.5 (this is more of a way of tricking the spline method into lowering its endpoints than a statement about the physical insight provided by the instruments). The outcome is shown over the standard quiet sun DEM of (Vernazza and Reeves 1978), which is often used as a proxy for a full-disk average DEM, in Figure 68.

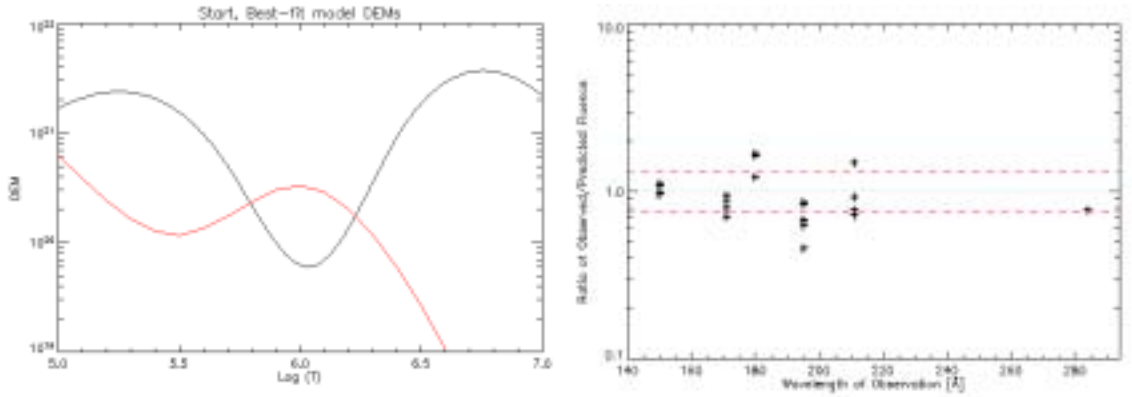

Figure 68. The best-fit three-point spline to the full-disk average pixel values shown in Table 18. At left, the recovered DEM (black) is plotted over the standard quiet-sun DEM (red). At right, the ratio of observed pixel values to those predicted by the recovered DEM are plotted; the red lines indicate ±25% error. The majority of the observations are well-reproduced by the recovered DEM.

Most striking about the recovered DEM is the dip at Log(T) = 6.0 – 6.1, precisely the temperature range at which the MSSTA and EIT instruments are most sensitive (and the temperature at which the standard DEM has its peak). This is a qualitatively persistent feature regardless of the placement of the spline points. Indeed, relaxing the conditions on the DEM reconstruction tends to make the dip even more pronounced. Figure 69 shows the effect of removing some of the constraints on the spline; the depletion of material at 1 MK becomes even more extreme, and the fit to the data becomes (slightly) better. Contrast these results with the ones shown in Figure 70, which uses the fluxes predicted by a flat DEM and the quiet sun DEM (with an overall normalization to optimize the fit). It is immediately apparent that the dip at 10^6 K improves the agreement between the model and the observations; in particular, it makes it possible to match the instruments at the cool (MSSTA 150 Å) and hot (EIT 284 Å) extremes of the overall kernel matrix. The slope of the observations at 171 Å, 180 Å, 195 Å and 211 Å (with increasing temperature of peak sensitivity) in Figure 70 indicates that the million-degree peak in the quiet-sun DEM is probably too hot; the cooler lobe of the recovered DEMs in peak is Figure 69 accounts for most of the emission seen in these bandpasses, and flattens out that slope.

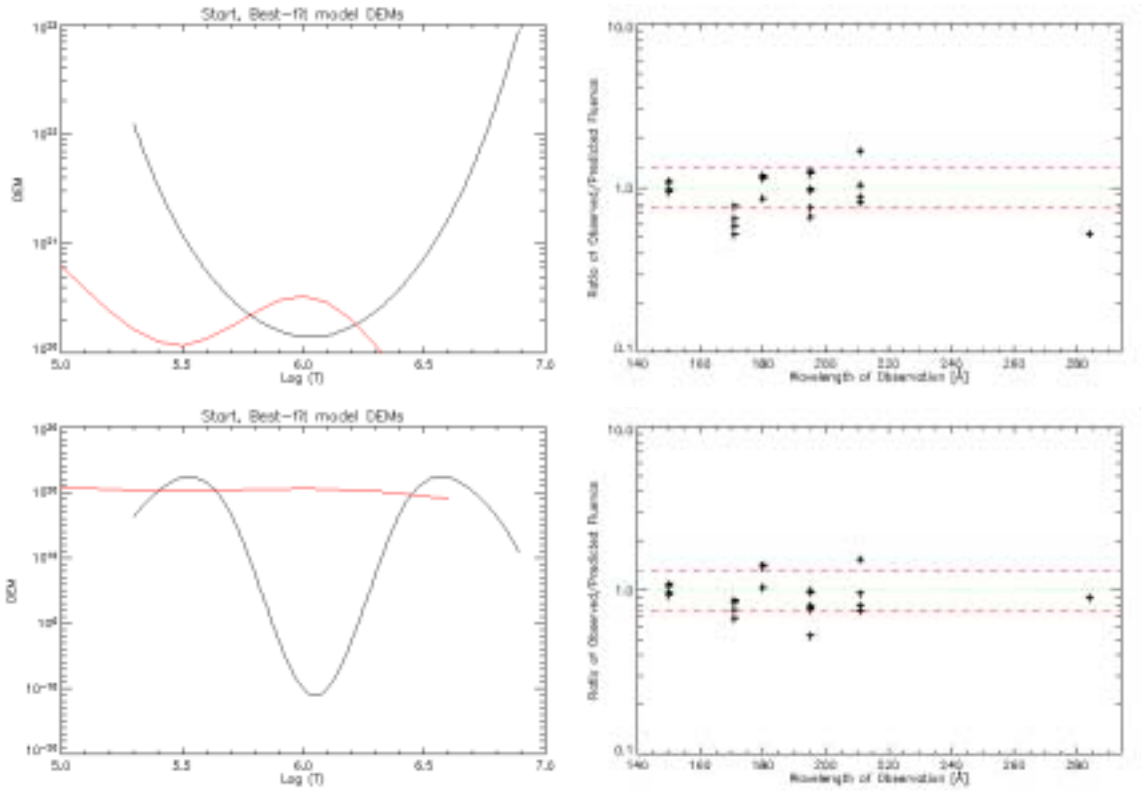

Figure 69. Relaxing the conditions on the spline only enhances the dip at 10^6 K. At top, the range was truncated at Log(T) = 5.3 – 6.9, and the endpoints allowed to float; at bottom, the endpoints were forced to go low, and the penalty function for excessively oscillatory DEMs was removed. Both result in slightly better fits to the observations, though qualitatively the DEMs and the reproduced observations are fairly similar to those shown in Figure 68.

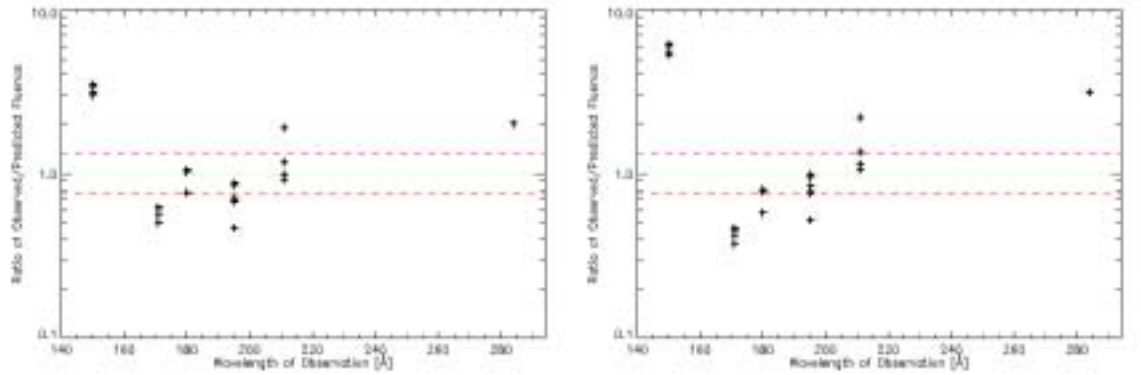

Figure 70. On the left, ratios of full-disk observations to those predicted by an optimized flat DEM of 3.1×10^{20} cm^{-5} K^{-1}; on the right ratios based on a scaled version of the standard quiet sun DEM. Both show noticeably worse agreement than that obtained from the distributions with little or no material at 10^6 K; the flat distribution is slightly preferred, although neither is entirely excluded by the data, given the possibility of systematic errors in the observations and temperature kernels.

The two-part plots in Figure 68 and Figure 69 are a rather unconventional view of observations fit by splined DEMs. Generally, a set of observations is plotted over the extracted DEM as error bars; for each bandpass (or emission line), the temperature at which the product of the temperature kernel (or contribution function) and DEM is a maximum is used as the x-coordinate of the point, and the ratio of observed to predicted pixel value times the predicted DEM is used as the y-coordinate (see, e.g., (Brosius, Davila et al. 1996)). The best-fit full-disk DEM is plotted in this fashion in Figure 71. This technique leads to confusion in terms of interpreting both the recovered DEM and interpreting multilayer observations in general, as it implies that each bandpass provides a measurement of the emission measure of plasma at a particular temperature, so that the DEM can be found simply by drawing a line through those points. In fact, both the x and y coordinates of the points shifts as the DEM is varied. However, this plot does verify that the best-fit DEM implies that the MSSTA is seeing very little million-degree material.

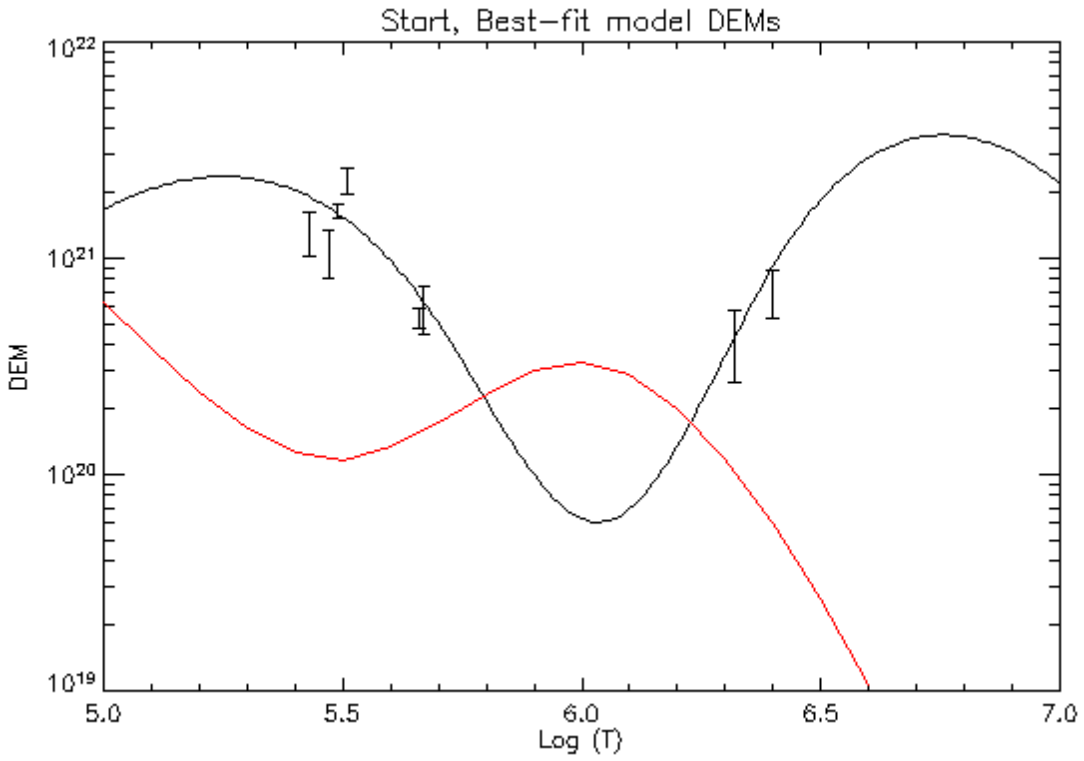

Figure 71. The best-fit full-disk DEM shown in Figure 68 is replotted. For each MSSTA and EIT bandpass, error bars are placed at the temperature at which the contribution in that bandpass is maximum. The size of the error bars is based on the scatter of multiple observations in each bandpass, and is not an estimate of systematic error. Plotting error bars in this way is somewhat misleading.

4.2.2 AREA-RESOLVED

Looking at averages over sub-regions of the solar disk has a number of advantages over treating the full-disk flux. We can avoid data gaps, ensuring that the measured average pixel values reflect the same lines of sight in each image. Furthermore, the extracted DEMs are easier to interpret; attempting to construct a model that reproduces an active-region DEM is more likely to put useful constraints on interesting physical parameters like heating rate and energy conductivity than trying to model a full-disk DEM. Ultimately, of course, we would like to extract DEMs for sub-regions of one arc-second or smaller, but the first step is to examine selected areas of a particular characteristic type of solar "terrain," and see what their recovered DEMs look like.

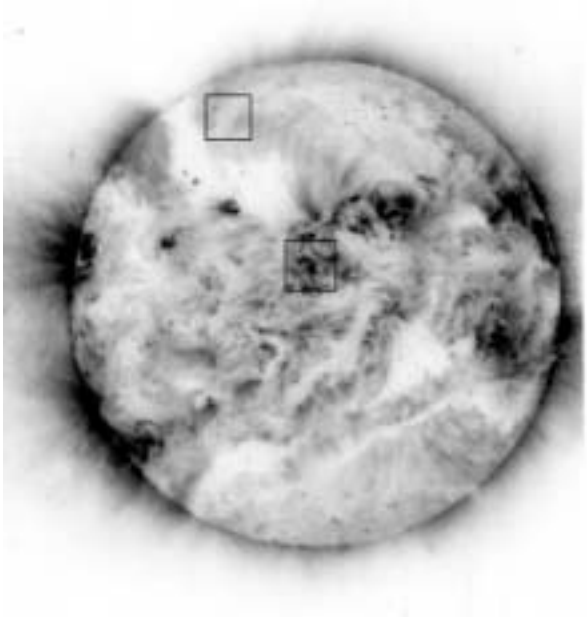

Figure 72. The active region and quiet region analyzed below were selected because of their placement on the disk, the fact that they were found in images in all 5 MSSTA and 3 EIT bandpasses, and because each is a fairly typical representative of a variety of solar feature. Here the selected sub frames are shown in boxes superimposed on an image from the MSSTA 195 Å telescope.

In Figure 72, the selected sub-regions are shown. These selections were made with reference to the MSSTA 195 Å telescope's 20-second exposure; pixels from the corresponding regions in the other MSSTA and EIT images were then identified and averaged by the DEM recovery routine. The results of the sub-region extraction are shown in Figure 73, with heliographic coordinates superimposed. The accuracy of the MSSTA-to-EIT co-alignment is somewhat problematic for the quiet region; however, any offset is unlikely to be a leading source of error.

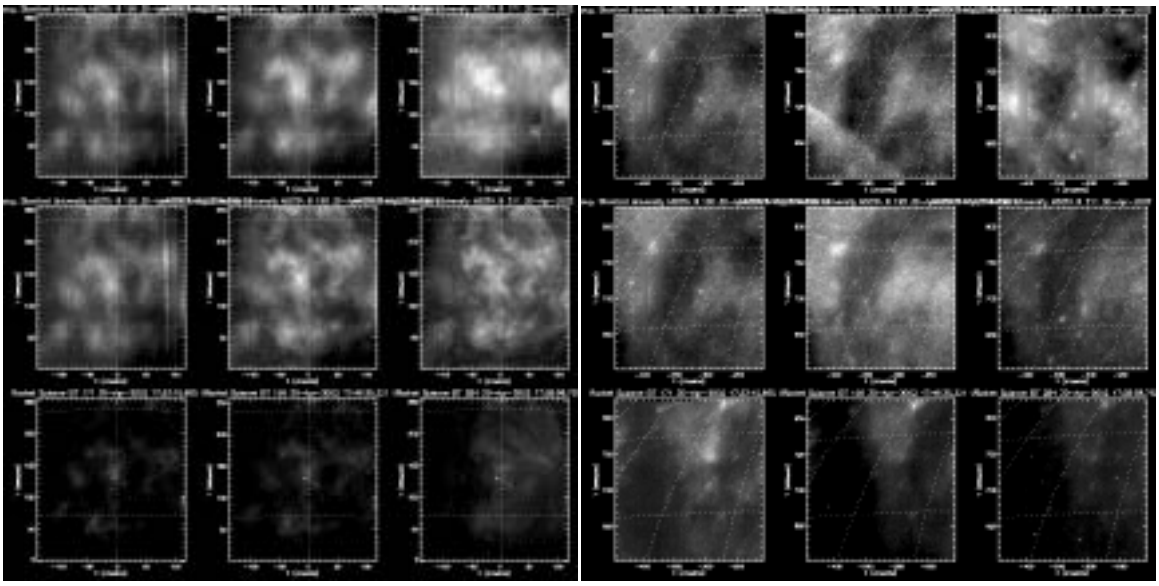

Figure 73. Sub-regions analyzed from the MSSTA III dataset. The left half of the frame shows a crop from an active region; the right half shows a region of diffuse quiet sun emission. Each sub-region is shown in 9 images. From left to right: (*top row*) MSSTA 180 Å, 150 Å, 171 Å; (*middle row*) MSSTA 180 A, 195 A and 211 A; (*bottom row*) EIT 171 Å, 195 A and 284 A. The 284 A quiet-region subframe contains a partial data gap, and was not used in the analysis.

Again, DEM recovery displayed great sensitivity to the definition of the temperature grid over which the spline was performed, casting the uniqueness of the solution into doubt. Figure 74 shows the results. Both recovered DEMs show depressed emission measure at 10^6 K, with most of the flux in all the bandpasses other than the 284 Å coming from cool material (T = 250,000 – 500,000K). However, despite the presence of the error bars plotted exclusively on the cool temperature lobe, the high-temperature lobe of the recovered DEM does contribute to the emission seen in the images and cannot be changed without affecting the fit. The agreement with observation is not especially good; the quiet region in particular is of by a factor of 2 on the 171 Å flux, and reproduces only two of the 5 wavelengths to within their nominal error bars. Still, the agreement is not unreasonable given the potential systematic errors associated with the images. The recovered DEMs provide substantially better agreement than either scaled versions of the relevant CHIANTI standard DEMs or optimized flat DEMs.

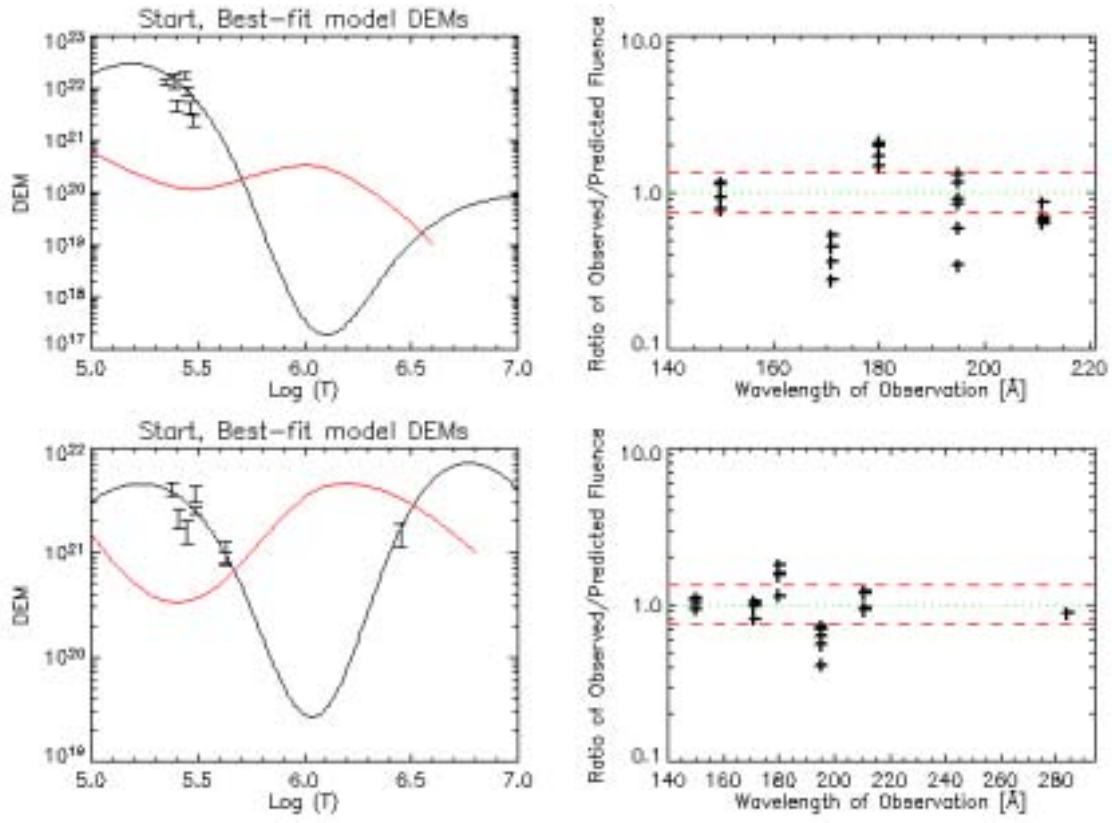

Figure 74. Results of sub-region DEM extraction. On top are the results from the quiet-sun region; the lower pair of frames are from the active region.

Carrying the sub-region DEM extraction to smaller spatial scales is certainly possible, but perhaps not as illuminating as might be hoped. Due to the imprecise co-alignment and marginal spatial resolution and signal-to-noise of some of the images, the smallest sub-region that can reliably be measured in multiple frames is about 30 arc-seconds across (so that the number of resolution elements on the full solar disk is of order 1000). Because of the alignment problems of the MSSTA 150 Å and 180 Å images, and data gaps in the EIT images, as many as half of these pixels will contain subsets of the full dataset. 30 arc-second-wide columns of sight through the corona are guaranteed to contain many unrelated structures, and thus the DEMs extracted from observations on this scale must be treated as ensemble averages or descriptions of global conditions in much the same way that the larger sub-regions discussed above were. Only by measuring multiple independent bandpasses at spatial scales of 1 arc-second or less will we have any hope of constraining plasma conditions in a single coronal loop (and, as discussed in Section 1.3, even if we could achieve a spatial resolution of 0.1 arc-second or smaller the problem of line-of-sight confusion in an optically thin plasma cannot be entirely avoided).

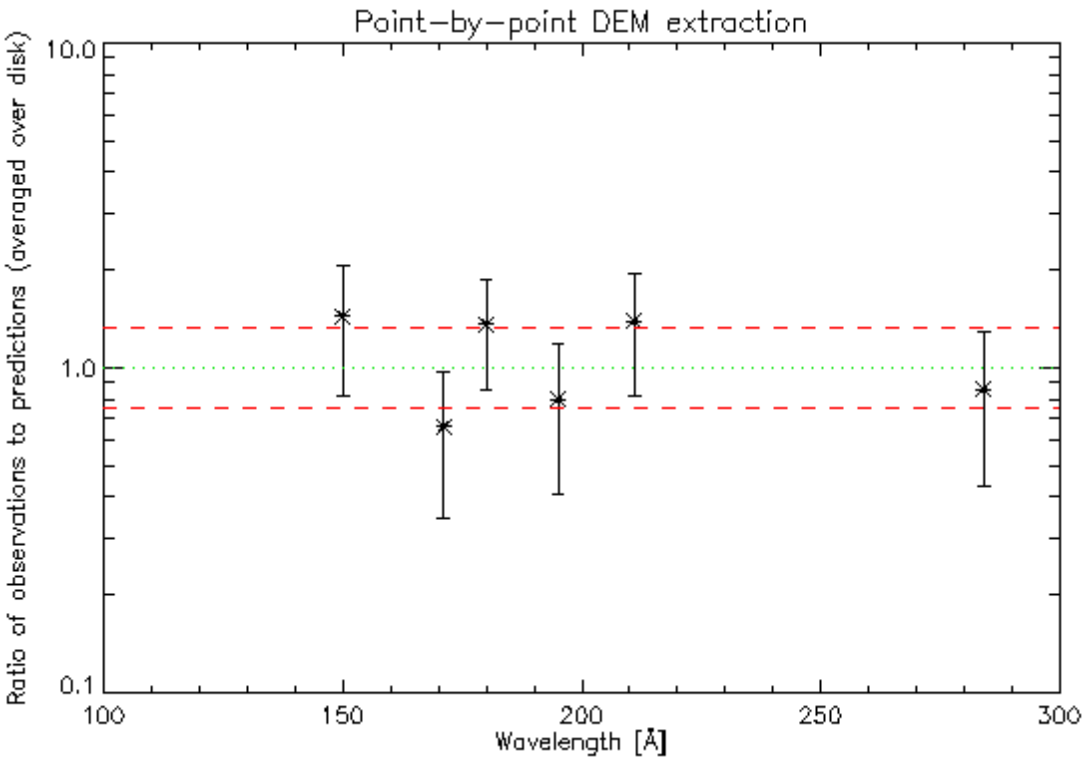

Figure 75. Pixel values predicted by DEM extraction on 30 arc-second square sub-regions compared with observations from those sub-regions (mean and RMS error over the full disk).

Nevertheless, a preliminary point-by-point DEM extraction was performed by mapping all the images to a regular 30-arcsecond grid and feeding the fluences for those points in the grid with full wavelength coverage to the DEM-fitting algorithm. The behavior of the fit under different splining procedures was not studied in detail; the run time of the algorithm for a single set of observations is of order 10 seconds, so extracting DEMs for the full disk takes several hours of computation. The overall agreement between observations and the predictions of the best-fit DEM is somewhat worse for the small pixels than for the more global regions, as might be expected. Figure 75 demonstrates this; for each bandpass, the RMS of the ratio of observations to predictions for all pixels and all exposures in that bandpass is plotted for the 8 EUV instruments included. The results are noticeably worse than what was found for the full-disk average DEM, and outside the nominal calibration error of most of the instruments. It is likely that these results could be improved by a more thorough optimization of the DEM modeling technique. In general, however, the pixel-sized DEMs display the same suppression of 10^6 degree material seen in the averaged regions.

4.2.3 RELIABILITY OF RESULTS

The first explanation for this phenomenon must be based on consideration of instrumental effects, particularly those effects that are not well-described by the calibration measurements listed in Chapter 2. For example, it is obvious from a cursory examination of the images that they contain significant scattered light (see, for example, the broad arcs of reflection from the Ritchey-Chrétien baffle tube overlying the disk in Figure 54), which our piecemeal approach to calibration does not account for. A white-light contribution to the flux observed in each pixel and each wave band would tend to suppress the expected peak in the DEM around one million degrees.

However, the general bimodal shape of our DEM results suggests that a pixel-by-pixel two-temperature model, like the one constructed by (Zhang, White et al. 1999) using just the three EIT bandpasses, could be made to fit the observations over the full disk. Constructing such a model based on the 8 bandpasses we observed would, in principle, be more reliable, less dependent on assumptions and initial conditions, and easier to interpret. However, before we conduct any analysis using the extracted DEMs, it is important to understand their limitations.

Re-examining Figure 69, we note that there is little to help us decide between the two recovered DEMs on the basis of their agreement with observations (that is, the right panels of the top and bottom rows, and the χ^2 of the fits, are difficult to distinguish). However, the DEMs bear only the most superficial qualitative resemblance to each other in that they are two-peaked, and have little emission measure at Log(T) = 6.0.

The problem is that the cubic spline tends to find solutions by driving the emission measure out to the edges of the region of sensitivity of the instrument temperature kernels; the χ^2 changes more smoothly in this regime, making it easier to find a solution. However, if good solutions without the dip at 10^6 K were to be found, they would have been uncovered; the fits were performed over a number of different combinations of end points and spline points, and always wound up with a two-temperature solution. Inspecting the error surface in parameter space confirms that more conventional-looking solutions like the CHIANTI standard DEMs consistently produce higher χ^2

values than the best-fit results. However, a wide variety of possible DEMs – including the two-temperature recovered solutions, the CHIANTI DEMs, and even flat multithermal plateaus such as those found by (Schmelz, Saba et al. 1999) – predict pixel values that are within a factor of 2 of all the observations in the MSSTA dataset, and thus none of these possibilities can be ruled out under the most pessimistic interpretation of the errors involved in the calibration of the instrument and the assumptions underlying the analysis.

4.3 Solar Physics Implications

Thus, my efforts at DEM recovery lead to two general conclusions: first, that the MSSTA observations are best fit by a model of the solar atmosphere that includes very little material at 10^6 K; and second, that neither the DEMs I found by analyzing the MSSTA observations, nor other DEMs obtained from similar analysis of similar observations, constitute compelling descriptions of the state of the solar plasma.

4.3.1 THE 1 MK DIP

If we believe the suggestion that there is very little million-degree material in the solar atmosphere, and that most of what is seen in the MSSTA images is much cooler plasma, we are forced to question a large body of observational evidence and an even larger body of received wisdom to the contrary. That images in the 171 Å bandpass depict the distribution of million-degree plasma is virtually gospel; the vast majority of studies of EUV spectroheliograms conclude that the plasma being observed is around $1 - 2 \times 10^6$ K (see, *e.g.*, (Lenz, DeLuca et al. 1999), (Aschwanden, Newmark et al. 1999) and numerous other works cited in Section 1.3). However, it has already been argued that any temperature measurements based on EUV filter ratios or forward modeling of plasma conditions are non-unique and thus highly suspect. More difficult to dismiss are the independent measures of coronal temperature obtained from microwave observations(White 1999), or spectral observations, including temperature sensitive emission line ratios (Feldman, Doschek et al. 1999); still, there are ambiguities in these observations, including discrepancies of a factor of two or more between microwave brightness temperatures and temperatures measured from EUV observations (Landi and Chiuderi Drago 2003).

It is interesting to note how much work is conducted by starting from the assumption that all plasma seen in a narrowband image is at the temperature at which the sensitivity of that imager peaks (in addition to the sources cited previously, see, *e.g.,* (Nightingale, Aschwanden et al.

1999), (Reale, Peres et al. 2000)). One of the central theoretical pillars of this sort of analysis of EUV images is the Unresolved Fine Structure model proposed by (Feldman 1983; Feldman, Widing et al. 1999) based on spectral observations. The UFS model states that the distribution of plasma at any temperature T_e in the solar atmosphere is dominated by nearly-isothermal structures that hold their peak temperature T_e along most of their lengths. Thus, while the heating source of structures at different temperatures may be the same, those structures are essentially unrelated to each other.

The UFS model discounts the importance of conduction in heating the lower corona and transition region, and holds that, in general, hot structures do not have visible footpoints at cooler temperatures. Evidence for this picture is provided by the observations of (Vourlidas, Klimchuk et al. 2001), who found that Hydrogen Lyman α emission is much too strong for back-heating from the corona (which they assumed to be isothermal, with temperature and density based on TRACE filter ratios) to account for more than a tiny fraction; and (Aschwanden, Schrijver et al. 2001), whose forward models of EUV-emitting loops included negligible conductive flux. In general, the assumption that narrowband multilayer images show isothermal, million-degree plasma leads to assumptions that favor isothermal models of coronal structures, which in turn imply little connectivity between the various temperature layers of the solar atmosphere. As noted in Section 1.3.4, this unified picture of the solar atmosphere also has implications for the location and nature of the coronal heating function.

However, it is possible to lose track of which results of the UFS model are based on observations and which are based on assumptions. There is some evidence of the importance of conduction from hot material to cool material: for example, (Berger, De Pontieu et al. 1999) found that the "moss" observed in TRACE 171 Å images was in fact composed of the roots of very hot structures seen as large coronal loops in Yohkoh SXT. Using forward modeling on MSSTA II data, (Martínez-Galarce, Walker et al. 2003) discovered a class of coronal funnels that could account for 171 Å and 1550 Å observations; these funnel models included solutions with peak temperatures ~500,000 K and strong emission in the 171 Å bandpass, in qualitative agreement with the two-temperature best-fit solutions presented in the previous section. As discussed in (Oluseyi, Walker et al. 1999)'s overview of solutions to the general quasistatic loop equation of (Vesecky, Antiochos et al. 1979), coronal loops with relatively low peak temperature can account for much of the observed diffuse EUV emission when ambipolar diffusion is taken into account;

these models imply substantially higher downward conduction than more conventional loop models.

Therefore, the fact that the MSSTA and EIT data can be reproduced by DEMs that consist of very little million-degree material, suggesting that most of the EUV flux is coming from cooler plasma, with contributions from a high-temperature component, is intriguing. In order to make this conclusion more quantitative, we would like to better-constrain the cool temperature component. As noted previously, it should be possible to construct a two-component DEM model for each pixel in the MSSTA/EIT dataset. Such a model would use four parameters (temperature and emission measure of the hot and cool components), similar to the models presented by (Zhang, White et al. 1999), but with the advantage of not being nearly as underconstrained. An effort to fit such a model to the points on a 30-arc-second grid is underway.

However, it is doubtful that such models will prove useful in the long run. They are not motivated by any compelling physical intuition: despite the fact that the corona is optically thin, and thus any line of sight will likely sample multiple different structures, there is no reason to believe that, in any large region, the solar atmosphere is well-described as a distribution of plasma at exactly two well-separated temperatures. Furthermore, the interpretation of the results is difficult (see the full-disk maps of the four model parameters in (Zhang, White et al. 1999), which cannot be said to provide much more insight into the state of the corona than the raw images). They are interesting primarily as an empirical result; that is, they are interesting because they can reproduce the observations.

4.3.2 DEM INDETERMINACY

The most important conclusion we are forced to draw, however, is that there are a lot of possible DEMs that can reproduce the observations. It would be absurd to state that the MSSTA and EIT data prove that there is no million-degree plasma in the solar atmosphere, and presumptuous to state even that they strongly suggest the absence of million-degree material. What we can say is that it is possible to account for MSSTA and EIT observations with a DEM that includes little or no million-degree material. Indeed, given the potential errors associated with the MSSTA and

EIT, it is impossible to rule out a two-temperature DEM, or a flat plateau, or even a modest peak at around 10^6 K. While the suite of narrowband images is not entirely useless as a constraint on the DEM – the overall level is fairly well-specified, and vast regions of DEM-space can be excluded for each pixel – it is not nearly as useful as might be hoped, or as is frequently assumed.

What can be done to improve the reliability of the results? First, it is possible in principle to extend the dataset by adding more multilayer bandpasses or including spectral observations. Ideally, we would like more bandpasses to constrain the behavior of the DEM at lower and higher temperatures, in order to exclude spline solutions with unrealistically high values at these temperatures. For example, the recovered quiet-sun DEM shown in Figure 74 includes a very high emission measure at T < 500,000K; the MSSTA observations are not sensitive enough to half-million degree material to reject this solution, but the large value of the DEM at Log(T) = 5.4 is inconsistent with the results of (Dupree 1972) and other spectral measurements of lines formed at transition region temperatures. From a practical standpoint, the set of data that could be used to extend the analysis of the MSSTA images is limited. TRACE images were taken simultaneously with the MSSTA flight, and these could be added to the fitting algorithm; but TRACE includes essentially the same 171 Å, 195 Å and 284 Å bandpasses as EIT, so its usefulness would be more as a cross-check than as an extension. Spectra from SUMER and CDS are available, but they cover such a small spatial region of the sun as to be largely inapplicable to the full-disk MSSTA and EIT images, and are offset in time by several hours. Ground-based observations in H α or Ca K are taken every day, but those lines are not optically thin and thus not susceptible to temperature kernel analysis. The same is true of the MSSTA 1550 Å and 1216 Å telescopes. Full-disk spectra from TIMED-SEE might be valuable in setting boundary conditions on the spline, but lack the spatial and spectral resolution to be included in the extraction routine.

However, as demonstrated in Figure 66, broader temperature coverage is less important than more accurate observations. We could certainly impose more constraints on the lower temperature ends of the DEM, or exclude some of solutions returned by the algorithm, on the basis of physical insight; that is an essential part of forward modeling. But the fact that such a wide range of diverse DEM distributions are capable of fitting the data means that it is there is no clear point at which we can draw the line between providing intelligent feedback to the fit and simply specifying the DEM *a priori*. As long as our measurements contain errors of 25% or more, there is little hope of identifying unique, well-constrained DEM distributions. We need to get the

error down to about 10% before we can begin to make definitive statements about the shape of the DEM. That 10% limit must include not only photon counting statistics, but systematic errors in the instrument calibration, imprecision and inaccuracies in the atomic physics of the emitting ions, and breakdowns or incorrect choices in the assumptions underlying a DEM analysis. The implications of this requirement for future missions are discussed in the final chapter.