








Incorporating Uncertainties in Atomic Data into the Analysis of Solar and Stellar Observations: A Case Study in Fe XIII

Xixi Yu¹, Giulio Del Zanna² , David C. Stenning¹ , Jessi Cisewski-Kehe³ , Vinay L. Kashyap⁴ , Nathan Stein⁵,
David A. van Dyk¹, Harry P. Warren⁶ , and Mark A. Weber⁴

¹ Statistics Section, Department of Mathematics, Imperial College London, London SW7 2AZ, UK

² DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK

³ Department of Statistics and Data Science, Yale University, New Haven, CT 06511, USA

⁴ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

⁵ Spotify, 45 W 18th St, New York, NY 10011, USA

⁶ Space Science Division, Naval Research Laboratory, Washington, DC 20375, USA

Received 2018 March 19; revised 2018 August 10; accepted 2018 September 2; published 2018 October 23

Abstract

Information about the physical properties of astrophysical objects cannot be measured directly but is inferred by interpreting spectroscopic observations in the context of atomic physics calculations. Ratios of emission lines, for example, can be used to infer the electron density of the emitting plasma. Similarly, the relative intensities of emission lines formed over a wide range of temperatures yield information on the temperature structure. A critical component of this analysis is understanding how uncertainties in the underlying atomic physics propagate to the uncertainties in the inferred plasma parameters. At present, however, atomic physics databases do not include uncertainties on the atomic parameters and there is no established methodology for using them even if they did. In this paper we develop simple models for uncertainties in the collision strengths and decay rates for Fe XIII and apply them to the interpretation of density-sensitive lines observed with the EUV (extreme ultraviolet) Imaging spectrometer (EIS) on *Hinode*. We incorporate these uncertainties in a Bayesian framework. We consider both a *pragmatic Bayesian* method where the atomic physics information is unaffected by the observed data, and a *fully Bayesian* method where the data can be used to probe the physics. The former generally increases the uncertainty in the inferred density by about a factor of 5 compared with models that incorporate only statistical uncertainties. The latter reduces the uncertainties on the inferred densities, but identifies areas of possible systematic problems with either the atomic physics or the observed intensities.

Key words: methods: statistical – Sun: corona

Supporting material: figure set

1. Introduction

Spectral observations of solar and stellar coronae, mostly taken in the X-ray, extreme ultraviolet (EUV), and ultraviolet (UV) part of the spectrum, are regularly combined with atomic data to infer fundamental plasma parameters such as electron temperatures and densities. This information is essential for constraining models of coronal heating. In general in astrophysics, reliable atomic data are essential for interpreting and modeling X-ray observations, as, for example, discussed in Kallman & Palmeri (2007). One key aspect about (the) modeling is the accuracy of the atomic data, an area which has recently received some attention, see e.g., Luridiana & García-Rojas (2012), Bautista et al. (2013), Loch et al. (2013), Chung et al. (2016). In these studies, guidelines to estimate uncertainties as a routine part of the computations of data have been provided, or some preliminary analysis based on comparisons between different calculations.

Over the years, the accuracy of spectral observations and atomic calculations has progressed hand in hand. Current space instruments now provide measurements with an accuracy of 20% or better. With the earliest observations and atomic data, discrepancies between measured and predicted line intensities of factors of two or more were common. In the past few years, due to large-scale atomic structure and scattering calculations (see, e.g., the reviews Badnell et al. 2016; Jönsson et al. 2017), atomic data have improved significantly. In a series of papers,

starting from Del Zanna et al. (2004), one of us (G.D.Z.) has benchmarked the available atomic data for several of the main ions against all the available experimental data, from laboratory sources to astrophysical spectra, and generally found good agreement (within 10%–20%) between observations and theoretical calculations. Several reasons for the remaining discrepancies have been identified. Sometimes lines were blended, sometimes their identifications were incorrect. In some other cases the radiometric calibration of an instrument was at fault, or the atomic calculations were incorrect.

The problem of interpreting astrophysical spectra is a complex one and it is tempting to account for any systematic uncertainties in the analysis with broad, ad hoc assumptions. One might assume, for example, an inflated uncertainty for each observed line intensity with the hope of encompassing the uncertainty associated with counting statistics, the radiometric calibration, and the atomic data. Such broad assumptions, however, do not account for obvious correlations within the analysis. For example, the uncertainty in the calibration clearly should be much smaller for two emission lines that are close in wavelength than for two lines that are far apart in wavelength. Similarly, the atomic data for strong lines that are the result of transitions from the ground state are likely to be more accurate than the atomic data for weak lines that are influenced by many different transitions (e.g., Foster et al. 2010).

Fortunately, recent reductions in the cost of computing have made it possible to carry out detailed statistical analysis on

complex systems. Lee et al. (2011), for example, have developed a Bayesian approach for sampling from a distribution of plausible calibration curves and incorporating this information into a comprehensive analysis of high-energy *Chandra* spectra. Along the same theme, here we consider the effect of uncertainties in the atomic data and the problem of propagating them to the determination of plasma parameters. Since the atomic calculations involve many thousands of parameters, we cannot sample the posterior in the usual way. Thus, we develop a method that relies on a relatively small number of realizations of the atomic data to compute posterior distributions for the parameters of interest. While there still exist several computational bottlenecks in the process, we present for the first time a framework that can be applied in general to this class of problems. To begin with, we deploy a simple, but realistic, model to describe the uncertainties in some of the atomic parameters for Fe XIII, use it to generate different realizations of the plasma emissivities, then apply this ensemble of atomic data to analyze the density-sensitive Fe XIII lines observed in the solar active region spectra using a Bayesian framework.

We are focusing here on Fe XIII for various reasons: it is one of the most widely used ions; several atomic calculations are available; the main lines, from the $3s^2 3p 3d$ configuration, are well identified and strong in active regions; are mostly free of blends; and fall in a spectral region where they are observed by the EUV Imaging Spectrometer on *Hinode* (EIS, Culhane et al. 2007) and the radiometric calibration is relatively well understood. EIS is routinely used to measure electron densities from coronal iron ions (e.g., see Watanabe et al. 2009; Young et al. 2009). Furthermore, considering lines from a single ion removes the uncertainties related to the ionization fractions and elemental abundance, greatly simplifying the problem. Future work will expand this analysis to the calculation of the uncertainties associated with emission measure distributions, which generally use observations of emission lines from different ionization stages and elements.

This paper is structured in the following way. In Section 2 we describe the Fe XIII emission lines of interest and present the traditional analysis of some representative observations. In Section 3 we develop the models for the uncertainties in the atomic data and describe how they are used to generate different realizations of the plasma emissivities. In Section 4 we present the analysis of a simple test case where we specify the plasma properties. In Section 5 we describe the application to actual observations. In Section 6 we conclude with a summary and a discussion of future work.

2. Modeling EIS Fe XIII Observations

To motivate our thinking about the data analysis, we consider the problem of determining the electron density in loop footpoints using observations of density-sensitive Fe XIII emission lines. The uncertainties in the atomic data and the analysis presented in the following sections are independent of the specific problem to which they are applied. This analysis could, for example, also be applied to measuring densities in million degree loops or in the diffuse corona above the limb. We feel, however, that considering a specific application makes the analysis both realistic and tractable.

Figure 1 illustrates a typical observation of a solar active region. The intense magnetic fields in the active region lead to the formation of 3–4 MK plasma on the relatively short loops in

the active region core (e.g., Warren et al. 2012; Del Zanna 2013; Del Zanna & Mason 2014). The footpoints of these high-temperature loops are bright in million degree emission lines, and the footpoints are often referred to as the “moss” because of their mottled appearance in high resolution images (e.g., Berger et al. 1999; Fletcher & De Pontieu 1999). These footpoint measurements provide information on the boundary conditions in these loops and are important for constraining models of coronal heating (e.g., Peres et al. 1994; Martens et al. 2000; Winebarger et al. 2008). Measurements of the electron density in the moss are of particular utility because they yield information on both the base pressure of the loop as well as the filling factor (Warren et al. 2008).

To further illustrate the concept of the moss, in Figure 2 we show Fe XIII 202.044 Å intensities computed from a simple, one-dimensional hydrodynamic loop model (Schrijver & van Ballegooijen 2005). Here a relatively large volumetric heating is assumed and a loop with an apex temperature of about 3 MK is produced. As expected, the Fe XIII emission comes from a relatively localized region near the footpoint of the loop. Thus, moss observations can yield information on individual loops. The emission at higher temperatures, in contrast, is generally an integration across many different loops along the line of sight, making the interpretation of such observations much more difficult.

As mentioned previously, the EIS instrument on *Hinode* observes many emission lines whose intensities can be combined to form density-sensitive ratios. Figure 1 also shows the spectral region observed with EIS near 202 Å, which is dominated by intense Fe XIII lines. A typical analysis of these lines involves fitting each of the spectral features with Gaussians to derive the line intensities and their corresponding statistical errors. For this work we consider the lines at 196.525, 200.021, 201.121, 202.044, 203.165, 203.826, and 209.916 Å (see Table 1) and fit them with single or multiple Gaussians, as appropriate. These lines will be discussed extensively in Section 3. We have randomly selected 1000 pixels from the EIS observations of the moss shown in Figure 1 for analysis. Note that the lines at 202.044 and 209.916 Å originate in the same upper level and they form a branching ratio that is independent of solar conditions. The other five lines form density-sensitive ratios with the 202.044 Å line.

Since it is not obvious how to reconcile all of the individual line ratios, we use a simple empirical model to interpret the observations. We make the standard assumption that the intensity of an emission line can be computed with

$$I_n = \epsilon_n(n_e, T_e)n_e^2 ds, \quad (1)$$

where $\epsilon_n(n_e, T_e)$ is the plasma emissivity, n_e and T_e are the electron density and temperature, and ds is the effective path length through the solar atmosphere (see, for example, Mariska 1992). First, we note that since the excitation energies of the levels associated with these wavelengths are very similar, the emissivity ratios used to evaluate the plasma densities are highly insensitive to changes in temperature. Furthermore, as illustrated in Figure 2, most of the Fe XIII emission in high-temperature loops are thought to originate in a narrow region near the footpoint over which the temperatures are close to the peak temperature of formation of the ion, ≈ 1.8 MK. We therefore adopt this value of the temperature and henceforth treat the plasma as isothermal. Of course, hydrodynamic




Figure 1. Observations of the NOAA active region 11785 on 2013 July 8 near 2 UT. The top panels show the photospheric magnetic field measured with HMI, million degree emission observed with the AIA Fe IX 171 Å channel, and high-temperature loops observed with XRT. The middle panels show EIS rasters in a series of emission lines that range in temperature of formation from 0.8 MK to 5.0 MK. The bottom panel shows an EIS spectrum near 202 Å from a single spatial pixel with the Fe XIII lines of interest highlighted. The EIS rasters are from an observation that began at 01:55 UT, and this field of view is indicated by the dotted lines in the top panels. The EIS full-CCD spectrum is from an observation that began at 00:20 UT; this field of view is indicated by the dashed lines in the top panels.

models show that there are gradients in temperature and density along the loop, where segments of high n_e occupy small ds , and segments of small n_e cover a large ds , so Equation (1) must be treated as an empirical description characterized by a representative density and effective path length.

With this empirical description, however, we can derive information about the solar atmosphere directly from the observations. The physical model shown in Figure 2 depends on additional assumptions about the loop geometry, the plasma composition, and the nature of the heating.

For this empirical model we can use the observed intensities and their corresponding statistical uncertainties, the computed plasma emissivities, and assumed temperature to infer the best-fit electron density and path length by performing a least-squares fit. The plasma emissivities for each line is computed using version 8 of the CHIANTI atomic database (Dere et al. 1997; Del Zanna et al. 2015).

The results of an example calculation are shown in Table 2, where we have taken the observed intensities from a single spectral pixel (arbitrarily chosen as #217) from an EIS full-CCD observation (EIS file `eis_l0_20130708_002042`) and applied Equation (1). The resulting best-fit parameters are $\log n_e = 9.68 \pm 0.01 \text{ (cm}^{-3}\text{)}$ and $\log ds = 8.67 \pm 0.02 \text{ (cm)}$. The error bars associated with the parameters are very small, suggesting that the parameters are very precisely determined. The uncertainties associated with the intensities, however, are also small and the standard method of determining the best-fit n_e and ds by minimizing χ^2 results in reduced $\chi^2 \approx 30$ for this case, indicating that the model is a poor fit to the data.

This example highlights the difficulty in interpreting many solar observations. Since the Sun is relatively close, we can obtain observations with high signal-to-noise. This bounty of photons, however, means that models generally do not pass rigorous statistical tests. This can be simply ignored or covered up by inflating the statistical errors with ad hoc assumptions.




Figure 2. Density measurements in the moss can be used in conjunction with hydrodynamic loop models to infer the properties of high-temperature active region loops. The bottom panel shows the temperature and density as a function of position derived from assuming steady, uniform heating. The middle panel shows the intensity of Fe XIII 202.044 Å as a function of position along the loop. The top panel shows a rendering of the loop as it would appear viewed from the side. See Warren et al. (2008) for additional details.

Table 1
Selected Fe XIII Lines Observed by *Hinode*/EIS

| $i-j^a$ | Identification | λ_{exp} (Å) | Notes |
|---------|--|----------------------------|-------------------------|
| 1–20 | $3s^2 3p^2 \ ^3P_0-3s^2 3p 3d \ ^3P_1$ | 202.044 | ... |
| 2–23 | $3s^2 3p^2 \ ^3P_1-3s^2 3p 3d \ ^3D_1$ | 201.126 | ... |
| 3–20 | $3s^2 3p^2 \ ^3P_2-3s^2 3p 3d \ ^3P_1$ | (209.916) | branching ratio |
| 3–25 | $3s^2 3p^2 \ ^3P_2-3s^2 3p 3d \ ^3D_3$ | 203.826 | self-blend ^b |
| 3–24 | $3s^2 3p^2 \ ^3P_2-3s^2 3p 3d \ ^3D_2$ | 203.795 | self-blend ^b |
| 7–60 | $3s 3p^3 \ ^3D_1-3s 3p^2 3d \ ^3F_2$ | 203.772 | self-blend ^b |
| 8–60 | $3s 3p^3 \ ^3D_2-3s 3p^2 3d \ ^3F_2$ | 203.835 | self-blend ^b |
| 3–23 | $3s^2 3p^2 \ ^3P_2-3s^2 3p 3d \ ^3D_1$ | (204.942) | branching ratio |
| 1–23 | $3s^2 3p^2 \ ^3P_0-3s^2 3p 3d \ ^3D_1$ | (197.431) | branching ratio |
| 2–24 | $3s^2 3p^2 \ ^3P_1-3s^2 3p 3d \ ^3D_2$ | 200.021 | ... |
| 2–19 | $3s^2 3p^2 \ ^3P_1-3s^2 3p 3d \ ^3P_2$ | 209.619 | ... |
| 2–22 | $3s^2 3p^2 \ ^3P_1-3s^2 3p 3d \ ^3P_0$ | 203.165 | blended |
| 4–26 | $3s^2 3p^2 \ ^1D_2-3s^2 3p 3d \ ^1F_3$ | 196.525 | ... |
| 2–21 | $3s^2 3p^2 \ ^3P_1-3s^2 3p 3d \ ^1D_2$ | (204.262) | blended |

Notes.

^a i and j are the indices of the lower and upper levels in the CHIANTI database.

^b Self-blend: multiple lines from the same ion that are close in wavelength.

The real deficiency in the analysis is taking the atomic data as fixed and without uncertainty. In reality, the uncertainties associated with the plasma emissivities are likely to be comparable to or larger than those from counting statistics,

Table 2
Modeling Fe XIII Line Intensities in the Moss^a

| Line | $I_{\text{obs}} \sigma_I$ | I_{model} | $ \Delta I /I(\%)$ |
|---------|---------------------------|--------------------|--------------------|
| 196.525 | 1473.1 ± 18.8 | 1443.6 | 2.0 |
| 200.021 | 1521.4 ± 29.1 | 1749.9 | 15.0 |
| 201.121 | 2373.2 ± 44.4 | 1987.0 | 16.3 |
| 202.044 | 2866.5 ± 53.6 | 2989.1 | 4.3 |
| 203.165 | 775.2 ± 42.5 | 767.5 | 1.0 |
| 203.826 | 9237.6 ± 142.6 | 8751.2 | 5.3 |
| 209.916 | 530.2 ± 56.4 | 516.2 | 2.6 |

Note.

^a Observed and modeled line intensities for a single spatial pixel (#217) using Equation (1). The best-fit density and path length are $\log n_e = 9.68 \pm 0.01$ and $\log ds = 8.67 \pm 0.02$. The intensities and their corresponding uncertainties are in units of $\text{erg cm}^{-2} \text{s}^{-1} \text{sr}^{-1}$.

and a proper data analysis must include a treatment of them. We now turn to estimating the uncertainties in the atomic data available for Fe XIII.

3. Uncertainties in the Atomic Data

The most recent (and largest) scattering calculation for Fe XIII is an R -matrix calculation carried out within the UK APAP network,⁷ which had a target of 749 levels up to $n = 4$ (Del Zanna & Storey 2012). The main focus of this calculation was to provide accurate data for the $n = 4 \rightarrow n = 3$ soft X-ray transitions. Indeed, new lines in this wavelength range were subsequently identified (Del Zanna 2012a). The scattering calculation was supplemented by a structure calculation which was used to calculate the radiative data, using either observed or empirically adjusted theoretical wavelengths. The scattering and radiative data produced in this calculation were recently made available within the CHIANTI database⁸ in version 8 (Del Zanna et al. 2015). We use these data as our baseline.

Storey & Zeippen (2010) previously carried out a similar scattering calculation (using the same R -matrix method and the same codes), the only difference being that it was aimed at improving the earlier calculations for the $n = 3$ levels. The target had a total of 114 fine-structure levels, and included only some $n = 4$ levels.

Del Zanna & Storey (2012) also performed separate calculations for the $n = 5, 6$ levels, but showed that cascading effects are small when considering the strong EUV lines emitted by the $n = 3$ levels. The same paper also showed that the intensities of the transitions from the $n = 3$ levels are close to those of the previous Storey & Zeippen (2010) model.

The Storey & Zeippen (2010) atomic data provided very good agreement between observed and theoretical intensities of the strongest EUV lines, as shown in one of the benchmark works by Del Zanna (2011), based on a variety of sources, including *Hinode*/EIS. Del Zanna (2011) also benchmarked other atomic data for this ion, calculated by Gupta & Tayal (1998) and Aggarwal & Keenan (2005). Various shortcomings in these calculations were found. On the other hand, excellent agreement (to within a relative 10%) was found for the main lines observed by *Hinode*/EIS in an active region moss area and the Storey & Zeippen (2010) atomic data, already

⁷ www.apap-network.org

⁸ www.chiantidatabase.org

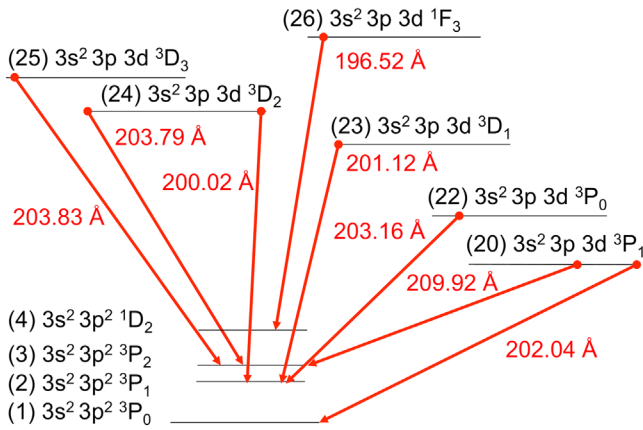



Figure 3. A simplified level diagram for the transitions relevant to the lines considered in this paper.

indicating an excellent accuracy in both the experimental and theoretical data, as we will also confirm below.

The Del Zanna (2011) benchmark work also reviewed all the previous identifications and determined which lines are likely to be blended, and hence avoided in our analysis. The main lines chosen for our study are listed in Table 1, in order of decreasing intensity. The main line is the straight decay to the ground state, which produces the emission line at 202.044 Å. We note that two transitions from the $3s\ 3p^2\ 3d\ ^3F_2$ were suggested to be blending the main density diagnostic line, already a self-blend at 203.8 Å. However, as noted by Del Zanna (2011), at the high densities that are considered here these lines do not have a significant contribution, so even if the identifications were incorrect the results presented here would still stand. The other lines (at 201.126, 200.021, 209.619, 203.165 Å) have a similar density sensitivity as the 203.8 Å one, except the 196.525 Å, which decays to a more excited level. The 203.165 Å was shown to be blended at low densities. Some lines (with wavelengths in brackets) were not considered since they are branching ratios (decays from the same upper level) with other lines we have included. Further benchmarks were carried out by Del Zanna (2012b).

The intensity of a spectral line is proportional to the population of the upper level and the spontaneous transition probability (the A -value). To assess which atomic rates affect a spectral line, it is therefore important to check which are the main populating mechanisms for each level. A simplified level diagram for the main transitions discussed in this section is provided in Figure 3, for a specific density. However, the main populating mechanisms for each atomic level normally vary with the density, so the issue can become quite complex to describe. Some details will be discussed in a separate paper, where also the various parameters that can affect an atomic calculation are reviewed. In summary, the 202.044 Å line is mainly populated by direct excitation from the $3s^2\ 3p^2\ ^3P_0$ ground state via a strong dipole-allowed transition. The various calculations provide the same rate, within a few percent. In turn, the population of the ground state decreases significantly as the population of the metastable levels increases. On the other hand, the populations of the other levels which produce the other lines in Table 1 are mainly driven by excitations from all the $3s^2\ 3p^2\ ^3P_{0,1,2}$ levels, although non-negligible contributions (typically 10%–30%) come also from cascading from higher levels.

Collision strengths at 2 MK – $3s^2\ 3p\ 3d$



Collision strengths at 2 MK, other $n=3$

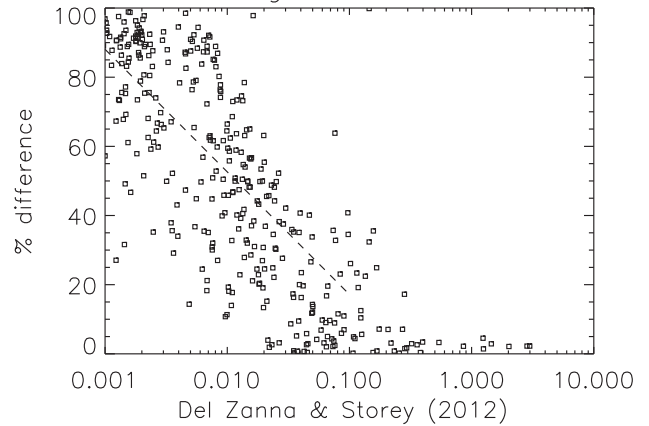


Figure 4. Percentage difference in the thermally averaged collision strengths (Storey & Zeppen 2010 vs. Del Zanna & Storey 2012) for a selection of transitions (top: to the $3s^2\ 3p\ 3d$ levels; bottom: to the other $n = 3$ levels). The dashed lines represent the approximate uncertainties used to generate alternative realizations of the atomic data for Fe XIII.

It is therefore important to first assess how accurate the rates of excitation from the $3s^2\ 3p^2\ ^3P_{0,1,2}$ to the $3s^2\ 3p\ 3d$ levels are. We have chosen to compare the latest values with those calculated by Storey & Zeppen (2010), because the two calculations were very similar, i.e., the main differences are caused by the size of the target and not by the method of the calculation. As already shown by Del Zanna & Storey (2012), the largest calculation provides very similar rates for the stronger lines, but significantly increased values for the weaker ones, as one would expect. We have considered only excitations from the $3s^2\ 3p^2\ ^3P_j$ and $3s^2\ 3p^2\ ^1D_2$ levels (the only ones with significant population) at the temperature of peak ion abundance in ionization equilibrium (2 MK, see Figure 4 top). As an estimate of the uncertainty in the strongest lines, with collision strengths above 1.0, we have taken 5%, which is well above the scatter of values. For the weaker lines, we have taken as an estimate the dashed line, i.e., a linear increase (up to a maximum of 50%).

We have then considered all the excitations to the remainder of the $n = 3$ levels calculated by Storey & Zeppen (2010), taking into account the different level orderings of the two calculations. In this case, we have taken a 10% uncertainty for the transitions above 0.1, and the linear increase shown in Figure 4 (bottom, up to a maximum of 50%).




Figure 5. Scatter plot of the thermally averaged collision strengths calculated by Del Zanna & Storey (2012) with the *R*-matrix codes (as in CHIANTI, present), vs. with the DW approximation. Dash and dot-dash lines indicate a $\pm 20\%$, 50% .

One possible estimate for all the $n = 3$ levels not included in Storey & Zeippen (2010) and all the $n = 4$ levels is to compare the full scattering calculation with the results of the distorted-wave (DW) calculation carried out by Del Zanna & Storey (2012), which does not include resonance enhancements (see Figure 5). We have taken a 20% uncertainty for the transitions above 0.01 and 50% for the weaker transitions.

The next step is to provide an estimate on the uncertainty of the *A*-values. As shown by Young (2004), different calculations can provide significantly different values. For our estimates, we have chosen to compare the Del Zanna & Storey (2012) *A*-values with those calculated by Young (2004) with the SUPERSTRUCTURE program (Eissner et al. 1974). An extended configuration set was used by Young (2004) to calculate radiative data for this ion. This data were made available within CHIANTI version 4 (Young et al. 2003) in 2003. Figure 6 shows comparisons of the *A*-values for all the transitions within the lowest 27 levels, which include the $3s^2 3p 3d$.

We have taken an uncertainty of 5% for transitions having an *A*-value above 10^{10} , 10% for those between 10^8 and 10^{10} , and 30% for weaker transitions. For the forbidden transitions within the ground configuration we have taken 10%.

We have modified the standard CHIANTI IDL routines distributed in SolarSoft (SSW, Freeland & Handy 1998) to assign to each transition an uncertainty in the *A*-value and in the excitation rate. We used the IDL function RANDOMN to randomly vary each rate within the estimated uncertainty. The distribution is normal, in the sense that if, e.g., an uncertainty is 10%, most values will vary within $\pm 20\%$. We then used the standard CHIANTI routine (`emiss_calc`) to calculate the line emissivities. For each of the seven chosen lines, we have added any Fe XIII lines within $\pm 0.1 \text{ \AA}$ to take self-blends into account. We have generated a total of 1000 realizations of the emissivities for each line, which are shown in Figure 7. The figure clearly shows how the spectral lines vary their emissivities as a function of the density. Figure 8 shows the variation of ratios with density. By attaching reasonable uncertainties to the atomic data we can generate realizations of the emissivities that capture this uncertainty. We can then use the ensemble of emissivities to characterize the uncertainties of the atomic data to infer physical parameters like plasma densities and column heights. Our methodology for combining

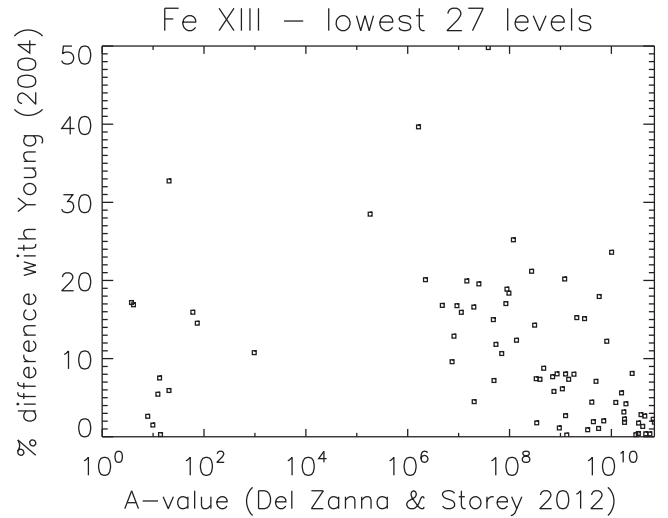


Figure 6. Percentage difference in the *A*-values calculated by Young (2004) and Del Zanna & Storey (2012) within the lowest 27 levels.

an ensemble of emissivities with observed data to account for uncertainties in atomic data is described in detail in Section 5.

4. Simulated Data Sets

The next step in this analysis is to generate sets of intensities from known densities and path lengths. This will allow us to test our ability to recover physical parameters from the Fe XIII intensities and to illustrate how the variations in the atomic data developed in the previous section led to variations in the inferred densities and path lengths. These intensities complement the set of observed intensities taken from the data illustrated in Figure 1.

The ratio curves shown in Figure 8 indicate that these lines are sensitive to density in the range of n_e of 10^8 – 10^{11} cm^{-3} . Thus, we randomly select 1000 densities uniformly on the interval $\log n_e = [8.5, 10.5]$. To continue with our theme of analyzing observations of active region moss we use the theoretical estimate of the moss path length from Martens et al. (2000) of

$$ds = \frac{2.56 \times 10^8}{P_0}, \quad (2)$$

where $P_0 = 2k_b n_e T_e$ is in dyne cm^{-2} , k_b is the Boltzmann constant, and ds is in cm. Again, we use the peak temperature of formation for Fe XIII, 1.8 MK, for this calculation. Note that this expression was derived for somewhat cooler emission and we do not expect it to track the Fe XIII path lengths exactly.

Each set of density and path length can be used to generate a set of intensities using statistics using the EIS pre-flight effective areas and the assumption of a 60 s exposure time and the $2''$ slit. Finally, we do not use these computed intensities directly when attempting to recover them with the model. We first apply a normally distributed random perturbation to each intensity, which mimics the variations in measured counts expected due to the finite exposure time.

An example set of simulated intensities is given in Table 3 for assumed values of $\log n_e = 9.90$ and $\log ds = 7.80$. If we use the standard set of CHIANTI emissivities to model these intensities we recover the input parameters almost exactly, $\log n_e = 9.90 \pm 0.01$ and $\log ds = 7.82 \pm 0.03$. As with the




Figure 7. Emissivities of the seven Fe XIII lines considered in this work. The emissivities are computed assuming a temperature of 1.8 MK, the temperature of formation for Fe XIII. The gray lines represent the 1000 realizations of the CHIANTI atomic data. The red curve is the default value from CHIANTI v.8. As discussed in Section 5, the blue curve is identified as being the most probable match to the observations (#471), and the green curve the second most probable (#368).

example set of observed intensities in Table 2, the uncertainties on these parameters are very small. Unlike the case with the observed intensities, however, we obtain a reduced χ^2 of order 1.

5. Inference

The standard method of χ^2 minimization (see Section 2) allows best-fit values of the plasma density n_e and column depth ds to be determined for a given pixel, under the assumption that the emissivity curves are completely and correctly specified. Now, equipped with 1000 data sets corresponding to randomly selected EIS pixels, we can consider the uncertainties in the fitted density and path length in each case that result from both statistical fluctuations in the observed intensities and the atomic data uncertainties incorporated in the ensemble of CHIANTI emissivities. To do so, we develop a Bayesian methodology that first assumes the observed data is uninformative regarding the atomic physics (the so-called *pragmatic Bayesian* method) and then incorporate the potential information in the observed data to learn about the atomic physics (the *fully Bayesian* method). The fully Bayesian method is a principled statistical analysis, while the pragmatic method makes simplifying assumptions that tend to overestimate the final uncertainty on the fitted density and path length. More details of the distinction between the two methods is discussed in Section 5.4. We start by providing an introduction to Bayesian inference in Section 5.1.

5.1. Bayesian Inference

We take a Bayesian approach in our statistical analysis because it enables us to build in the complex hierarchical dependencies engendered by atomic uncertainties. Such an approach offers a probability-based formalism for combining information from our prior knowledge and the current data. This requires both a *prior distribution*, which quantifies the uncertainty in the values of the unknown model parameters

before the data is observed, and a *likelihood function*—the distribution of the data given the model parameters. The likelihood function allows us to assess the viability of a parameter value given the observed data under a proposed statistical model. The likelihood function is combined with the prior distribution to yield the *posterior distribution*, which quantifies the uncertainty in the values of the unknown model parameters taking account of the observed data. If we let X and ψ represent generic data and unknown model parameters, respectively, Bayes’ theorem provides the posterior distribution as

$$p(\psi|X) = \frac{L(\psi|X) p(\psi)}{p(X)}, \quad (3)$$

where $L(\psi|X)$ is the likelihood of X given ψ (sometimes written as $p(X|\psi)$) and $P(\psi)$ the prior distribution of ψ . The term $P(X)$ is a normalizing constant necessary to make $p(\psi|X)$ a proper probability distribution. (The term $p(X)$ is sometimes referred to as the “evidence” in the astrophysics literature.) The posterior distribution, which combines information in the data with our prior knowledge, is our primary statistical tool for deriving parameter estimates and their uncertainties.

To perform a Bayesian analysis of the Fe XIII intensities, we start by defining notation and terminology in Section 5.2. We specify the likelihood function and the prior distribution in Section 5.3. In Section 5.4, we derive the posterior distribution under two sets of assumptions, which result in the aforementioned pragmatic Bayesian and fully Bayesian approaches. In Sections 5.5 and 5.6 we discuss our model-fitting routines, separate pixel-by-pixel and simultaneous analyses, where we consider the 1000 pixel data sets individually and simultaneously. In Sections 5.7 and 5.8 we apply our methodologies to the simulated and the observed intensities, respectively.




Figure 8. Theoretical ratios of six Fe XIII lines with respect to Fe XIII 202.044. See Figure 7 for additional details.

Table 3
Modeling Fe XIII Line Intensities in the Moss^a

| Line | $I_{\text{obs}} \sigma_I$ | I_{model} | $ \Delta I /I(\%)$ |
|---------|---------------------------|--------------------|--------------------|
| 196.525 | 751.0 ± 9.4 | 750.7 | 0.0 |
| 200.021 | 678.7 ± 11.8 | 694.7 | 2.4 |
| 201.121 | 750.7 ± 14.8 | 748.6 | 0.3 |
| 202.044 | 1012.5 ± 20.6 | 1011.9 | 0.1 |
| 203.165 | 313.7 ± 13.7 | 296.2 | 5.6 |
| 203.826 | 3533.1 ± 52.2 | 3498.8 | 1.0 |
| 209.916 | 168.4 ± 21.0 | 174.8 | 3.8 |

Note.

^a An example set of intensities computed from assumed values of $\log n_e = 9.90$ and $\log ds = 7.82$. The observed intensities include a normally distributed random perturbation. The best-fit density and path length are $\log n_e = 9.90 \pm 0.01$ and $\log ds = 7.82 \pm 0.03$. The format of the table is the same as Table 2.

5.2. Notation

Suppose that in each of $K = 1000$ pixels we observe the intensities of each of the J spectral lines with wavelengths $\Lambda = \{\lambda_1, \dots, \lambda_J\}$. Let $I_{k\lambda}$ be the observed intensity of the line with wavelength $\lambda \in \Lambda$ in pixel $k \in \{1, \dots, K\}$, $\sigma_{k\lambda}$ its known standard deviation, $D_k = (I_{k\lambda_1}, \dots, I_{k\lambda_J})$, and $\mathcal{D} = \{D_1, \dots, D_K\}$.

We also have a collection of $M = 1000$ realizations of the plasma emissivities, denoted by \mathcal{M} ,

$$\mathcal{M} = \{\epsilon_\lambda^{(m)}(n_{ek}, T_{ek}), \lambda \in \Lambda, m = 1, \dots, M\},$$

where n_{ek} and T_{ek} are the electron density and temperature for pixel k and m indexes the emissivity realization (i.e., emissivity curve, $\epsilon_\lambda^{(m)}(n_{ek}, T_{ek})$), with $m = 1$ corresponding to the default CHIANTI emissivities.

The expected intensity of the line with wavelength λ in pixel k can be rewritten (from Equation (1)) as $\epsilon_\lambda(n_{ek}, T_{ek})n_{ek}^2 ds_k$, where ds_k is the path length through the solar atmosphere for pixel k . Let $\theta_k = (\log n_{ek}, \log ds_k)$ be the plasma parameters in pixel k , and $\Theta = (\theta_1, \dots, \theta_K)$.

5.3. Statistical Model

The first step in specifying our statistical model is to construct the likelihood function. We model the intensities $I_{k\lambda}$ given m , n_{ek} , and ds_k as a normal (i.e., Gaussian) distribution,

$$I_{k\lambda}|m, n_{ek}, ds_k \stackrel{\text{indep}}{\sim} \mathcal{N}(\epsilon_\lambda^{(m)}(n_{ek}, T_{ek})n_{ek}^2 ds_k, \sigma_{k\lambda}^2), \quad (4)$$

for $\lambda \in \Lambda$, where $\mathcal{N}(\mu, \sigma^2)$ is a normal distribution with mean μ and variance σ^2 . We suppress the conditioning on the $\sigma_{k\lambda}$ throughout for notational simplicity. Thus, the likelihood function of D_k given emissivity index, m , and plasma parameters, θ_k , is

$$\begin{aligned} L(m, \theta_k|D_k) &= p(D_k|m, \theta_k) \\ &= \prod_{j=1}^J \mathcal{N}(I_{k\lambda_j} | \epsilon_{\lambda_j}^{(m)}(n_{ek}, T_{ek})n_{ek}^2 ds_k, \sigma_{k\lambda_j}^2), \end{aligned} \quad (5)$$

where $\mathcal{N}(x|\mu, \sigma^2)$ is the density of a normal distribution with mean μ and variance σ^2 evaluated at x . Note that we focus on methods that treat the emissivity index m as an unknown parameter, whose prior is specified below, whose posterior we estimate to determine the most likely emissivity realizations among those in \mathcal{M} , and whose uncertainties affect both the fit and error bars of θ_k .

Next, we specify the joint prior distribution on the unknown model parameters. For $\log n_{ek}$ and m we specify a continuous uniform distribution and a discrete uniform distribution, respectively,

$$p(\log n_{ek}) = \frac{1}{5} \quad \text{for } 7 \leq \log n_{ek} \leq 12, \quad (6)$$

$$p(m) = \frac{1}{M} \quad \text{for each } m = 1, \dots, M. \quad (7)$$

This choice of prior on m stipulates that the 1000 realizations of emissivity curves in \mathcal{M} are all a priori equally likely to be the true emissivity. As the realizations were generated by attaching reasonable uncertainties to the atomic data as described in Section 3, the atomic data uncertainties are contained in $p(m)$

and are thus captured by the corresponding posterior distribution. Therefore, the 1000 realizations of emissivity curves can also be considered as a sample of 1000 draws from an implicit prior distribution.

For $\log ds_k$, however, a uniform prior, $p(\log ds_k) \propto 1$, yields an improper posterior distribution because the likelihood converges to a positive constant as $\log ds_k$ goes to $-\infty$. Therefore, we specify a Cauchy distribution for $p(\log ds_k)$,

$$\log ds_k \sim \text{Cauchy}(\text{center} = 9, \text{scale} = 5). \quad (8)$$

which is a broad, fat-tailed distribution covering all conceivable values for the path length that we expect based on all sets of Fe XIII intensities, with an example set of intensities shown in Table 2.

We assume the parameters are a priori independent so that the joint prior distribution is

$$\begin{aligned} p(m, \theta_k) &= p(m) p(\theta_k) \\ &= p(m) p(\log n_{ek}) p(\log ds_k). \end{aligned} \quad (9)$$

Here θ is indexed by k , but m is not. This reflects the fact that, although θ_k vary among the pixels, we expect the true emissivity (i.e., the true value of m) to be an underlying physical quantity that is the same for all pixels.

We consider two ways to fit the plasma parameters, Θ , given the observed or simulated intensities, \mathcal{D} , while accounting for atomic uncertainty, \mathcal{M} . First we can analyze each pixel separately in a sequence of pixel-by-pixel analyses. Although this may yield different estimates of m , the index of the preferred emissivity curve among the pixels, it allows us to see if the intensities of each pixel give consistent information as to the best emissivity curve(s). Alternatively, we can simultaneously analyze the intensities from all the pixels to arrive at an overall estimate of the most likely emissivity curve. Using this strategy, uncertainty can be quantified with a list of the most likely emissivity realizations from \mathcal{M} (or their indices, m) along with their associated posterior probabilities.

We consider both the separate pixel-by-pixel and simultaneous analyses, and for each develop both pragmatic and fully Bayesian approaches. Specifically, Section 5.4 develops the pragmatic and fully Bayesian approaches to the pixel-by-pixel analyses and Section 5.5 describes the algorithms used to deploy these approaches. The simultaneous analysis and its algorithm are discussed in Section 5.6.

5.4. Pragmatic and Fully Bayesian Methods for Separate Pixel-by-pixel Analysis

Given the likelihood function in Equation (5) and the prior distribution in Equation (9), the joint posterior distribution for m and θ_k under the separate pixel-by-pixel analyses is

$$p(m, \theta_k|D_k) = \frac{L(m, \theta_k|D_k) p(m, \theta_k)}{p(D_k)}, \quad (10)$$

where $p(D_k) = \sum_{m=1}^M \int L(m, \theta_k|D_k) p(m, \theta_k) d\theta_k$.

Then the marginal posterior distribution $p(\theta_k|D_k)$ can be obtained by summing over m ,

$$p(\theta_k|D_k) = \sum_{m=1}^M p(m, \theta_k|D_k). \quad (11)$$

In this way, we are able to infer θ_k accounting for uncertainties of the atomic data via the ensemble in \mathcal{M} .

5.4.1. Pragmatic Bayesian Method

For the pragmatic Bayesian method, as described by Lee et al. (2011), we assume that the observed intensities are uninformative as to the most likely emissivities. That is, we do not take into account the information in the intensities for narrowing the uncertainty in the choice of emissivity realizations. Mathematically, this assumption can be written $p(m|D_k) = p(m)$, i.e., m and D_k are independent. Thus, the pragmatic Bayesian joint posterior distribution of m and θ_k is

$$p(m, \theta_k|D_k) = p(\theta_k|D_k, m) p(m|D_k) \quad (12)$$

$$= p(\theta_k|D_k, m) p(m), \quad (13)$$

and the marginal posterior distribution of θ_k (from Equation (11)) is

$$\begin{aligned} p(\theta_k|D_k) &= \sum_{m=1}^M p(m, \theta_k|D_k) \\ &= \sum_{m=1}^M p(\theta_k|D_k, m) p(m). \end{aligned} \quad (14)$$

The pragmatic Bayesian method accounts for atomic uncertainty in a conservative manner. The assumption that $p(m|D_k) = p(m)$ ignores information in the intensities, D_k , that may reduce uncertainty of atomic data represented by m and hence of θ_k . We now consider methods that allow D_k to be informative for m .

5.4.2. Fully Bayesian Method

In contrast to the pragmatic Bayesian method, the fully Bayesian method, as described by Xu et al. (2014), incorporates the potential information in the data (i.e., the intensities) to learn about m . The fully Bayesian joint posterior distribution of m and θ_k is given in Equation (12) and the marginal posterior distribution of θ_k is given by

$$\begin{aligned} p(\theta_k|D_k) &= \sum_{m=1}^M p(m, \theta_k|D_k) \\ &= \sum_{m=1}^M p(\theta_k|D_k, m) p(m|D_k) \end{aligned} \quad (15)$$

where each $p(\theta_k|D_k)$ is normalized so that $\sum_{m=1}^M p(m|D_k) = 1$.

Using Bayes' theorem, we can directly compute the probability of each emissivity realization, m , given the data in each pixel separately,

$$p(m|D_k) = \frac{p(D_k|m)}{\sum_{m=1}^M p(D_k|m)}. \quad (16)$$

This is the marginal posterior probability among those emissivity realizations in \mathcal{M} . Equation (16) holds because each of the m has the same prior probability (see Equation (7)).

The Bayesian posterior distribution in Equation (16) allows the observed intensities to be informative for the atomic physics, following the principles of Bayesian analysis (Xu et al. 2014). It enables us to use the intensities to determine which emissivity realizations are more or less likely and averages over (posterior) uncertainty in emissivity realizations.

5.5. Algorithms for the Separate Pixel-by-pixel Analyses

5.5.1. Algorithms for Pragmatic Bayesian in the Separate Pixel-by-pixel Analyses

The Metropolis–Hastings (MH) algorithm (e.g., Hastings 1970) is a general term for a family of Markov chain simulation methods that are useful for sampling from Bayesian posterior distributions. Let $p(\psi|X)$ be the target posterior distribution, using the notations in Section 5.1. A proposed ψ^* is sampled from a proposal distribution $q(\psi^*|\psi^{(t)})$ at iteration $t + 1$. Calculating the acceptance probability, $\rho = \frac{p(\psi^*|X)q(\psi^{(t)}|\psi^*)}{p(\psi^{(t)}|X)q(\psi^*|\psi^{(t)})}$, we set $\psi^{(t+1)} = \psi^*$ with probability $\min(\rho, 1)$ and set $\psi^{(t+1)} = \psi^{(t)}$ otherwise.

To obtain a Monte Carlo (MC) sample of (m, θ_k) from the pragmatic Bayesian posterior in Equation (13), we first obtain an MC sample of the emissivity index, $\{m^{(1)}, \dots, m^{(L)}\}$, from its prior distribution, Equation (7). For each $m^{(\ell)}$, with $\ell = 1, \dots, L$, we can then sample $\{\theta_k^{[\ell, t]}, t = 1, \dots, T\}$ from $p(\theta_k|m^{(\ell)}, D_k)$ using the MH algorithm. This requires that we specify the proposal distribution $q(\theta^*|\theta^{(t)})$. To do so, we first compute the value of θ_k that maximizes $\log p(\theta_k|m^{(\ell)}, D_k)$, i.e., the maximum a posteriori (MAP) estimates, $\hat{\theta}_k$, along with the 2×2 Hessian matrix evaluated at the mode $\hat{\theta}_k$, $H(\hat{\theta}_k)$, for each $m^{(\ell)}$. We then use $t_4(\theta_k|\hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})$ as the MH proposal distribution, where $t_\nu(x|\mu, \Sigma)$ is the density of a multivariate t distribution with ν degrees of freedom, mode μ , and scale matrix Σ , evaluated at x . This type of MH sampler is known as an independence sampler (Gilks et al. 1996). We run MH for T iterations, the last of which is taken as the MC sample corresponding to $m^{(\ell)}$, i.e., $\theta_k^{(\ell)} = \theta_k^{[\ell, T]}$.

5.5.2. Algorithms for Fully Bayesian in the Separate Pixel-by-pixel Analyses

In the fully Bayesian separate pixel-by-pixel analyses, our aim is to obtain an MC sample from the joint posterior distribution, Equation (12), and we propose three basic strategies for doing this: (i) two-step MC with MH, described in Section 5.5.3 and Appendix A, (ii) two-step MC with a Gaussian approximation, described in Appendix B, and (iii) Hamiltonian MC (HMC), described in Appendix C. Specifically, the first strategy uses the MH algorithm while the second strategy makes a Gaussian approximation to the conditional distribution of θ_k given the sampled emissivity realization m , respectively. Comparing the three strategies, the two-step MC with MH is preferred because of the accuracy of estimates with moderate computation time, while two-step MC with a Gaussian approximation may be faster (but less accurate) and HMC can be more accurate (but slower) under certain conditions.

5.5.3. Implementation of Two-step MC with MH for Fully Bayesian in the Separate Pixel-by-pixel Analyses

In order to implement the fully Bayesian method and to obtain an MC sample of θ_k via Equation (15), we first evaluate Equation (16) for each m where

$$p(D_k|m) = \int L(m, \theta_k|D_k) p(\theta_k) d\theta_k \quad (17)$$

is the Bayesian evidence conditional on a given emissivity. For each sampled θ_k , we need only evaluate the likelihood for

$m = 1, \dots, M$, and then renormalize the M likelihood values by this weighted sum, which can be achieved via a two-step sampling as described in this section.

The 2D integral in Equation (17) can be evaluated numerically using the grid generated from the Trapezoidal Quadrature Rule (TQR), which is suitable for finite domain quadrature.⁹ The Product Rule is also used in the construction of multivariate grids, which leads to an evenly designed grid.

The 2D quadrature can then be expressed as

$$\begin{aligned} & \int L(m, \theta_k|D_k) p(\theta_k) d\theta_k \\ &= \sum_{i,j} w_{i,j} L(m, \log n_{ek}^{(i)}, \log ds_k^{(j)}|D_k) p(\log n_{ek}^{(i)}, \log ds_k^{(j)}), \end{aligned} \quad (18)$$

where nodes $(\log n_{ek}^{(i)}, \log ds_k^{(j)})$ and weights $(w_{i,j})$ are defined by the chosen quadrature rule.¹⁰ The integral range of the two parameters is $(\hat{\theta}_k - 3 \times \text{sdev}_k, \hat{\theta}_k + 3 \times \text{sdev}_k)$ where sdev_k is a vector of the square root of the diagonal elements in variance-covariance matrix $(-H(\hat{\theta}_k))^{-1}$.

Having evaluated Equation (16) at each m , we can obtain an MC sample of the emissivity index, $\{m^{(1)}, \dots, m^{(L)}\}$. For each $m^{(\ell)}$ we sample from $p(\theta_k|D_k, m^{(\ell)})$ using an independence sampler exactly as described in Section 5.5.1. For each $m^{(\ell)}$, we run the independence sampler for T iterations to obtain the MC sample corresponding to $m^{(\ell)}$, $\theta_k^{(\ell)} = \theta_k^{[\ell, T]}$. The detailed two-step MC with MH (S_{MH}) is given in Appendix A.

5.6. Simultaneous Analysis

When we consider all the K -pixel intensities together in a simultaneous analysis using the fully Bayesian method, the likelihood function of m and Θ given \mathcal{D} , and the prior distribution of m and Θ are, respectively,

$$L(m, \Theta|\mathcal{D}) = \prod_{k=1}^K L(m, \theta_k|D_k) \quad (19)$$

and

$$p(m, \Theta) = p(m) \prod_{k=1}^K p(\theta_k). \quad (20)$$

Thus, the joint posterior distribution of m and Θ can be expressed as

$$p(m, \Theta|\mathcal{D}) = \frac{L(m, \Theta|\mathcal{D})p(m, \Theta)}{p(\mathcal{D})}, \quad (21)$$

where $p(\mathcal{D}) = \sum_{m=1}^M \int L(m, \Theta|\mathcal{D})p(m, \Theta) d\Theta$. Similarly, treating m as an unknown parameter, we express the left-hand side of Equation (21) as

$$p(m, \Theta|\mathcal{D}) = p(\Theta|\mathcal{D}, m) p(m|\mathcal{D}), \quad (22)$$

and we conduct statistical inference by obtaining an MC sample from this joint posterior distribution.

⁹ Package “mvQuad” provides a collection of methods for (potentially) multivariate quadrature in R, and is available at <https://cran.r-project.org/web/packages/mvQuad/>.

¹⁰ TQR and Product Rule are used in the construction of multivariate grids, where level = 5 is a subcommand in the grid creating commander, which represents accuracy level, typically number of evaluation points for the parameters in each dimension.

First we can use all the data simultaneously to obtain the marginal posterior probability of each emissivity realization m ,

$$p(m|D) = \frac{\prod_{k=1}^K p(D_k|m)}{\sum_{m=1}^M \prod_{k=1}^K p(D_k|m)}, \quad (23)$$

and sample $m^{(\ell)}$, for $\ell = 1, \dots, L$, with weights given by the marginal posterior probabilities in Equation (23) so that those favored by the data are sampled more frequently. The computation of $p(D_k|m)$ for each k and m is discussed in Section 5.5.3.

For each sampled m , we sample θ from its conditional posterior distribution

$$\begin{aligned} p(\Theta|D, m) &\propto L(m, \Theta|D) p(\Theta) \\ &= \prod_{k=1}^K L(m, \theta_k|D_k) p(\theta_k) \\ &= \prod_{k=1}^K \prod_{j=1}^J N(I_{k\lambda_j} | \epsilon_\lambda^{(m)}(n_{ek}, T_{ek}) n_{ek}^2 ds_k, \sigma_{k\lambda_j}^2) \\ &\quad \times p(\log n_{ek}) p(\log ds_k), \end{aligned} \quad (24)$$

as these K -pixel data sets were randomly selected from the observations indicated in Section 2, so that we can safely assume conditional independence among them.

Similarly, an MH sampler is used to obtain a correlated MC sample, $\{\Theta^{[t]}, t = 1, \dots, T\}$, from $p(\Theta|m^{(\ell)}, \mathcal{D})$. A $t_4(\theta_k|\hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})$ proposal distribution is used for each pixel independently and separately to make the computation more efficient. With this proposal distribution, we run the MH for T iterations over all the K -pixel intensities and obtain the MC sampler corresponding to $m^{(\ell)}$, $\Theta^{(\ell)} = \Theta^{[t, T]}$. The detailed two-step MC with MH via simultaneous analysis ($\mathcal{S}_{\text{MHsimul}}$) is given in Appendix D.

5.7. Application to Simulated Intensities

Here we illustrate both the separate pixel-by-pixel and the simultaneous analyses, mentioned in Sections 5.5 and 5.6, with a simulated case, using $K = 1000$ simulated sets of intensities for each of $J = 7$ spectral lines with known density and path lengths as described in Section 4. This will allow for the comparison of our inferred values with known values.

We run the separate pixel-by-pixel and simultaneous analyses described in Sections 5.5.3 and 5.6. For both analyses, 30 MH samplers, determined by constructing autocorrelation plots in this setting (Xu et al. 2014), are drawn for each sampled emissivity realization $m^{(\ell)}$, and the last MH sampler is taken as an MC sampler. There are 8000 MC samplers drawn in each simulation.

The comparison of the relative posterior probability $p(m|D_k)$ for each emissivity index and for each pixel, in separate pixel-by-pixel analyses, is shown in the left panel of Figure 10. The emissivity realization with index 1 occupies almost all of the probability. Similarly, in the simultaneous analysis, the posterior probability of the emissivity realization with index 1 is nearly to one. Both analyses recover the fact that all of the simulated sets of intensities are computed from the default CHIANTI atomic data (the emissivity realization with index 1) instead of the perturbed atomic data as described in Section 4.

Table 4

MSE Between the Fitted Values and the True Values for Both Parameters ($\log n_e, \log ds$) via Both Separate Pixel-by-pixel and Simultaneous Analyses

| | $\log n_e$ | $\log ds$ |
|--------------------------------|------------------------|------------------------|
| \mathcal{S}_{MH} | 1.345×10^{-5} | 4.936×10^{-5} |
| $\mathcal{S}_{\text{MHsimul}}$ | 6.748×10^{-7} | 2.241×10^{-6} |

Comparing the results from separate pixel-by-pixel and simultaneous analyses using their mean square errors (MSEs), a measure of how well the fitted values explain the given set of observations, Table 4 shows simultaneous analysis achieves smaller MSE values and indicates the more data we have, the smaller MSE is achieved, i.e., simultaneous analysis gives a better explanation of the given set of observations (i.e., intensities).

5.8. Application to Observed Intensities

Here we demonstrate the effects of the different types of analyses by applying them to a real data set, the EIS full-CCD observations of an active region used as an exemplar in Table 2 (EIS file eis_10_20130708_002042). This data set comprises sets of measured intensities of $J = 7$ spectral lines of Fe XIII in $K = 1000$ distinct, independent pixels. The results are shown in Figure 9 for the same pixel as exemplified in Table 3. The joint posterior probability density distribution $p(\theta_k|D_k)$ computed using the pragmatic and fully Bayesian methods are shown as contour plots, and marginalized 1D posterior densities $p(\log n_{ek}|D_k)$ and $p(\log ds_k|D_k)$ are shown as curves along the corresponding axes. The estimates of $\log n_{ek}$ and $\log ds_k$ computed via the standard analysis, i.e., the χ^2 minimization of Equation (1), are marked with straight lines. Notice that the pragmatic Bayesian method inflates the error bars relative to the standard method as it accounts for the atomic data uncertainties. The fully Bayesian method shrinks the error bars relative to the pragmatic Bayesian method and shifts the best estimate since it selects a subset of the full range of atomic uncertainties that are consistent with the data. The standard method underestimates the uncertainties in all cases, and is shifted relative to the fully Bayesian estimate.

The comparison of the relative posterior probability $p(m|D_k)$ for each emissivity index and for each pixel, in separate pixel-by-pixel analyses, is shown in the right panel of Figure 10. There are two dominant emissivity realizations that have a combined posterior probability of over 0.99 using the two-step MC with MH. An example of the posterior probability of the two dominant emissivity realizations given Pixel 593 is shown in Table 5. Similarly, in the simultaneous analysis, the posterior probability of the emissivity curve with index 471 is exactly one. It indicates that the emissivity realizations reveal a consistent feature of the solar atmosphere.

The computational time is considered in terms of (i) the elapsed time and (ii) the sum of the user and system times, which is a closer measure to real clock time. For the separate pixel-by-pixel analyses, the computation time over all 1000 pixels is about (i) 14.5 hr and (ii) 41.0 hr, respectively, for the two measures of computational time. For the simultaneous analysis, both time measurements are about 6.0 hr. These computation times consist of both the quadrature part and sampling part; the computation of the quadrature part is exactly the same for both




Figure 9. Comparisons of the inferred density n_e and path length ds using different methods for pixel #217 (for the intensities listed in Table 2). The results from the different methods used are color coded, with blue representing the standard method, red the pragmatic Bayes method (Section 5.4.1), and green the fully Bayes method (Section 5.4.2), with the latter two computed for all pixels simultaneously (see Section 5.6). The contour plots (with levels at $0.01\times$, $0.1\times$, and $0.5\times$ the maximum) show where the majority of the mass of the joint probability distributions of (n_e, ds) fall, and their marginalized distributions along each axis is shown to the top and to the right of the corresponding axis. The results from the standard analysis is shown along with the histograms as straight lines (solid for the best-fit and dashed denoting the $\pm 1\sigma$ errors on the best-fit obtained from the default CHIANTI emissivity functions), extending into the contour plot region. The best-fit value from standard analysis is also marked on the contour plot with a “+” sign, with the arms of the symbol corresponding to the sizes of the error bars. The standard deviations of the marginalized posterior densities, as well as the 90% equal-tail bounds for both the pragmatic and full Bayes cases are listed in the legend. As expected, the density and path length are highly correlated. The standard method underestimates the uncertainties, and the pragmatic Bayes method inflates them due to atomic data uncertainties. The fully Bayes method strikes a balance between atomic data uncertainties and how well the data are fit, shrinking the error bars relative to pragmatic Bayes and shifting the estimates.

(The complete figure set (1000 images) is available.)

the separate pixel-by-pixel and simultaneous analyses with a computation time of 1.2 hr for both measurements.

6. Conclusions and Discussion

We have presented the first comprehensive treatment of atomic physics uncertainties in the analysis of solar spectra. To make this analysis tractable, we have considered the relatively simple problem of inferring the electron density and path length from a set of observed Fe XIII intensities and a simple model for the emission (see Equation (1)). For this work we have used observed Fe XIII intensities from the EIS spectrometer on the *Hinode* satellite. If we consider only the uncertainties due to counting statistics, we obtain very small error bars on the electron density and path length, suggesting that the parameters are very precisely determined by the observations.

An essential component of this analysis is a model that we have constructed for the uncertainties in the collisional excitation and spontaneous decay rates. These rates are needed to compute the plasma emissivities that relate the observed intensities with the physical parameters of the plasma. This model for the uncertainties reflects the fact that for many transitions, such as those between the lower levels in Fe XIII, these rates appear to have converged. For other transitions, however, the rates are still highly uncertain. We have modified the CHIANTI software to produce self-consistent realizations of the atomic data based on this model for the uncertainties.

We have used a Bayesian framework to interpret the observed intensities in the context of the different realizations of the atomic data. A pragmatic Bayes approach, where each realization of emissivities is considered to be equally likely, yields uncertainties in the electron density and path length that

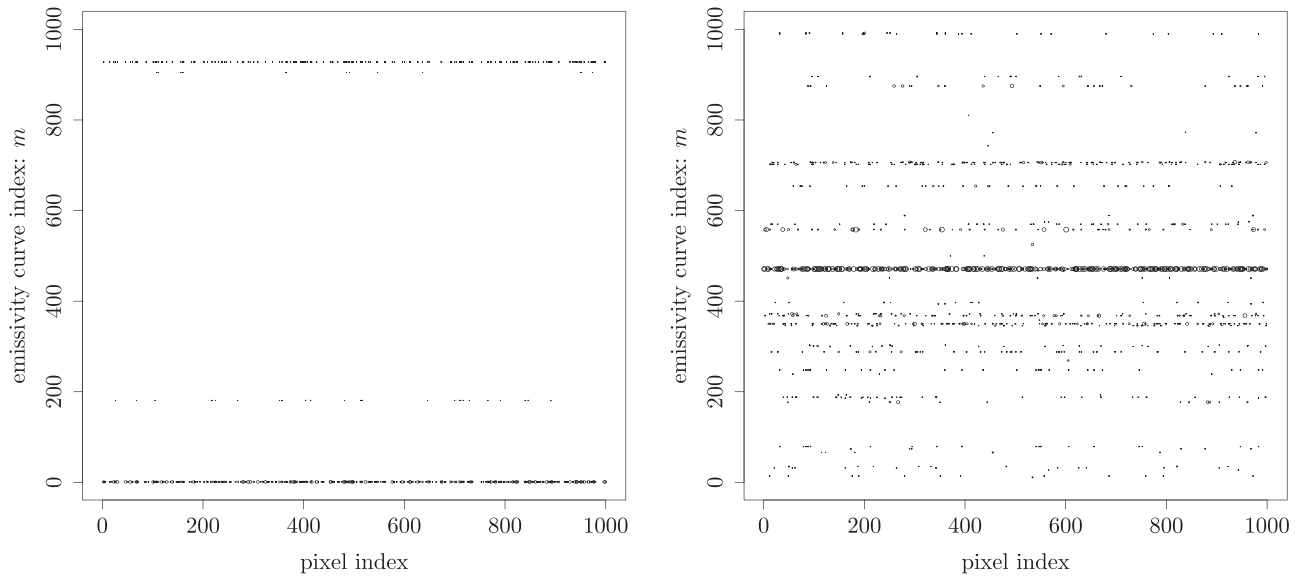


Figure 10. Selecting the optimal emissivity curves with a separate pixel-by-pixel fully Bayesian analysis. The x -axis and the y -axis represent the index of the pixels and the index of the emissivity curves, respectively. For each pixel, the relative posterior probability is plotted along a vertical column for the emissivity indices, where index 1 represents the default CHIANTI emissivities. The size of the dots represents the relative values of the posterior probability ($p(m|D_k)$ for the emissivity index m and pixel data set D_k) assigned to each emissivity index for a given pixel. The analyses carried out for the simulated data set (left; generated using default $m = 1$, and showing only $p(m|D_k) > 0.06$) and for a real data set (right; showing only $p(m|D_k) > 0.1$) are shown.

Table 5

The Posterior Probability of the Two Dominant Emissivity Realizations Given Pixel 593, $p(m|D_{593})$, Using the Two-step MC with MH via Both Separate Pixel-by-pixel and Simultaneous Analyses

| m | S_{MH} | $S_{MHsimul}$ |
|--------|----------|---------------|
| 471 | 0.894 | 1.000 |
| 368 | 0.105 | 0.000 |
| others | <0.001 | 0.000 |

are about a factor of 5 larger than the uncertainty implied by counting statistics alone. A fully Bayesian approach, where we allow the observed intensities to update the uncertainty in the emissivity curves, reduces the uncertainties in the plasma parameters, but also suggests that a different realization of the atomic data is more likely than the default CHIANTI calculation. This indicates some combination of systematic errors in the atomic physics, instrument calibration, and the observed intensities.

The methodology that we have developed here is both labor intensive and computationally demanding. Nevertheless, we believe that it represents a breakthrough in how atomic data uncertainties are brought into an analysis. Future improvements to the methodology and the structure of atomic databases will no doubt improve the process and make it more accessible. The extension to other emission lines would require an evaluation of the uncertainties in the collisional excitation and spontaneous decay rates similar to those described in Section 3 for each ion. Other uncertainties, such as those for the ionization and recombination rates needed to compute the ionization balances, also need to be addressed if lines from different ionization stages are considered. Once these uncertainty models are determined, we can only generate discrete realizations of the atomic data. This necessitates a brute force approach to computing the posterior which includes a sum over all of the realizations. The more common scenario is that the posterior is a continuous function of the parameters, which can be sampled

more easily. It is clear, however, that the uncertainty in the atomic data is often the dominant source of error in the analysis of solar spectra. Thus, this effort is essential to a rigorous analysis of the data.

Some constraints and the uncertainties in the atomic data could, in principle, be extracted from an analysis of the probability distribution $p(m|D)$ of the different realizations. In practice, however, our ability to consider this inverse problem is severely limited by the mismatch between the very large number of rates that go into calculating the level populations: the modeled line emissivities depend on 56,394 rates and their associated uncertainties. In principle, if all the transitions produced by the main levels in the ion could be observed, some constraints could be established. However, we only observe a very small number of emission lines.

Finally, we stress that the analysis presented here cannot overcome any limitations in the model used to interpret the observations. In this work, for example, we have assumed that the observed emission can be described by a simple model with a single density, temperature, and path length. Despite its simplicity, this model reproduces the observed intensities remarkably well. The path lengths, however, are relatively long ($ds \sim 10^9$ Mm) compared to the path lengths expected for the moss (see Equation (2)). It is likely that the observed emission is a combination of high density, short path length emission from the moss and low density, long path length emission from the overlying corona. To keep the analysis simple we have avoided using a more complex model. However, it would be necessary to consider more complex emission measure distributions if we seek to interpret the plasma parameters derived from the observations.

The authors acknowledge the generous support from the International Space Science Institute for hosting discussions among the ‘‘Improving the Analysis of Solar and Stellar Observations’’ international team. The full team consisted of Harry Warren (NRL; PI), Mark Weber (CfA; co-PI), Inigo Arregui (Inst. Astr. de Canarias), Frederic Auchere (Inst. Astr.

Spatiale), Connor Ballance (Queen’s Univ), Jessi Cisewski (Yale), Giulio Del Zanna (Cambridge), Veronique Delouille (Royal Obs. Belgium), Adam Foster (CfA), Chloe Guennou (Columbia), Vinay Kashyap (CfA), Fabio Reale (OAPA Palermo), Randall Smith (CfA), Nathan Stein (UPenn/Spotify), David Stenning (IPA/SAMSI/Imperial), and David van Dyk (Imperial). G.D.Z acknowledges support from STFC via the the University of Cambridge DAMTP astrophysics grant. The UK APAP network was funded during 2012–2015 by STFC (grant No. ST/J000892/1). H.P.W. was supported by NASA’s *Hinode* project. X.Y., D.S., and D.v.D. were supported by a Marie-Skodowska-Curie RISE (H2020-MSCA-RISE-2015-691164) Grant provided by the European Commission. M.W. was supported under the *Hinode*/XRT contract NNM07AB07C from MSFC/NASA to SAO. V.L.K. was supported by NASA contract NAS8-03060 to the *Chandra X-ray Center*, and acknowledges travel support from the Smithsonian Competitive Grants Program for Science Fund 40488100HH0043. CHIANTI is a collaborative project involving George Mason University, the University of Michigan (USA) and the University of Cambridge (UK). *Hinode* is a Japanese mission developed and launched by ISAS/JAXA, with NAOJ as domestic partner and NASA and STFC (UK) as international partners; it is operated by these agencies in co-operation with ESA and the NSC (Norway).

Software: mvQuad (Weiser 2016), CHIANTI (Version 4) (Young et al. 2003; Dere et al. 1997), CHIANTI (Version 8) (Del Zanna et al. 2015; Dere et al. 1997), SolarSoft (Freeland & Handy 1998), Stan (Carpenter et al. 2016).

Appendix A

A.1. Separate Analyses: Two-step MC with MH

For Pixel k , i.e., the k th set of intensities, the two-step MC with MH (\mathcal{S}_{MH}) proceeds for $\ell = 1, \dots, L$ with

Step 1: Sample $m^{(\ell)} \sim p(m|D_k)$ via Equation (16).

Step 2: For $t = 1, \dots, T - 1$,

Step 2.1: Sample $\theta_k^{\text{prop}1} \sim t_4(\theta_k|\hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})$ and compute

$$\rho = \frac{p(\theta_k^{\text{prop}1}|D_k, m^{(\ell)}) t_4(\theta_k^{\text{prop}1}|\hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})}{p(\theta_k^{[t]}|D_k, m^{(\ell)}) t_4(\theta_k^{\text{prop}1}|\hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})}. \quad (25)$$

Step 2.2: Set

$$\theta_k^{[t+1]} = \begin{cases} \theta_k^{\text{prop}1}, & \text{with probability } \min(\rho, 1) \\ \theta_k^{[t]}, & \text{otherwise} \end{cases}. \quad (26)$$

Step 3: Set $\theta_k^{(\ell)} = \theta_k^{[T]}$.

For simplicity at each iteration, if the sampled emissivity index in Step 1 is the same as the previous draw, we do not need to iterate MH to sample θ_k in Step 2 since we already have a good proposal distribution for the same target distribution. Moreover, if there does exist one dominant emissivity curve, e.g., there exists m^* such that $p(m^*|D_k) \geq 0.9999$, we only need to sample this m^* all the time.

Appendix B

B.1. Separate Analyses: Two-step MC with Gaussian Approximation

This is an alternative method to sample $p(m, \theta_k|D_k)$ based on Equations (15) and (16). As in Section 5.5.3, we can evaluate Equation (16) at each m and obtain an MC sample of the emissivity index, $\{m^{(1)}, \dots, m^{(L)}\}$. For each $m^{(\ell)}$, instead of using the MH algorithm, we can then sample from $p(\theta_k|D_k, m^{(\ell)})$ by considering an approximate algorithm via Gaussian approximation.

We can conduct a Gaussian approximation to $p(\theta_k|D_k, m^{(\ell)})$ with mean equal to the MAP estimates, $\hat{\theta}_k$, and variance-covariance matrix $(-H(\hat{\theta}_k))^{-1}$. Specifically, the Gaussian approximation distribution $\mathcal{N}(\theta_k|\hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})$ has the same mode and curvature as the target conditional distribution $p(\theta_k|D_k, m^{(\ell)})$. Thus, the two-step MC with Gaussian approximation (\mathcal{S}_G) proceeds for $\ell = 1, \dots, L$ with

Step 1: Sample $m^{(\ell)} \sim p(m|D_k)$ via Equation (16).

Step 2: Sample $\theta_k^{(\ell)} \sim \mathcal{N}(\theta_k|\hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})$, where $\hat{\theta}_k$ depends on $m^{(\ell)}$.

Similar to Appendix A, if there is one dominant emissivity curve, we only need to sample this dominant one all the time.

B.2. Results from the Simulated Set of Intensities and the Observed Intensities

Here we illustrate a two-step MC with Gaussian approximation using a simulated case and a realistic case as described in Sections 5.7 and 5.8.

For a two-step MC with Gaussian approximation, the same as the two-step MC with MH in Section 5.7, TQR and Product Rule are used in computing multivariate quadrature in Equation (17). Once we obtain an MC sample of emissivity index via Equation (16), a Gaussian approximation is conducted to $p(\theta_k|D_k, m^{(\ell)})$ for each sampled $m^{(\ell)}$ and each pixel D_k as described in Appendix B.1. There are 8000 MC samplers drawn for each pixel.

The plots to compare the relative posterior probability for each emissivity index and for each pixel are identical to those in the simulated and realistic cases in Sections 5.7 and 5.8.

In the realistic case, the computation time over all 1000 pixels is 8.0 hr or 20.7 hr, with respect to the elapsed time or the sum of user and system times, respectively. It consists of both the quadrature part and the sampling part, where the computation time of the quadrature part is the same as with the two-step MC with MH, 1.2 hr for both measurements.

Appendix C

C.1. Separate Analyses: Hamiltonian Monte Carlo

Another alternative method to obtain an MC sample from the joint posterior distribution in Equation (10) via the separate analyses is to start by obtaining a sample from their marginal posterior distribution,

$$\theta_k^{(1)}, \dots, \theta_k^{(L)} \sim p(\theta_k|D_k).$$

First, we rewrite

$$p(\theta_k|D_k) \propto L(\theta_k|D_k) p(\theta_k), \quad (27)$$

where

$$\begin{aligned} L(\theta_k|D_k) &= \sum_{m=1}^M L(m, \theta_k|D_k) p(m|\theta_k) \\ &= \frac{1}{M} \sum_{m=1}^M L(m, \theta_k|D_k) \\ &= \frac{1}{M} \sum_{m=1}^M \prod_{j=1}^J \mathcal{N}(I_{k\lambda_j} | \epsilon_\lambda^{(m)}(n_{ek}, T_{ek}) n_{ek}^2 ds_k, \sigma_{k\lambda_j}^2), \end{aligned} \quad (28)$$

since the prior independent assumption, $p(m|\theta_k) = p(m) = 1/M$, and the observation independent assumption among lines of wavelengths.

Evaluating $p(\theta_k|D_k)$ in this way we can use the Stan¹¹ software package (Carpenter et al. 2016) to obtain $\{\theta_k^{(1)}, \dots, \theta_k^{(L)}\}$ via HMC to sample directly from its marginal posterior distribution, Equation (27). However, we must analytically marginalize over m , via Equation (28), since it cannot accommodate discrete parameters.

With these MC samples $\{\theta_k^{(1)}, \dots, \theta_k^{(L)}\}$ in hand, we can sample m from its conditional posterior distribution,

$$\begin{aligned} p(m|\theta_k^{(\ell)}, D_k) &= \frac{p(m) L(m, \theta_k^{(\ell)}|D_k)}{\sum_{\tilde{m}=1}^M p(\tilde{m}) L(\tilde{m}, \theta_k^{(\ell)}|D_k)} \\ &= \frac{L(m, \theta_k^{(\ell)}|D_k)}{\sum_{\tilde{m}=1}^M L(\tilde{m}, \theta_k^{(\ell)}|D_k)}, \end{aligned} \quad (29)$$

for $\ell = 1, \dots, L$.

C.2. Sampling Multimodal Posterior Distributions with Stan

The simulation obtained in Appendix C.1 results in bimodal posterior distributions for $\log n_e$ and $\log ds$ for a couple of pixel data sets. Specifically, the two modes correspond to the two different emissivity curves. The resulting relative size of the two modes does not match the actual posterior distributions indicating that the HMC algorithm has trouble in jumping between the modes. This multiple-mode problem may be due to an insufficient number of emissivity curves because our set of emissivities sample the full uncertainty range sparsely. To solve this problem, we have experimented with adding a few strategically chosen synthetic emissivity curves to the set \mathcal{M} and the augmented set of curves is denoted by \mathcal{M}^{aug} , where \mathcal{M} is a subset of \mathcal{M}^{aug} , i.e., $\mathcal{M} \subset \mathcal{M}^{\text{aug}}$. These tend to connect the modes and allow HMC to jump between modes. We can then remove the samples associated with the synthetic emissivity curves to get MC samples purely from the original target.

We run the algorithm described in Appendix C.1 with \mathcal{M} replaced by \mathcal{M}^{aug} . For each sampled value of $\theta_k^{(\ell)}$, $\ell = 1, \dots, L$, we compute $p(m|\theta_k^{(\ell)}, D_k)$ for each $m \in \mathcal{M}^{\text{aug}}$, with \mathcal{M} replaced by \mathcal{M}^{aug} in Equation (29), and sample a value of m , say $m^{(\ell)}$, from it. Once we have these sample values of m , $m^{(\ell)}$, for $\ell = 1, \dots, L$, we can then extract the samples of θ_k that

correspond to the non-synthetic emissivity curves to get MC samples purely from the original target, i.e., consider the conditional posterior distribution $p(m|\theta_k^{(\ell)}, D_k)$ for each $m \in \mathcal{M}$.

This creative method of adding synthetic emissivity curves in HMC can be generalized to all pixel data sets. If all the multiple-mode pixels have two modes and these two modes depend on the two same emissivity curves, the same synthetic emissivity curves can be added into the original ones and the above procedure can be repeated to all pixel data sets.

C.3. Results from the Simulated Set of Intensities and the Observed Intensities

Here we illustrate HMC with Stan through a simulated case and a realistic case as described in Sections 5.7 and 5.8.

For HMC with Stan (\mathcal{H}), a few strategically chosen synthetic emissivity curves are added, as described in Appendices C.1 and C.2. There are 5 chains running, 4000 iterations each, and the first half of the iterations of each chain are discarded as burn-in.

In the simulated case, the comparison of the relative posterior probability $p(m|D_k)$ for each emissivity index and for each pixel shows the emissivity curve with index 1 occupies almost all of the probability which also recovers the fact that all of the simulated sets of intensities are computed from the actual CHIANTI atomic data (the emissivity curve with index 1 instead of the perturbed atomic data as described in Section 4).

In the realistic case, once we run HMC with Stan as described in Appendix C.1, bimodal distributions appear for several of the pixels. The two modes correspond to two different emissivity curves with index 471 and 368, i.e., Emis_{471} and Emis_{368} . Moreover, the relative size of the two modes does not match the actual posterior distribution as shown in the left column of Figure 11. Therefore, a few strategically chosen synthetic emissivity curves are added to the original set and the augmented set is

$$\mathcal{M}^{\text{aug}}/\mathcal{M} = \{w_1 * \text{Emis}_{471} + w_2 * \text{Emis}_{368}\}$$

where $(w_1, w_2) = (0.75, 0.25)$, $(0.50, 0.50)$, and $(0.25, 0.75)$. The HMC with Stan is run once more with \mathcal{M} replaced by \mathcal{M}^{aug} as described in Appendix C.2. Samples of $\theta_k^{(\ell)}$, $\ell = 1, \dots, L$, are obtained as shown in the middle column of Figure 11. For each sampled value of $\theta_k^{(\ell)}$, we compute $p(m|\theta_k^{(\ell)}, D_k)$ for each $m \in \mathcal{M}^{\text{aug}}$, via Equation (29), and sample a corresponding $m^{(\ell)}$ from it. Considering the conditional posterior distribution $p(m|\theta_k^{(\ell)}, D_k)$ for each $m \in \mathcal{M}$, we can then extract the samples $\theta_k^{(\ell)}$ that correspond to the non-synthetic emissivity curves to get MC samples purely from the original target as shown in the right column of Figure 11. The computation time over all 1000 pixels is 51.4 hr or 135.5 hr, with respect to the two ways of measuring the computation time, the elapsed time or the sum of user and system times, respectively.

Appendix D

D.1. Simultaneous Analysis

The two-step MC with MH via simultaneous analysis ($\mathcal{S}_{\text{MHsimul}}$) proceeds for $\ell = 1, \dots, L$ with

Step 1: Sample $m^{(\ell)} \sim p(m|D)$ via Equation (16).

¹¹ Stan is a probabilistic modeling language developed by Andrew Gelman and collaborators. It interfaces with the most popular data analysis languages like R, Python, etc., and is available at mc-stan.org.

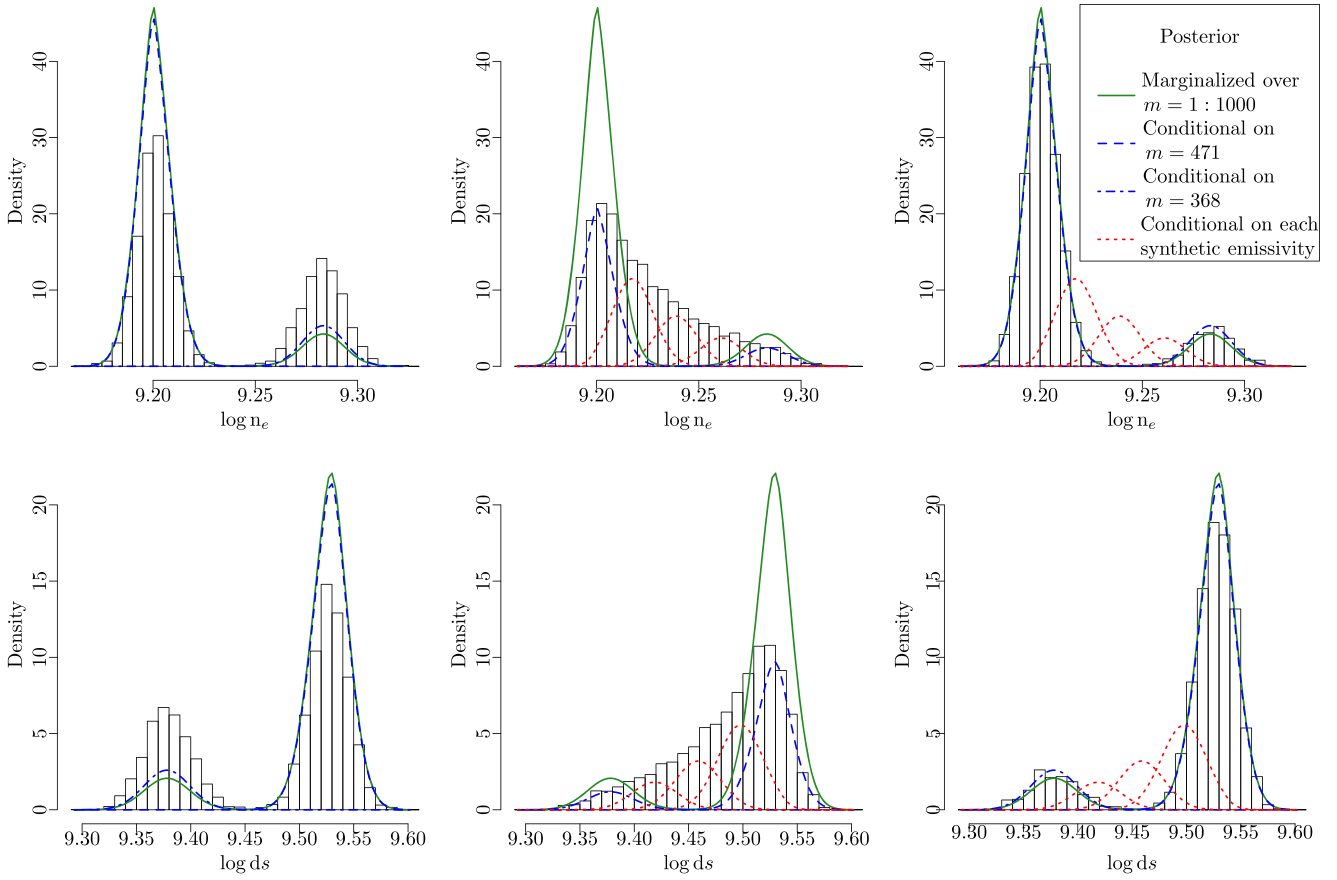


Figure 11. Demonstrating the augmented emissivity curves method to correct HMC/Stan analysis. The posterior probability densities are shown as calculated for the pixel 593 data set for the $\log n_e$ (top row) and $\log ds$ (bottom row). In all cases, the exact distributions, marginalized over the ensemble set of emissivity curves, are shown as solid green curves. The distributions conditional on one of the two most likely emissivity candidates, Emis_{471} and Emis_{368} , are shown as the dashed and dot-dashed blue lines, respectively, and are renormalized with their corresponding posterior weights. The histogram represents the posterior density distributions computed using HMC via Stan. Going from left to right in each row, the left panels depict the problem that HMC produces modes of different heights compared to the exact calculation; the middle panels show its deformation as the augmented emissivity sample \mathcal{M}^{aug} is used where the posterior distributions conditional on each of the augmented emissivities are shown as the red dotted curves and are renormalized with their corresponding posterior weights; and the right panels show the corrected versions after removing the augmented emissivities.

Step 2: Proceed for $t = 1, \dots, T$,

Step 2.1: For each pixel $k = 1, \dots, K$, sample $\theta_k^{[\text{prop}]}$ \sim $t_4(\theta_k | \hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})$ and set $\Theta^{[\text{prop}]} = (\theta_1^{[\text{prop}]}, \dots, \theta_K^{[\text{prop}]})$

Step 2.2: Compute

$$\rho = \frac{\prod_{k=1}^K p(\theta_k^{[\text{prop}]} | D_k, m^{(\ell)}) \cdot \prod_{k=1}^K t_4(\theta_k^{[t]} | \hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})}{\prod_{k=1}^K p(\theta_k^{[t]} | D_k, m^{(\ell)}) \cdot \prod_{k=1}^K t_4(\theta_k^{[\text{prop}]} | \hat{\theta}_k, (-H(\hat{\theta}_k))^{-1})}. \quad (30)$$

Step 2.3: Set

$$\Theta^{[t+1]} = \begin{cases} \Theta^{[\text{prop}]}, & \text{with probability } \min(\rho, 1) \\ \Theta^{[t]}, & \text{otherwise} \end{cases}. \quad (31)$$

Step 3: Set $\Theta^{(\ell)} = \Theta^{[T]}$.

Similar to the separate analyses, for simplicity at each iteration, if the sampled emissivity index in Step 1 is not updated, we do not need to iterate MH to sample each θ_k in Step 2. If one dominant emissivity curve exists, we only need to sample the dominant all the time.

The results in Figure 12 compare the fitted value using a two-step MC with MH to the true value of both parameters $\log n_e$ (right) and $\log ds$ (left) via both separate pixel-by-pixel

(top row) and simultaneous (bottom row) analyses. The gray lines represent the vertical error of one standard deviation. The dashed line represents equality, where the fitted value is identical to the true value. Compared with the separate pixel-by-pixel analyses, it shows that the error bars are smaller around the truth when we use the simultaneous analysis than when we use one pixel data set at a time. The results in the plots illustrate that as more data are used in the analysis by simultaneously analyzing those pixels, incorporating the uncertainty in the atomic physics calculations results in more accurate fitted values.

Appendix E

E.1. Comparison of Algorithms and Output Data Analysis

To obtain an MC sample of the parameters, $\log n_{ek}$ and $\log ds_k$, via the separate pixel-by-pixel analyses with joint posterior distribution in Equation (10), three algorithms were implemented for (the) fully Bayesian model on each of the 1000 pixel observed data sets: \mathcal{S}_G in Appendix B.1, \mathcal{S}_{MH} in Section 5.5.3, and \mathcal{H} in Appendix C.1.




Figure 12. Comparison of best-fit values and actual input for the simulated data set. The comparisons are shown for both $\log n_e$ (left column) and $\log ds$ (right column). Calculations are performed using a fully Bayesian two-step MC with MH for each pixel data set separately (top row) and for all the pixel data sets simultaneously (bottom row). The red dots represent the difference between the best-fit value and the actual input and the horizontal dashed lines represent the line of equality. The gray dots represent a vertical error of ± 1 standard deviation for the fitted values that incorporates atomic data uncertainty. Notice that the uncertainties are reduced when all the K -pixel data sets are used simultaneously.

Our aim is to find which algorithm provides a more accurate simulation to the target posterior distribution and is the best to use for statistical inference. From a statistical point of view, we assume the HMC, which might give the best result, as the base line, and to see whether these two two-step MC samplers provide better inference or not.

The first test statistic we consider is the z -statistic, which is the difference in posterior mean between the sample values from \mathcal{S}_G or \mathcal{S}_{MH} and from \mathcal{H} divided by the standard deviation of \mathcal{H} because HMC is assumed to be the base line, indicating how far away that estimate is from the mean in standard units, i.e.,

$$z_{\text{score}}^i = \frac{\text{mean}_{\mathcal{S}_i} - \text{mean}_{\mathcal{H}}}{\text{sd}_{\mathcal{H}}}, \text{ for } i = G \text{ or MH.} \quad (32)$$

Figure 13 shows the histograms of z -scores for both parameters, $\log n_e$ (top row) and $\log ds$ (bottom row), in two comparisons (left column: \mathcal{S}_G to \mathcal{H} , right column: \mathcal{S}_{MH} to \mathcal{H}), respectively, considering all the 1000 pixels. Looking at the worst case scenarios, the most extremes we see from the comparison on the left-hand side is about 0.12 to 0.25 of standard deviation off, which corresponds to Pixel 36, 87, 302, 453, 650, and 934. The comparison on the right-hand side indicates that most extremes are

about 0.15 of standard deviation off occurring at Pixel 302 and 364. The vertical lines correspond to the z -scores values of these extracted pixels. This suggests that we need to look at the full posterior distributions for those extreme pixels and for the three algorithms more closely, which can be found in Appendix E.3, to get some insights.

The second test statistic to compare is the ratio of standard deviations between \mathcal{S}_G or \mathcal{S}_{MH} and \mathcal{H} , i.e.,

$$\frac{\text{sd}_{\mathcal{S}_i}}{\text{sd}_{\mathcal{H}}}, \text{ for } i = G \text{ or MH,} \quad (33)$$

which essentially gives the relative size of confidence intervals that we compute.

Figure 14 shows the histograms of the ratio of standard deviations for both parameters, $\log n_e$ (top row) and $\log ds$ (bottom row), in two comparisons (left column: \mathcal{S}_G to \mathcal{H} , right column: \mathcal{S}_{MH} to \mathcal{H}), respectively, considering all the 1000 pixels. The most extremes we see from the comparison on the left-hand side corresponds to Pixel 634 and 779. The comparison on the right-hand side indicates the most extremes occurring at Pixel 396, 418, 634, and 779. The vertical lines correspond to the ratio values of these extracted pixels. An

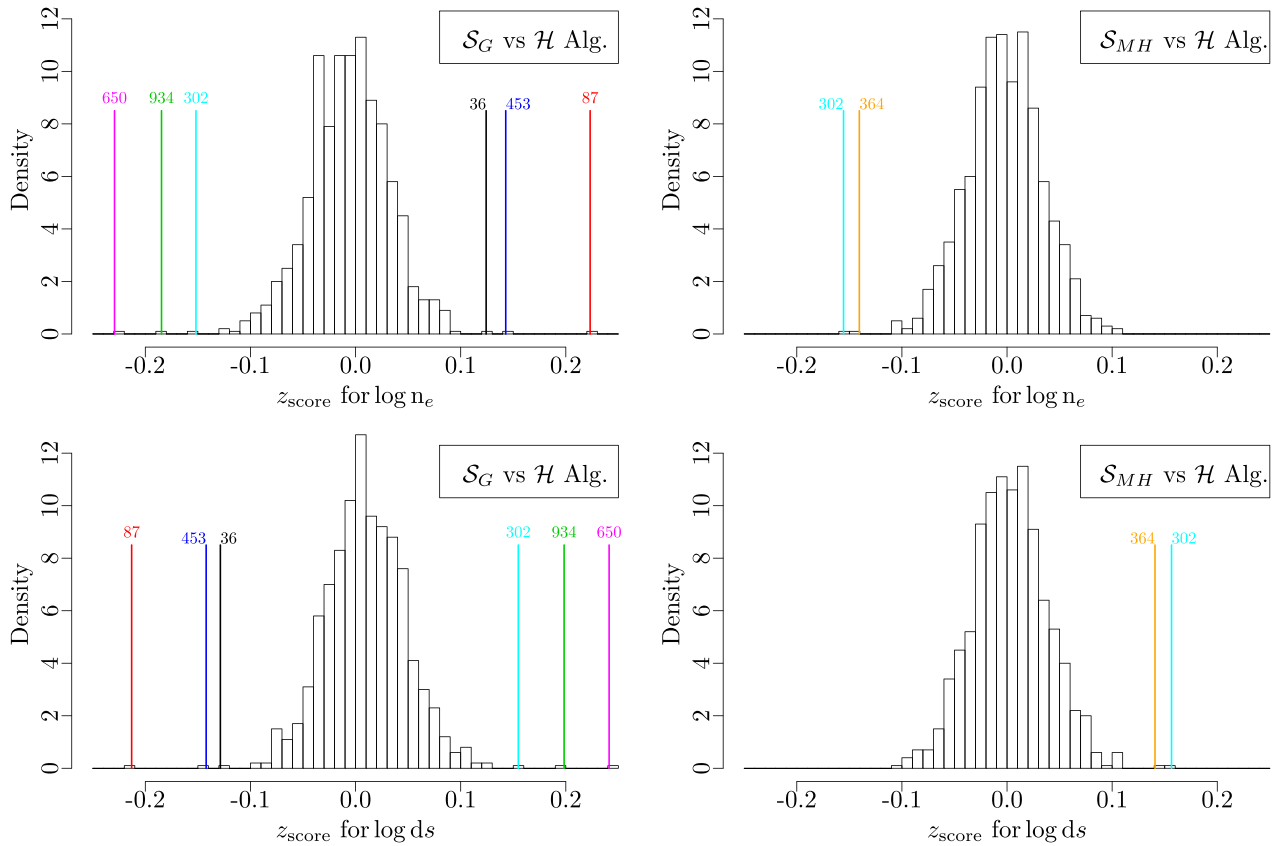


Figure 13. Selecting the extreme pixels via z -scores (unit difference in posterior mean) that are computed using the outputs of the three algorithms, \mathcal{S}_G , \mathcal{S}_{MH} , and \mathcal{H} under a fully Bayesian method and separate pixel-by-pixel analyses. The histograms represent the z -scores for both parameters, $\log n_e$ (top row) and $\log ds$ (bottom row), in two comparisons (left: \mathcal{S}_G to \mathcal{H} , right: \mathcal{S}_{MH} to \mathcal{H}), respectively, considering all the 1000 pixels. The vertical lines correspond to the values of the pixel indices, top left: 650, 934, 302, 36, 453, 87, top right: 302, 364, bottom left: 87, 453, 36, 302, 934, 650, bottom right: 364, 302, from left to right.

example of their posterior distributions for the three algorithms can be found in Appendix E.3.

E.2. Parallelization

To improve the efficiency of the code, we parallelize the 1000 pixels into 20 or 10 completely separate processes when pre-processing emissivities (i.e., obtaining the posterior probability of each emissivity curve) or sampling θ , for all the three algorithms. The `doParallel` package is used to provide a mechanism to execute for each loops in parallel within each process, where a multi-core backend is registered and a four worker cluster (of a 64-bit 2.5 GHz CPU with 128 GB of RAM) is created and used. Specifically, in the source builds, we set the number of processors to use for the build to the number of cores on our machine we want to devote to the build, which is thirty-two. We also set the maximum allowed number of additional R processes allowed to be run in parallel to the current R processes, which is thirty-two as well. For \mathcal{H} , each pixel is run with multiple cores and four pixels are run at the same time. For \mathcal{S}_G or \mathcal{S}_{MH} , we run each pixel with a different core and thirty-two multi-core backends are used in parallel.

E.3. The Posterior Values of the Parameters for the Three Algorithms and for the Extracted Pixels

By comparing the three algorithms using the two test statistics mentioned in Appendix E.1, several extreme pixels are picked out from each comparison.

Figure 15 show the histograms of the posterior values of the parameters $\log n_e$ (left) and $\log ds$ (right) conditional on all 1000 emissivity curves and the certain extracted pixel data set, respectively. The sampling algorithms used are \mathcal{S}_G algorithm (top row), \mathcal{S}_{MH} algorithm (middle row), and \mathcal{H} algorithm (bottom row). Three more synthetic emissivity curves are conditioned when using \mathcal{H} as described in Appendix C.3.

For Pixel 364 (the left panel of Figure 15), which are extracted from the right column of Figure 13, having used the synthetic emissivity curves, it is still not a very great job of jumping between the modes for the \mathcal{H} algorithm in this bimodal case.

For Pixel 396 (the middle panel of Figure 15), it is the histograms of the \mathcal{H} algorithm that do not quite get into the tail that makes the standard deviation from the \mathcal{H} algorithm relatively small and filters this pixel out from the right column of Figure 14.

Similarly, for Pixel 650 (the right panel of Figure 15), which are extracted from the left column of Figure 13, the \mathcal{S}_G algorithm does not do a great job at recovering the actual posterior with a noticeable discrepancy in the mode.

E.4. Discussion

As an example of the posterior distribution of emissivity curves given for Pixel 593, we get these two dominant emissivity curves from both two-step MC samplers and HMC. Considering the two of them, it is nearly all the probability up to 0.99, as is shown in Table 6.

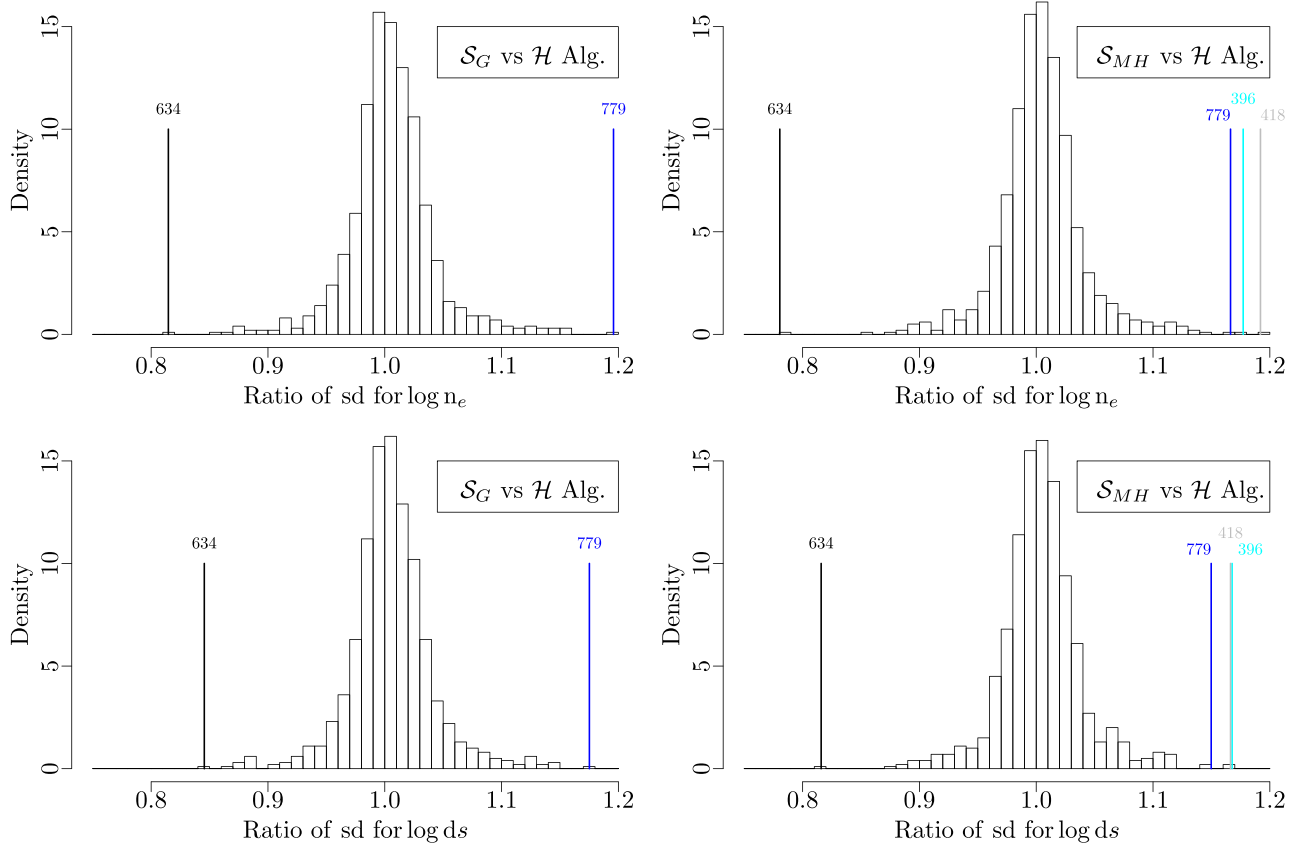


Figure 14. Selecting the extreme pixels via the ratio of standard deviations that are computed using the outputs of the three algorithms, S_G , S_{MH} , and \mathcal{H} under the Fully Bayesian method and separate pixel-by-pixel analyses. The histograms represent the ratio of standard deviations for both parameters, $\log n_e$ (top row) and $\log d_s$ (bottom row), in two comparisons (left: S_G to \mathcal{H} , right: S_{MH} to \mathcal{H}), respectively, considering all the 1000 pixels. The vertical lines correspond to the values of pixel indices, top left: 634, 779, top right: 634, 779, 396, 418, bottom left: 634, 779, bottom right: 634, 779, 418, 396, from left to right.

Table 6

The Posterior Probability of the Two Dominant Emissivity Curves Given Pixel 593 Using Two-step MC and HMC via Separate Pixel-by-pixel Analyses under a Fully Bayesian Method

| m | HMC with Stan (\mathcal{H}) | 2stepMC with S_G or S_{MH} |
|--------|---------------------------------|--------------------------------|
| 471 | 0.860 | 0.894 |
| 368 | 0.138 | 0.105 |
| others | <0.0019 | <0.0016 |

Table 7

Computation Time for the Three Algorithms in Realistic Case under a Fully Bayesian Method and Separate Pixel-by-pixel Analyses

| Algorithm | Elapsed (hr) | Sum of User and System (hr) |
|---------------|--------------|-----------------------------|
| S_{MH} | 14.5 | 41.0 |
| S_G | 8.0 | 20.7 |
| \mathcal{H} | 51.4 | 135.5 |

The computation time, in the realistic case, over all 1000 pixels for the three algorithms are shown in Table 7. Two ways of measuring the computation time are included, the elapsed time and the sum of user and system times. Moreover, the computation time of both the two-step MC samplers consist of both the quadrature part and sampling part, where the computation times of the quadrature part is 1.2 hr for both measurements.

We actually have two basic strategies for obtaining an MC sample from the joint posterior distribution, HMC and two-step MC sampler. By comparing the histograms of the posterior values, there is definitely an issue with the Gaussian assumption (two-step MC with Gaussian approximation) where the MC samplers are not matching very well with the actual posterior and it is more conservative. The HMC algorithm looks appropriate but occasionally does not give the relative size of the mode correctly, even after adding synthetic emissivity curves. For all of the 1000 pixels, the MC samplers generated from the two-step MC with MH match the density line of the actual posterior very well and this algorithm takes a moderate computation time. Therefore, being more accurate and significantly faster, a two-step MC with MH is preferable to use for statistical inference.

In our experiment, there are $J = 7$ spectral lines with corresponding wavelengths being considered; we call the two that are not close to the others in wavelength, 196.525 and 209.916 versus 200.021–203.826 Å, “extreme wavelengths.” We have experimented with one of the two-mode-case pixels (Pixel 593), where the two extreme wavelengths are removed one at a time from the analysis and the three algorithms mentioned in Section 5.5.3, Appendices B.1, and C.1 are repeated. Whether we consider the two extreme wavelengths or not, the resulting MC samplers have a good match to their actual posterior distributions; however, the shape of the actual posterior distribution differs dramatically when including wavelength 196.525 Å compared to when it is excluded from

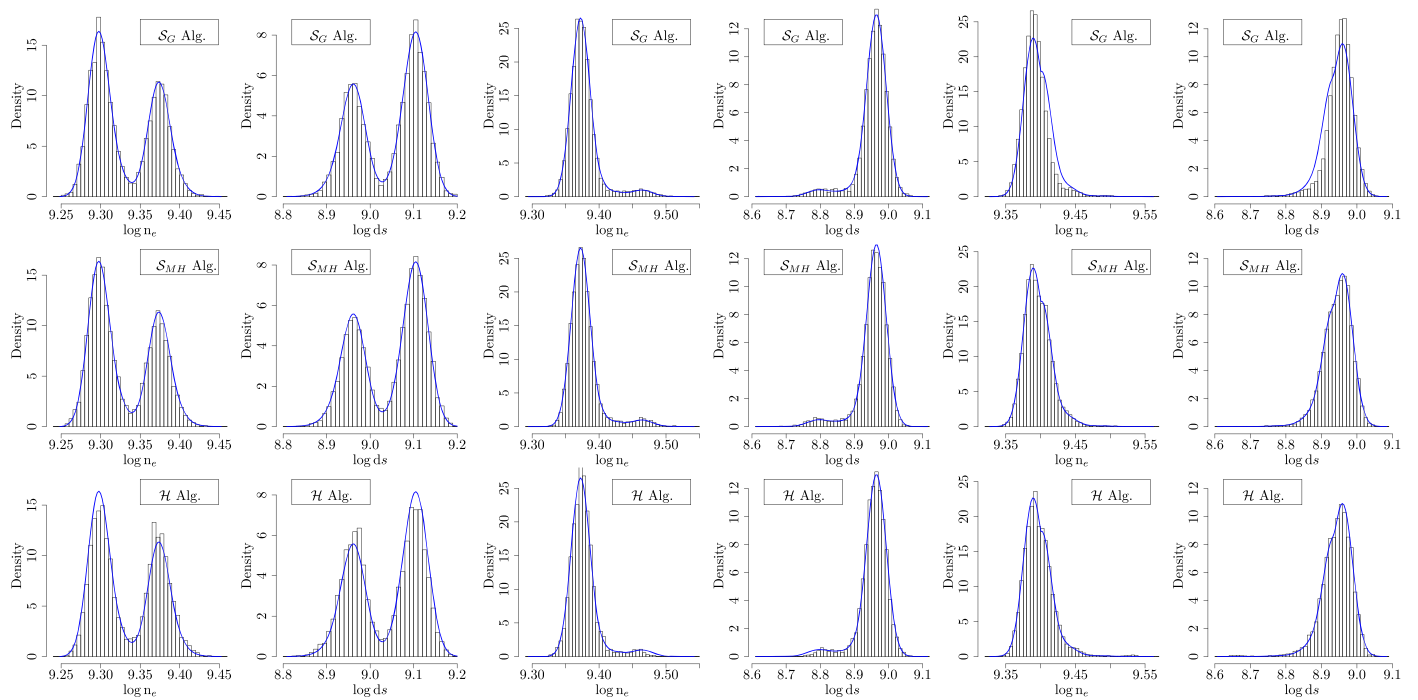


Figure 15. Selecting the optimal algorithm from the full posterior distribution for the selected extreme pixels under a fully Bayesian method. The histograms represent the posterior distributions of $\log n_e$ (left column of each panel) and $\log ds$ (right column of each panel) conditional on all 1000 emissivity curves and Pixel 364 (left panel), or 396 (medium panel), or 650 (right panel) in the real observed data set, respectively. The sampling algorithms compared are S_G (first row), S_{MH} (second row), and \mathcal{H} (third row). The actual conditional posterior distributions for each parameter are shown as the blue lines.

the analysis. Because including the extreme wavelengths did not impact the ability of the MC samplers to recover the actual posterior distributions, we used the seven-wavelength data set in all of the experiments.

ORCID iDs

Giulio Del Zanna <https://orcid.org/0000-0002-4125-0204>
 David C. Stenning <https://orcid.org/0000-0002-9761-4353>
 Jessi Cisewski-Kehe <https://orcid.org/000-0002-9656-2272>
 Vinay L. Kashyap <https://orcid.org/0000-0002-3869-7996>
 Harry P. Warren <https://orcid.org/0000-0001-6102-6851>

References

- Aggarwal, K. M., & Keenan, F. P. 2005, *A&A*, **429**, 1117
 Badnell, N. R., Del Zanna, G., Fernández-Mencheró, L., et al. 2016, *JPhB*, **49**, 094001
 Bautista, M. A., Fivet, V., Quinet, P., et al. 2013, *ApJ*, **770**, 15
 Berger, T. E., de Pontieu, B., Fletcher, L., et al. 1999, *SoPh*, **190**, 409
 Carpenter, B., Gelman, A., Hoffman, M., et al. 2016, *JOSS*, **20**, 1
 Chung, H.-K., Braams, B. J., Bartschat, K., et al. 2016, *JPhD*, **49**, 363002
 Culhane, J. L., Harra, L. K., James, A. M., et al. 2007, *SoPh*, **243**, 19
 Del Zanna, G. 2011, *A&A*, **533**, A12
 Del Zanna, G. 2012a, *A&A*, **546**, A97
 Del Zanna, G. 2012b, *A&A*, **537**, A38
 Del Zanna, G. 2013, *A&A*, **558**, A73
 Del Zanna, G., Berrington, K. A., & Mason, H. E. 2004, *A&A*, **422**, 731
 Del Zanna, G., Dere, K. P., Young, P. R., Landi, E., & Mason, H. E. 2015, *A&A*, **582**, A56
 Del Zanna, G., & Mason, H. E. 2014, *A&A*, **565**, A14
 Del Zanna, G., & Storey, P. J. 2012, *A&A*, **543**, A144
 Dere, K. P., Landi, E., Mason, H. E., Monsignori Fossi, B. C., & Young, P. R. 1997, *A&AS*, **125**, 149
 Eissner, W., Jones, M., & Nussbaumer, H. 1974, *CoPhC*, **8**, 270
 Fletcher, L., & De Pontieu, B. 1999, *ApJL*, **520**, L135
 Foster, A. R., Smith, R. K., Brickhouse, N. S., Kallman, T. R., & Witthoef, M. C. 2010, *SSRv*, **157**, 135
 Freeland, S. L., & Handy, B. N. 1998, *SoPh*, **182**, 497
 Gilks, W. R., Richardson, S., & Spiegelhalter, D. J. 1996, *Markov Chain Monte Carlo in Practice*, **1**, 19
 Gupta, G. P., & Tayal, S. S. 1998, *ApJ*, **506**, 464
 Hastings, W. K. 1970, *Biometrika*, **57**, 97
 Jönsson, P., Gaigalas, G., Rynkun, P., et al. 2017, *Atoms*, **5**, 16
 Kallman, T. R., & Palmeri, P. 2007, *RvMP*, **79**, 79
 Lee, H., Kashyap, V. L., van Dyk, D. A., et al. 2011, *ApJ*, **731**, 126
 Loch, S., Pindzola, M., Ballance, C., et al. 2013, in *AIP Conf. Ser.* 1545 8th Int. Conf. on Atomic and Molecular Data and Their Applications, ed. J. D. Gillaspay, W. L. Wiese, & Y. A. Podpaly (Melville, NY: AIP), **242**
 Luridiana, V., & García-Rojas, J. 2012, in *IAU Symp.* 283, *Planetary Nebulae: An eye to the Future* (Cambridge: Cambridge Univ. Press), **139**
 Mariska, J. T. 1992, *The Solar Transition Region* (Cambridge: Cambridge Univ. Press)
 Martens, P. C. H., Kankelborg, C. C., & Berger, T. E. 2000, *ApJ*, **537**, 471
 Peres, G., Reale, F., & Golub, L. 1994, *ApJ*, **422**, 412
 Schrijver, C. J., & van Ballegoijen, A. A. 2005, *ApJ*, **630**, 552
 Storey, P. J., & Zeippen, C. J. 2010, *A&A*, **511**, A78
 Warren, H. P., Winebarger, A. R., & Brooks, D. H. 2012, *ApJ*, **759**, 141
 Warren, H. P., Winebarger, A. R., Mariska, J. T., Doschek, G. A., & Hara, H. 2008, *ApJ*, **677**, 1395
 Watanabe, T., Hara, H., Yamamoto, N., et al. 2009, *ApJ*, **692**, 1294
 Weiser, C. 2016, *mvQuad: Methods for Multivariate Quadrature* (R package version 1.0-6), <http://cran.R-project.org/package=mvQuad>
 Winebarger, A. R., Warren, H. P., & Falconer, D. A. 2008, *ApJ*, **676**, 672
 Xu, J., Van Dyk, D. A., Kashyap, V. L., et al. 2014, *ApJ*, **794**, 97
 Young, P. R. 2004, *A&A*, **417**, 785
 Young, P. R., Del Zanna, G., Landi, E., et al. 2003, *ApJS*, **144**, 135
 Young, P. R., Watanabe, T., Hara, H., & Mariska, J. T. 2009, *A&A*, **495**, 587