

Using Bayesian analysis and Gaussian processes to infer electron temperature and density profiles on the Mega-Ampere Spherical Tokamak experiment

G. T. von Nessi and M. J. Hole

Citation: [Review of Scientific Instruments](#) **84**, 063505 (2013); doi: 10.1063/1.4811378

View online: <http://dx.doi.org/10.1063/1.4811378>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/rsi/84/6?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Supersonic helium beam diagnostic for fluctuation measurements of electron temperature and density at the Tokamak TEXTOR](#)

Rev. Sci. Instrum. **83**, 065107 (2012); 10.1063/1.4707150

[Temperature and density characteristics of the Helicity Injected Torus-II spherical tokamak indicating closed flux sustainment using coaxial helicity injection](#)

Phys. Plasmas **15**, 082501 (2008); 10.1063/1.2968205

[Electron density measurements in the National Spherical Torus Experiment detached divertor region using Stark broadening of deuterium infrared Paschen emission lines](#)

Rev. Sci. Instrum. **77**, 10F127 (2006); 10.1063/1.2336456

[Electron density and temperature determination using the concept of particle confinement time uniqueness](#)

Rev. Sci. Instrum. **76**, 053508 (2005); 10.1063/1.1899403

[Evolution of the tokamak ISTTOK plasma density and electron temperature radial profiles determined by heavy ion beam probing](#)

Rev. Sci. Instrum. **70**, 947 (1999); 10.1063/1.1149448

Nor-Cal Products



Manufacturers of High Vacuum
Components Since 1962

- Chambers
- Motion Transfer
- Flanges & Fittings
- Viewports
- Foreline Traps
- Feedthroughs
- Valves



www.n-c.com
800-824-4166

Using Bayesian analysis and Gaussian processes to infer electron temperature and density profiles on the Mega-Ampere Spherical Tokamak experiment

G. T. von Nessi^{a)} and M. J. Hole

Research School of Physical Sciences and Engineering, The Australian National University, Canberra ACT 0200, Australia

(Received 13 February 2013; accepted 3 June 2013; published online 25 June 2013)

A unified, Bayesian inference of midplane electron temperature and density profiles using both Thomson scattering (TS) and interferometric data is presented. Beyond the Bayesian nature of the analysis, novel features of the inference are the use of a Gaussian process prior to infer a mollification length-scale of inferred profiles and the use of Gauss-Laguerre quadratures to directly calculate the depolarisation term associated with the TS forward model. Results are presented from an application of the method to data from the high resolution TS system on the Mega-Ampere Spherical Tokamak, along with a comparison to profiles coming from the standard analysis carried out on that system. © 2013 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4811378>]

I. INTRODUCTION

The incoherent Thomson scattering (TS) of laser light off of electrons is the foundation for standard diagnostics used to measure electron temperature and density in fusion devices.¹ Recently, the Mega-Ampere Spherical Tokamak (MAST) TS system was upgraded to measuring 130 points across the midplane at a frequency of up to 240 Hz.²

In this work, Bayesian methods are used to infer both T_e and n_e profiles using both TS and interferometric data coming from diagnostics on MAST. The primary advantage of this analysis is that an absolute calibration of the time-integrated TS laser pulse energy is not required to infer either the thermal electron temperature nor density profiles.

A stochastic model, based on Gaussian processes (GPs),³ is used to model the correlation between TS observation points. The noise, signal-variance, and length-scale associated with this model are inferred as nuisance parameters in addition to the T_e and n_e profiles. This provides a strong benefit of the presented approach, in that signal noise is strongly decoupled from the inferred profiles; and thus, helps to mitigate the probability of previous errors,⁴⁻⁶ where signal noise was given a physical interpretation, from reoccurring. Moreover, the standard perturbative methods for approximating the depolarisation term in the TS forward model, (cf. Naito *et al.*⁷) are bypassed in this analysis, in favour of a more direct and flexible numerical integration calculation using Gauss-Laguerre quadratures. Caching of depolarisation calculations is employed to negate the higher computational cost of this approach.

The paper is structured as follows. Section II gives a brief overview of Bayesian inference using diagnostic data and the application of GPs to model the correlation between observation points. Next, the TS and interferometric forward models are presented with a description of the depolarisation computation. Section IV present profiles inferred for a

high-performance MAST discharge along with a comparison to profiles calculated from the MAST standard analysis. Finally, conclusions and possible extensions are discussed.

II. OVERVIEW OF BAYESIAN INFERENCE

The objective of any inference scheme is to statistically infer a vector of model parameters, denoted $\bar{\lambda}$, given a vector of diagnostic data and associated uncertainties, $\bar{\mu}$ and $\bar{\sigma}$, respectively. In the Bayesian perspective, inference centres around Bayes' formula:

$$\mathcal{P}(\bar{\lambda}|\bar{\mu}, \bar{\sigma}, \bar{I}) = \frac{(\prod_i \mathcal{P}(\mu_i|\bar{\lambda}, \sigma_i, \bar{I})) \mathcal{P}(\bar{\lambda})}{\mathcal{P}(\bar{\mu}, \bar{\sigma}, \bar{I})}, \quad (1)$$

where \bar{I} denotes background assumptions. To keep notation uncluttered, \bar{I} is dropped for the rest of the paper, with background assumptions being explicitly indicated where appropriate. The parts of Eq. (1) and their application/interpretation in the context of diagnostic data are well documented in the literature⁸⁻¹¹ and will not be discussed in detail here.

TS and interferometer uncertainties are assumed to be pairwise uncorrelated, with each observation having an associated likelihood of the assumed form

$$\mathcal{P}(\mu_i|\bar{\lambda}, \sigma) = \mathcal{N}(\mu - \mathcal{F}(\bar{\lambda}), \sigma^2), \quad (2)$$

where $\mathcal{N}(\mu, \sigma^2)$ is a Gaussian distribution of mean μ and σ^2 the variance; $\mathcal{F}(\bar{\lambda})$ is the forward model associated with the given diagnostic. Justifications for this form of the likelihood are given elsewhere.^{8,10}

A GP is employed to define the covariance matrix reflecting a spatial correlation between TS observation points

$$K_{ijk} := \zeta_k^2 \exp\left(-\frac{(R_i - R_j)^2}{2\tau_k^2}\right) + \eta_k^2 \delta_{ij}, \quad (3)$$

where R_i and R_j represent the radial coordinate of the i th and j th TS observation point along the midplane; and the k index indicates correspondence to electron temperature or density measurements, with a value of 0 or 1, respectively. In

^{a)}greg.vonnessi@anu.edu.au

Eq. (3) ζ_k , τ_k , and η_k are commonly referred to as hyper-parameters:³ non-physical quantities which help characterise the prior and/or likelihood in the overall inference. In this analysis, these are scalar quantities uniquely associated with each profile: τ_k reflects the average radial length-scale over which the profile is changing; ζ_k is the signal variance, which serves to decouple the average profile gradient from the length-scale; and η_k is the average scalar noise on the profile. With the expression in Eq. (3), the prior for the inference is proportional to

$$\begin{aligned} \mathcal{P}(\vec{\lambda}) &= \mathcal{P}(\vec{T}_e, \vec{n}_e, \tau_0, \tau_1, \zeta_0, \zeta_1, \eta_0, \eta_1) \\ &\propto \mathcal{N}(\vec{T}_e^T K_0^{-1} \vec{T}_e) \mathcal{N}(\vec{n}_e^T K_1^{-1} \vec{n}_e) \prod_{k=0}^1 \\ &\quad \times (\mathbf{1}_{[0,10]}(\vec{\tau}_k) \mathbf{1}_{[0,10]}(\vec{\zeta}_k) \mathbf{1}_{[0,10]}(\vec{\eta}_k)), \end{aligned} \quad (4)$$

where \mathcal{N} are zero-mean Gaussian distribution with their covariance matrix shown in the argument, and $\mathbf{1}_{[a,b]}(x)$ indicates a uniform distribution of the variable x over the closed interval $[a, b]$. The upper bounds for the τ_k , ζ_k , and η_k uniform distributions were empirically selected so as to not preclude any physically attainable profiles on MAST, as determined by analysing 36 different shot/time slices.

Equation (4) shows the prior as a product of uniform and zero-mean Gaussian distributions. As the covariance matrices for the Gaussians over T_e and n_e are themselves characterised by *inferred* hyper-parameters, the prior in Eq. (4) embodies a stochastic model of the spatial correlation between different TS observation points. For a fixed set of hyper-parameters, one can think of this prior as favouring T_e and n_e profiles convolved with Gaussian kernels (i.e., mollifications¹²) of fixed widths corresponding to τ_k in the inference. As the hyper-parameters themselves are inferred with minimal constraint, the analysis is also able to infer the average length-scale, signal variance, and noise variance for both the T_e and n_e profiles, as intrinsically held by the data. Further details on these points can be found in the study by Rasmussen and Williams.³

III. FORWARD MODEL

In the MAST TS system, Thomson scattered light from each observation point is spectrally divided into four bands via a polychrometer. Each filtered band is then focused onto an avalanche photodiode (APD), translating the integrated intensity over the spectral band into a voltage signal. When integrated over the TS laser pulse length, a quantity is produced which is sensitive to both the thermal electron temperature and density at the associated observation point

$$V_{TS} = C_S n_e E_L \int \frac{\phi(\lambda)}{\phi(\lambda_L)} \frac{S(\lambda_s, \lambda_L, \theta, T_e)}{\lambda_L} d\lambda, \quad (5)$$

where C_S encompasses a collection of known, fixed system constants; E_L is the integrated laser energy; $\phi(\lambda)$ is the polychrometer response functions; λ_L is the TS laser wavelength; λ_s is the wavelength of the scattered photons; θ is the scattering angle; and $S(\lambda_s, \lambda_L, \theta, T_e)$ the standard Selden expression.^{2,13} The Selden relation relates the intensity of scattered light off of thermal electrons to T_e and is standard

in the literature; but the relation is recalled here for completeness (note the variable transformations in Eqs. (9) and (10):

$$S(\epsilon, \theta, 2\alpha) = S_Z(\epsilon, \theta, 2\alpha) q(\epsilon, \theta, 2\alpha), \quad (6)$$

$$S_Z(\epsilon, \theta, 2\alpha) = \frac{\exp(-2\alpha x)}{2K_2(2\alpha)(1+\epsilon)^3} [2(1-\cos\theta)(1+\epsilon) + \epsilon^2]^{-1/2}, \quad (7)$$

$$\begin{aligned} q(\epsilon, \theta, 2\alpha) &= 1 + \frac{2x}{y} \exp(2\alpha x) \left(y^2 \int_x^\infty \frac{\exp(-2\alpha\xi)}{(\xi^2 + u^2)^{3/2}} d\xi \right. \\ &\quad \left. - 3 \int_x^\infty \frac{\exp(-2\alpha\xi)}{(\xi^2 + u^2)^{5/2}} d\xi \right), \end{aligned} \quad (8)$$

$$\epsilon := \frac{\lambda_s - \lambda_i}{\lambda_i}, \quad 2\alpha := \frac{m_e c^2}{T_e}, \quad u = \frac{\sin\theta}{1 - \cos\theta}, \quad (9)$$

$$x := \left(1 + \frac{\epsilon^2}{2(1 - \cos\theta)(1 + \epsilon)} \right)^{1/2}, \quad y := \frac{1}{(x^2 + u^2)^{1/2}}, \quad (10)$$

where $q(\epsilon, \theta, 2\alpha)$ represents the relativistic depolarisation correction term.^{2,7,13}

For TS systems, the expressions presented by Naito *et al.*⁷ are normally used to approximate $q(\epsilon, \theta, 2\alpha)$. In this work, however, Gauss-Laguerre quadratures are utilised to provide a more direct and flexible calculation of the depolarisation correction. The fundamental difference in this approach is that it is a non-perturbative calculation, which contrasts the approximations of Naito *et al.*⁷ that are based on Taylor expansions. Indeed, changing variables according to $v = 2\alpha(\xi - x)$ in Eq. (8) gives

$$\begin{aligned} q(\epsilon, \theta, 2\alpha) &= 1 + \frac{4\alpha x}{y} \left(y^2 \int_0^\infty \frac{\exp(-v)}{\left(\left(\frac{v}{2\alpha} + x\right)^2 + u^2\right)^{3/2}} dv \right. \\ &\quad \left. - 3 \int_0^\infty \frac{\exp(-v)}{\left(\left(\frac{v}{2\alpha} + x\right)^2 + u^2\right)^{5/2}} dv \right), \end{aligned} \quad (11)$$

which can be integrated directly using a Gauss-Laguerre quadrature.¹⁴ The initial computational cost of the quadrature construction is offset by the fact that the integrals in Eq. (11) are computed via pre-calculated quadrature poles and weights. Gauss-Laguerre quadratures are very accurate/efficient for calculating integrals of the form seen in Eq. (11) and allow for arbitrary levels of accuracy to be specified by simply changing the number of quadrature points.¹⁴

The primary issue with a forward model based on Eq. (5) is that it requires an absolute calibration of E_L , in addition to the calibrations reflected in the value of C_S . Indeed, observations modelled by Eq. (5), taken across a midplane chord can only serve to constrain the n_e profile up to a scaling constant, if E_L is otherwise unknown. Fortunately, MAST has a midplane CO₂ interferometer that provides a line-integrated

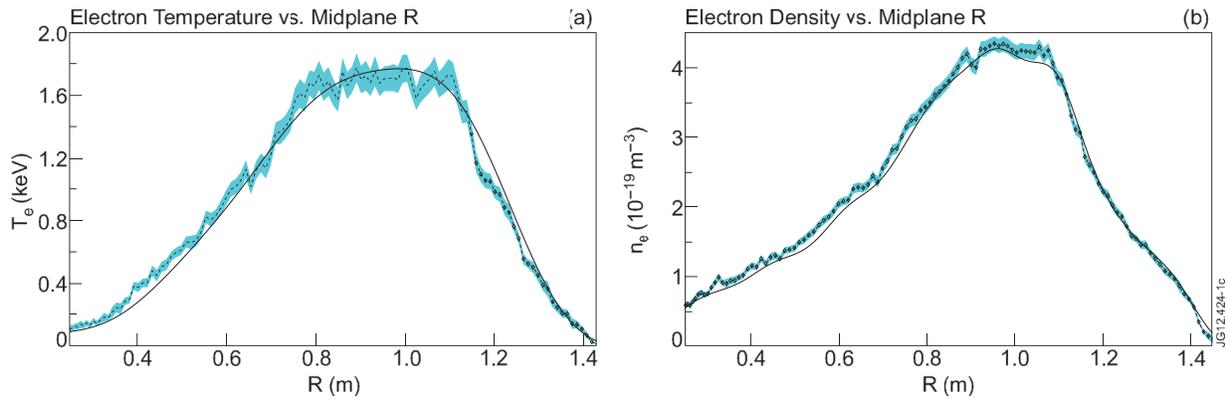


FIG. 1. Comparison of inferred T_e and n_e profiles inferred from TS data for shot #24 600 at 280 ms, with T_e and n_e profiles corresponding to (a) and (b), respectively. The heavy lines indicate profiles inferred using the Bayesian analysis described in Sec. II. The dashed lines marked with diamonds and coloured region reflect MAST scheduler output for the corresponding profiles and uncertainty, respectively. Uncertainties associated with the expectation of the Bayesian inferred profiles are too small to resolve visually and have been suppressed from both figures.

measurement of n_e in the midplane. Using this data, a simple coordinate transformation is employed to construct the integrated electron density along the TS laser's line of sight (also in the midplane). Thus, there is an effective observation/forward model available for the line integrated electron density

$$O_{IF} = \int_{TS} n_e d\ell, \quad (12)$$

where the integral is along the TS laser's line-of-sight; and O_{IF} indicates the observations coming from the interferometer. It is clear from Eq. (12), that the interferometer observation can constrain the scaling of the n_e profile, if the shape of the profile is known. Combining both the TS and interferometer observations, thus can effectively constrain all aspects of the n_e profile. Moreover, as E_L is the same for all TS observation points, it is inferred as a nuisance parameter (i.e., an unphysical parameter which is integrated out in the inference of physical model parameters) that can also absorb any error in the value used for C_S , with n_e still being well-constrained by both TS and the line-integrated interferometric observation. This serves to greatly reduce the errors seen on the inferred profiles, even when E_L is given a physically unconstraining uniform prior, as in this analysis (see Sec. IV).

IV. RESULTS

To demonstrate the analysis detailed above, an inference and subsequent comparison is made against profiles coming from the standard TS analysis carried out on MAST. This comparison is made on discharge 24 1600 at 280 ms, which is a L-mode discharge in a DnD configuration with 3.35 MW of co-injected NBI heating. The Bayesian inference was carried out using a 1000 point Gauss-Laguerre quadrature to calculate $q(\epsilon, \theta, 2\alpha)$, with posterior moments taken from sampling statistics obtained via a specialised implementation of Skilling's nested sampling (NS) algorithm detailed elsewhere¹⁵ (see Sivia and Skilling¹⁰ for details on NS). The sampling results were also independently validated by comparison with samples generated from a Hamiltonian Markov

Chain Monte Carlo (HMCMC) algorithm (see MacKay¹⁶ for details on HMCMC).

Figure 1 shows that the inferred, expectation of the profiles show generally, very good agreement with the profiles coming from standard analysis. The most striking difference between both sets of profiles is the difference in uncertainties. Indeed, the nature of the unified Bayesian inference indicates that the profiles are very well constrained. This can be understood by noting that T_e and n_e are being inferred from four APD measurements (nominally) at every observation point, which are in addition to the global constraints provided by interferometry and the prior, rendering the inference as being strongly over-determined. The presence of global constraints and over-constraining local observations make the inference very robust against APD errors, mis-calibrations or signal loss. In contrast, the MAST analysis uncertainty reflects a maximum entropy result (i.e., with minimal prior assumptions relative to the forward model used) and is expected to have higher uncertainties than results utilising the prior in Eq. (4) (or any analogy thereof).

As mentioned in Sec. II, this inference yields mollified profiles, along with the associated hyper-parameters represented in Eq. (3). Thus, the results produced have a different interpretation (i.e., they need to be understood in relation to their respective hyper-parameters) than the profiles from the standard MAST analysis. Indeed, Fig. 1 can only be viewed as a sanity check for the global shape/scale of the T_e/n_e profiles, as the fine-scale structure information coming from the Bayesian analysis is largely contained in the inferred hyper-parameters. In this context, it is not surprising to see the inferred profiles lying outside the errors bars associated with the MAST analysis. One of the main advantages for such a comparison, is that one can qualitatively see how much of the profile structure can be attributed to a scalar noise term. Quantitatively, one can compare the scalar noise term, η_k , of the profile to the random errors modelled for the diagnostic to validate uncertainties produced by explicitly propagating error through the forward model itself. Finally, one may adjust the prior on any of the hyper-parameters in the inference to reflect *a priori* knowledge of the length-scales or random noise seen in the system; this can even be taken to the point

TABLE I. Table of sampling statistics of nuisance parameters used in the inference of T_e and n_e profiles from TS data for discharge #24 600 at 280 ms.

Nuisance parameter (Unit)	Sampling expectation w/ 95% confidence interval
E_L (J)	$(1.345 \times 10^{-1})^{+1.703 \times 10^{-3}}_{-1.561 \times 10^{-3}}$
τ_0 (m)	$(2.296 \times 10^{-1})^{+6.655 \times 10^{-3}}_{-8.101 \times 10^{-3}}$
ζ_0	$1.006^{+2.822 \times 10^{-3}}_{-5.427 \times 10^{-3}}$
η_0	$(6.698 \times 10^{-3})^{+3.209 \times 10^{-3}}_{-6.073 \times 10^{-3}}$
τ_1 (m)	$(9.974 \times 10^{-2})^{+1.761 \times 10^{-3}}_{-1.081 \times 10^{-2}}$
ζ_1	$1.511^{+3.039 \times 10^{-3}}_{-2.839 \times 10^{-3}}$
η_1	$(6.650 \times 10^{-3})^{+5.353 \times 10^{-3}}_{-5.580 \times 10^{-3}}$

where these hyper-parameters are even fixed at particular values. Of course, such priors will effect the structure of the inferred profiles; e.g., fixing a relatively small length-scale and scalar noise will yield profiles which have a structure closer to those produced by the MAST analysis. Exploration of physics using these more informed priors on the hyper-parameters is a current research endeavour.

The small uncertainties in Fig. 1 are mostly due to the inference being over constrained, as four APD signals (with relatively small uncertainties) are present at each TS observation point with only T_e and n_e needing to be inferred at said point. Even with the addition of the E_L nuisance parameter and the six hyper-parameters associated with the GP stochastic model, the ratios of inference parameters to observations is still around 1:2 for most MAST discharges. While the prior associated with the GP stochastic model does inform the posterior, the hyper-parameters being model parameters themselves with uniform priors make the GP stochastic model non-selective (i.e., effectively having a uniform probability measure) over the entire space of profiles. Thus, the GP stochastic model and prior by themselves are not able to strongly influence the size of the uncertainties seen in Fig. 1. In theory, one may use a completely uninformative Jeffreys prior on the GP stochastic model hyper-parameters; but utilisation of this prior ultimately requires that posterior expectations be calculated at each sampling step, which effectively render the inference as being computationally intractable, using currently available tools.

Table I shows sampling expectations and uncertainties associated with the nuisance parameters presented in Eqs. (3) and (5). While the values in Table I are treated as nuisance parameters, they have been included to give the reader some context for the length scales, signal, and noise variance that are typically encountered with MAST TS data. Again, the uncertainties on these values are small relative to the inferred values, as the inference is over-constrained to the point where uncertainties due to degeneracies in highly likely model parameter configurations are all but eliminated.

Finally, as the inference outlined in this paper is non-analytic in nature, having approximately 270 model parameters for a given discharge, it is slower than the standard analysis. Indeed, to calculate statistical moments of the profiles using NS takes approximately 30 minutes per inference

(on average), on a 2.2 GHz processor with 8 GB of memory. It is a current research focus to speed up this inference by developing a parallelised version of sampling algorithm.

V. CONCLUSIONS

A new method for the unified Bayesian inference of thermal electron temperature and density profiles has been demonstrated, which is also able to infer length-scales and scalar noise parameters intrinsically contained within the diagnostic data. By employing GPs, the average length-scale, signal variance, and noise variance are inferred as nuisance parameters, which directly yield profiles where noise is minimised. Finally, a new approach to calculating the depolarisation correction in the TS forward model is presented, which is simple to implement and affords easy adjustment of calculation accuracy.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Rory Scannell at CCFE for his helpful discussions regarding the MAST TS system and Dr. David Muir, also at CCFE, for providing aid with MAST data access.

This work was jointly funded by the Australian Government through International Science Linkages Grant No. CG130047, the Australian National University, the RCUK Energy Programme under Grant No. EP/I501045, and the European Communities under the contract of Association between EURATOM and CCFE. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

- ¹I. H. Hutchinson, *Principles of Plasma Diagnostics*, 2nd ed. (Cambridge University Press, 2005).
- ²R. Scannell, M. J. Walsh, P. G. Carolan, A. C. Darke, M. R. Dunstan, R. B. Huxford, G. J. McArdle, D. Morgan, G. Naylor, T. O’Gorman, S. Shibaev, N. Barratt, K. J. Gibson, G. J. Tallents, and H. R. Wilson, “Design of a New Nd:YAG Thomson scattering system for MAST,” *Rev. Sci. Instrum.* **79**(10), 10E730 (2008).
- ³C. E. Rasmussen and C. K. L. Williams, *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning Series (The MIT Press, 2005).
- ⁴N. J. L. Cardozo, F. C. Schüller, C. J. Barth, C. C. Chu, F. J. Pijper, J. Lok, and A. A. M. Oomens, “Plasma filamentation in the Rijnhuizen tokamak RTP,” *Phys. Rev. Lett.* **73**(2), 256 (1994).
- ⁵B. Ph. van Milligen, I. G. J. Classen, and C. J. Barth, “Revision of TV Thomson scattering data analysis and detection of profile structure,” *Rev. Sci. Instrum.* **74**(9), 3998 (2003).
- ⁶N. J. L. Cardozo, I. G. J. Classen, C. J. Barth, C. C. Chu, A. J. H. Donné, R. J. E. Jaspers, J. Lok, A. A. M. Oomens, F. J. Pijper, and F. C. Schüller, “Erratum: Plasma Filamentation in the Rijnhuizen Tokamak RTP [Phys. Rev. Lett. **73**, 256 (1994)],” *Phys. Rev. Lett.* **90**(14), 149901 (2003) (erratum).
- ⁷O. Naito, H. Yoshida, and T. Matoba, “Analytic formula for fully relativistic Thomson scattering spectrum,” *Phys. Fluids B-Plasma Phys.* **5**(11), 3 (1993).
- ⁸G. T. von Nessi, M. J. Hole, J. Svensson, and L. C. Appel, “Evidence cross-validation and Bayesian inference of MAST plasma equilibria,” *Phys. Plasmas* **19**(1), 012506 (2012).
- ⁹J. Svensson, A. Werner, and JET-EFDA Contributors, “Current tomography for axisymmetric plasmas,” *Plasma Phys. Controlled Fusion* **50**(8), 085002 (2008).
- ¹⁰D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial*, 2nd ed. (Oxford University Press, New York, 2006).

- ¹¹E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).
- ¹²L. C. Evans, *Partial Differential Equations*, 2nd ed., Graduate Studies in Mathematics Vol. 19 (American Mathematical Society, 2010).
- ¹³R. Scannell, "Investigation of H-mode edge profile behaviour on MAST using Thomson scattering," Ph.D. dissertation, Oxford, 2007.
- ¹⁴F. B. Hildebrand, *Introduction to Numerical Analysis*, 2nd ed., Dover Books on Advanced Mathematics (Dover Publications, 1987).
- ¹⁵G. T. von Nessi, M. J. Hole, and MAST team, "A unified method for inference of tokamak equilibria and validation of force-balance models based on Bayesian analysis," *J. Phys. A: Math. Theor* **46**(18), 185501 (2013).
- ¹⁶D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*, 1st ed. (Cambridge University Press, 2003).