Referee Report on the Doctoral Thesis

"Forward and Backward Modelling of Spectroscopic Diagnostics in Fusion Plasmas"
By Ing. Matěj Tomeš

Referee:
dr. Maarten De Bock,
Visible Spectrosopy Diagnostic Responsible Officer
ITER Organization*
Route de Vinon-sur-Verdon - CS 90 046 - 13067 St Paul Lez Durance Cedex – France
maarten.debock@iter.org

* The review of this doctoral thesis falls outside the scope of the professional activities of the referee for the ITER organization. The views and opinions expressed herein do not necessarily reflect those of the ITER organization.

Table of Contents

1 PURPOSE 3
2 FORWARD & BACKWARD MODELLING OF THOMSON SCATTERING 3
  2.1 GENERAL ASSESSMENT 3
  2.2 DETAILED/MINOR COMMENTS 4
  2.3 QUESTIONS 4
3 NEON CONCENTRATION INFERENCE IN THE JET DIVERTOR 5
  3.1 GENERAL ASSESSMENT 5
  3.2 DETAILED/MINOR COMMENTS 6
  3.3 QUESTIONS 7
4 TOWARDS FORWARD MODELLING OF SYNCHROTRON RADIATION 7
  4.1 GENERAL ASSESSMENT 7
  4.2 DETAILED/MINOR COMMENTS 8
  4.3 QUESTIONS 8
5 OVERALL CONCLUSION 9

1 Purpose

The purpose of this referee report is to provide an assessment of the doctoral thesis "Forward and Backward Modelling of Spectroscopic Diagnostics in Fusion Plasmas" by Ing. Matěj Tomeš.

This report structured according to the structure of the doctoral thesis itself: an assessment of the Thomson scattering modelling, of the divertor neon concentration modelling and of the synchrotron radiation of runaway electrons. Each section is split up into an overall assessment (including novelty, scientific relevance and potential future applications), detailed comments on specific pages, figures or equations and a number of questions. A final section gives a short overall assessment of the thesis as a whole.

2 Forward & Backward Modelling of Thomson Scattering

2.1 General assessment

The first content chapter of the doctoral dissertation describes both a novel backward model and forward model for the Thomson scattering technique of measuring the electron temperature and density of a fusion plasma. A short description of the conventional backward model is also provided. All descriptions are clear and to the point.

For the backward models, the concerns with the standard backward model are clearly highlighted and the advantages of a Bayesian probabilistic model are as well. The Bayesian probabilistic backward model uses existing Bayesian methodology and an existing python library (both of which are clearly and correctly referenced). A Bayesian analysis has already previously been implemented for the case of Thomson Scattering; e.g. on JET (Sehyun Kwak et al 2020 Nucl. Fusion 60 046009) and W7-X (S.A. Bozhenkov et al 2017 JINST 12 P10004) using the Minerva framework. This earlier work, however, used a simplified approach, relaying on purely Gaussian processes; whereas the work presented in the doctoral dissertation provides a more generalized approach. In particular, the idea to use both time dependent signal and uncertainty functions based on generalized Gaussian functions in order to capture the unexpected tail in the experimental Thomson Scattering signal struck me as creative.

The Minerva framework is also fairly closed and the work presented here is the first, as far as my knowledge goes, in an open source Python environment that would be easier to integrate with the forward model discussed later in the chapter, which is also open source and developed using the (Python and open source) Cherab/Raysect framework.

Nonetheless, it might have been good to refer the earlier work on Bayesian inference of Thomson Scattering measurements.

The forward model discussed in the latter part of the chapter is, to the best of my knowledge, the first one that uses both raytracing and a modular, device independent approach. The former not only allows to include a detailed model of the optical components of the collection optics, it also allows to take into account the effect of reflections. The latter allows to transfer the modelling with relative ease to other fusion devices and allows to add improved (for better accuracy) or simplified (for calculation speed) version of the physics description without impacting the rest of the code.

The modularity is intrinsic to the Cherab/Raysect framework used. The work presented in this doctoral dissertation include the creation of not just a new Thomson Scattering module for the Cherab/Raysect framework, but also the introduction of several laser related classes that will allow other contributors to built additional modules for laser-based diagnostics – Thomson scattering or otherwise. This is of significant benefit to the larger fusion community. The choice for contributing to the Cherab/Raysect framework has the additional benefit that it allows to include detailed models for background radiation developed by third parties (for whom the ‘background’ of Thomson scattering might have been the prime physical process to investigate). This way a common flaw of proprietary codes – to put all focus on the main process of interest and oversimplify secondary/background processes – can be avoided.

The inclusion of a detailed – and verified – description of the complete optical design of the collection optics, instead of just including a simple sightline geometry as is done in other Thomson Scattering models, has significant benefits: it allows to verify to what extend the (less computational intensive) sightline approximation is valid and how the real life optical design of a Thomson Scattering diagnostic can be optimized in order to approximate a simplified sightline model as close as possible. This becomes more and more relevant as fusion experiments evolve into larger fusion reactors with complex optical labyrinths needed to shield neutron radiation.

In conclusion, the work in this chapter is of high scientific quality and relevance, contains several original contributions and opens possibilities for further applications.

2.2 Detailed/minor comments

Section 2.3.1: I found the explanation of standard backwards model a bit too concise or simplified. E.g. it states the derivation of the electron temperature T_e needs no calibration, but at least the filter shapes as function of wavelength need to be calibrated/measured. This becomes clear later on in the chapter, so I understand that what was meant is that T_e does not need an "absolute" calibration in contrast to the electron density n_e. It might have been more careful to state no “absolute” calibration is needed for T_e.

Equation 2.8: The use of both $B_{p,ch}$ in $W_{p,ch}(t)$ as constant term and B as exponent in the exponential term of $U_{p,ch}(t)$ is a bit confusing. I'm quite sure they are independent, but the second B is not explained explicitly, so I cannot be 100% sure. It is also unclear if the scale parameter A in $U_{p,ch}(t)$ is the same as the $A = 0.2\ T_{duplex}$ used to define ΔU in $U_{p,ch}(t)$ and the α scale parameters in $W_{p,ch}(t)$.

Figure 2.15: The parameters $B_{p,ch}(t)$ and B (exponent in exponential term of $U_{p,ch}(t)$) seems missing as either probabilistic or deterministic parameter inputs.

Figure 2.21 and text related to it on page 39: The statement that "significant improvement was achieved in the edge region $Z>0.28m$" cannot be visually concluded from figure 2.21. The scale of the vertical axis is just too large to see the statistical predictions clearly.

Page 43: The last paragraph on the density is, understandably, very short. The confidence with which the hypothesis is presented that the failure of a statistical treatment of the density profiles is due to the pulse shape deformation, stands in contrast with the very limited evidence that is presented. A bit more caution and an indication of what could be done to investigate the hypothesis would have been a better conclusion of the section than the fairly blunt "This also supports the hypothesis".

Page 48 to 50: the start of section 2.5, end of 2.4.2 and caption of figure 2.25 are a bit of a duplication, all explaining how figure 2.25 was generated.

Page 59: I realize I am a bit nitpicking here, but the collaboration with ITER has identifier "LGA-2021-A-76". The given "LGA-2015-M08" is the overarching Memorandum of Understanding for any collaboration between IPP Prague and ITER.

2.3 Questions

1. The section on the backward model seem to indicate that the non-gaussian tail of the pulses is reasonably systematic. Hence why did you choose to model the pulse shape by a (generalized) Gaussian with a larger uncertainty for the tail, rather than to model the pulse shape as a (generalized) gaussian with an extra tail term?

2. If I understand it correctly, the standard backward analysis was used to analyse the forward model for the benchmarking. However, the forward model has no time dependence of the laser pulses so the pulse shape fitting was skipped. I assume this is also why the newly developed Bayesian statistical backwards model was not used? Would it be possible to include the pulse shape (e.g. as post-processing after the raytracing) in the forward model? And what would be the advantages or disadvantages of doing so?

3. In current tokamaks the simplified forward models for TS provide very adequate results. Hence the novel model in Cherab/Raysect you developed as part of your PhD work is certainly a nice-to-have, but not strictly a need-to-have. Could you give some reasons why your Cherab/Raysect model becomes much more of a need-to-have for more reactor oriented machines like ITER, SPARC, DEMO ...?

4. How would you go about introducing the depolarization effects in the forward model for machines like ITER that have core temperatures well above 10keV? Not only in the model, but also in propagating the (de)polarization using the ray-tracing?

3 Neon concentration inference in the JET divertor

3.1 General assessment

In this chapter a forward model of divertor KT3 spectroscopic diagnostic in CHERAB (using ray tracing) is presented, as well as a backward model to extract Neon concentrations and methodology for validating the backward model.

The chapter starts with a fairly extensive overview of the multiple radiation processes that can take place in a tokamak divertor. This is of course not new work, but nevertheless justified to be included in the dissertation because if the complexity and its impact on both forward and backward modelling.

The subsequent description of the forward model for the KT3 KT3 spectroscopic diagnostic is clear and concise. The forward model itself is, to my knowledge, first of a kind for the JET divertor spectroscopy. It is, however, less innovative or detailed than the Thomson Scattering model developed for chapter one. Nonetheless, the simplifications made to the forward model due to limited input data and constraints on computational load are well explained and adequate. Moreover, this forward model was a conditio sine qua non for the backward model presented later in the chapter.

The discussion of the large discrepancies between forward model simulated Neon emission lines (based on SOLPS) and actual measured Neon emission lines is a bit short in my opinion, especially because the subject of the chapter is on the Neon concentration measurements (and not on the much better matching Balmer series). This could of course be an issue with the SOLPS input data and not so much with the forward model. Impurity emission prediction is notoriously difficult due to the many unknowns and hence assumptions that have to be made. Nevertheless, a longer, more detailed discussion might have be useful.

The next section introduces the validation methodology developed for the backward model. The method assigns weights to each SOLPS cell, that defines how much that cell is contributing to the signal of each channel sightline crossing that cell. The weights are determined by both the fraction of the cell covered by the channel and the local radiative power coming from the cell. Because plasma parameters (like electron density, neon concentration, …) are defined per cell, one can assign the weight of a cell for a given channel to the parameter value of that cell. And one can assess how much a certain parameter “contributes” to the signal of a channel.

The methodology is creative and provides useful insights. A minor point of criticism is that the description of the method sometimes lacked some clarity. Although the initial description can be followed quite easily, the change from weights of a cell for a given channel to weight of a plasma parameter for a given channel was not immediately obvious to me. Also the implementation for ratios (e.g. the neon line ration of figure 3.19) is not obvious: I assume in such case the sum of the radiative power for both lines of which the ratio is investigated is taken as the power term in the weight, but that is not mentioned in the text.

The zero transport backward model developed is based on the one developed for the ASDEX Upgrade tokamak (which is correctly referenced). The novel contribution comes from the application of the forward model using ray tracing introduced at the start of the chapter (and the fact that JET is a different machine than ASDEX-U is). The description and justification of the assumptions used are clear. What is a bit less clear is the how the backward statistical model was built up, which parameters were considered deterministic and which ones probabilistic. A flow diagram as used in figure 2.15 for the Thomson scattering model might have helped.

The backward model was subsequently applied to forward modelled spectra based on known SOLPS simulated data and verified using the validation methodology described earlier in the chapter. The discrepancies are well documented, and the discussion points out likely causes in the the necessary simplifications applied in the backward model. Subsequently the backward model is applied to experimentally measured spectra. Given the discrepancies already identified when verifying against known SOLPS input data, as well as the large discrepancies in neon line intensities between the forward model (based on SOLPS input data that does show reasonably agreement for the Balmer series) and experimental spectra, it would have been good to provide a clear warning about the expected validity of the inferred neon concentrations from experimental spectra. In the discussion at the end of the chapter, the fact quantitative data should be treated with great care is mentioned, but it might have been good to also highlight that also at the start of the section.

3.2 Detailed/minor comments

Page 61: In the 2nd paragraph of the introduction there's an empty reference [.] for the process of radiative cooling.

Figure 3.1: There’s a typo in the caption. I guess the n_e for the dashed CX trace should be 10^{20} m^{-3}

Page 66 and figure 3.3: Only PLT and PRB are visible on the figure (and PLT is clearly dominating). I assume that PLS and PRC, which are also mentioned in the text, are negligible, but as they are not in the figure that cannot be derived from it. A mention that PLS and PRC are negligible and were not included in the figure for that reason would have been more rigorous.

Page 80-81 and figure 3.15: With $W_{i,ch}$ defined in (3.9) as the weight of cell i to channel ch, it is not immediately clear how to interpret figure 3.15. The text on page 80-81 defines the weight distribution per channel as function of the density n_e, but does not explain how one gets there starting from the weight of cells (which is the only things discussed in before page 80-81). I assume the spatial density distribution of figure 3.14 is used to map cells covered by a channel to density values in those cells, hence assigning the weight of the cell to the density value of that cell. And as such figure 15 defines which density has a higher weight in the measured intensity of the channel. But this is not immediately clear from the text.

Page 84 and figure 3.16: It is not clear why all channels except 10 to 16 exhibit a bimodal distribution. For channel 17 and upwards some bimodality can be seen, but no below channel 10.

Figure 3.27: The caption states the blue distributions are “inferred” and the orange are the “actual” distributions. But the legend in the figure itself states “inferred” for the orange data points and just states “distribution” for the blue data points (while both are distributions). This could have been more clear. Also, the orange distributions in figure 3.27 corresponds to the blue distributions in figure 3.20; keeping consistent colours between the figures might have helped clarity.

Page 94: The statement in the one but last sentence that "For shot number 97490 ..., there was an observed increase in intensity near R=2.8m of approximately 50%" is (a) not clearly indicating with respect to what there is this 50% increase and (b) is not obvious from figures 3.29 and 3.30 (other shots show very similar, albeit scaled down, radiance profiles).

Page 99-100: There might be a typo in last sentence of page 99, first of page 100: I think HFS should be LFS. I think this because Figure 3.37 shows a discontinuity at R=2.8m which corresponds to the LFS, not to the HFS. I cannot really distinguish a clear discontinuity near the HFS of tile 5 (around 2.56m).

Figures 3.39 and 3.40: The different vertical scales make it hard to verify that the electron density indeed increases and electron temperature decreases for increasing Neon puff. This is a general comment for many figures within the thesis: trying to keep vertical scales the same on figures that are to be compared would help (or if the ranges covered differ too much from figure to figure, then at least scale them by an integer factor).

3.3 Questions

1. The CHERAB model currently does not allow to combine Doppler (temperature) and Stark (density dominated) broadening effects of spectral emission lines. Is that a fundamental limitation of the architecture of CHERAB? If not, how would you address the creation of a module that can combine both types of line broadening effects?

2. The transistion from cell contribrution weights to weighted quantity distributions was handled in a very concise manner in the thesis. Could you elaborate in more detail how one can get from weight of a 2D grid of cells *i* to channel a *ch* to the weight distribution of a 1D quantity like e.g. the electon density? And what radiative power would you consider in case of deriving the weight distribution for a line ratio?

3. Based on the validation with SOLPS simulations it is said that the zero transport backward model could underestimate the neon concentration by approximately 3 to 4 times. The neon concentrations derived with this method do however yield concentrations that go up to 25%, which - assuming a factor 3-4 underestimation - would lead to potentially a close 100% neon concentrations defined as n_{Ne}/n_e, with a significant fraction of Ne^{2+} (which would break quasi-neutrality). Could you elaborate on how this apparent non-physical situation could be explained? I.e. Are there other effects that could counter-balance the expected underestimation?

4 Towards Forward Modelling of Synchrotron Radiation

4.1 General assessment

This chapter describes the initial development of a synchrotron radiation forward model that uses a 6D phase-space distribution of runaway electrons as main input and uses that with a newly developed Cherab/Raysect emission module, such that the (reverse) ray tracing can be applied. The use of a 6D phase-space distribution, decoupled from the synchrotron radiation modelling is a novel approach compared to the particle tracing methods used by other synchrotron radiation simulations. It also opens up a lot of interesting applications, including – but not limited to –investigation of the potential pick-up of synchrotron radiation by multiple diagnostic configurations for a 1-time calculated distribution or using parametrized distributions as input to allow sensitivity scans are even Bayesian statistical backward models. The potential power of this novel method for synchrotron radiation simulation was clearly highlighted.

The simplification of the formulas describing synchrotron radiation seems at odds with the inclusion of more complex 6D phase-space distribution as input and the inclusion of reflections. But it can be understood in terms of reducing computational load and as a first step to build further upon. Furthermore, the comparison of the forward model to an actual measurement shows already remarkable agreement, despite the above-mentioned simplification. Obviously a more in-dept study of the potential impact of the simplifications as well as a more rigorous comparison with experimental data would be needed for firm conclusions, but given the fact this is only a first step in the development of this forward model, these initial outcomes are very encouraging.

4.2 Detailed/minor comments

Equation (4.8), (4.11) and (4.12): Shouldn't the c^2 be c? In (4.7) the term $(E/(mc^2))^2$ is dimensionless, but the replacing term $c^2/(lambda\ v_g)$ has the dimension of a velocity.

Page 126, figure 4.11, 4.14 and 4.15: The text highlights the big difference between figure 4.11 and 4.14 due to toroidal asymmetry, and then figure 4.15 was added for clarity, adding both figures next to one and other. However, the figure 4.15a which should be identical to 4.11, looks quite different from 4.11.

4.3 Questions

1. There are 2 main simplifications in this model: the assumption the emission direction coincides directly with the velocity direction of the electron emitting and the fact that the spectrum is a delta function at the wavelength directly defined by the electron's energy. Which of those two simplifications is expected to have the biggest effect and would hence have priority in further developing of the model?

2. For runaway detection in larger machines it is crucial to be fast, such that mitigation action can be taken before the runaway beam drills a hole in e.g. a cooling channel. Could you elaborate on how your model could help identifying methods based on Synchrotron radiation to that would allow to quickly identify a runaway electron synchrotron signature in fast (likely non-imaging) detection systems.

5 Overall conclusion

The work presented in this doctoral thesis builds upon both preexisting theoretical concepts and software libraries, but the implementation for Thomson Scattering, divertor spectroscopy and synchrotron radiation of runaway electrons is innovative, has high scientific relevance and applications for diagnostic design, diagnostic data analysis and even potentially for the control and protection of fusion reactors.

The application of the concepts of statistical backward modelling showed the ability of the author for creative scientific work. Examples in each of the 3 topics are:

- The treatment of the tail in the Thomson signal using a time varying uncertainty function.
- The validation method for neon divertor spectroscopy where the complex 2D nature of the divertor is mapped onto a limited number of observation channels.
- The decoupling of the runaway phase space distribution from the synchrotron radiation modelling itself.

In terms of clarity in the presentation of the work the chapters on Thomson Scattering and synchrotron radiation were of excellent quality. The chapter on neon concentration inference was harder to follow. This is in part due to the inherent complexity of and the many unknowns in divertor modelling and spectroscopy. Nevertheless, that particular chapter could have benefit from a bit less condensed writing.

Finally in the Thomson Scattering section is would have been good to refer to recent early work on Bayesian statistical analysis for JET and W7-X. Nonetheless, I still consider the work presented here sufficiently novel and different from the abovementioned work.

In conclusion, in my opinion this doctoral thesis is worthy of a PhD-degree.