Bayesian Analysis of Electron Cyclotron Emission Measurements at Wendelstein 7-X

vorgelegt von
M.Sc.
Udo Höfel
ORCID: 0000-0003-0971-5937

von der Fakultät II – Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
— Dr. rer. nat. —

genehmigte Dissertation

Promotionsausschuss:
Vorsitzender: Prof. Dr. Holger Stark
Gutachter: Prof. Dr. Dieter Breitschwerdt
Gutachter: Prof. Dr. Robert Wolf
Gutachter: Prof. Dr. Hans-Jürgen Hartfuß
Tag der wissenschaftlichen Aussprache: 4. Februar 2020

Berlin, 2020

I sought the fount of fire in hollow reed
Hid privily, a measureless resource
For man, and mighty teacher of all arts.

———————————————————————

Aeschylus, Prometheus Bound

Abstract

E LECTRON CYCLOTRON EMISSION spectroscopy (ECE) is a standard diagnostic technique on the optimized stellarator Wendelstein 7-X (W7-X) that can record data with a high resolution in time. The spatial assignment via the blackbody emission of a plasma layer depends on the optical depth and thus plasma parameters as well as the magnetic field along the line of sight. The measurements with a multichannel radiometer contain a large amount of information about the electron temperature profile, as well as being very sensitive to the magnetohydrodynamic equilibrium at W7-X.

First, the diagnostic was comissioned and absolutely calibrated. At W7-X this is achieved by an optical system identical to the plasma measuring system, which alternately measures room temperature and liquid nitrogen temperature of a microwave blackbody radiator by means of a rotating mirror. The signal difference associated with the temperature change then permits the determination of the calibration factors, the accuracy of which represents the most important source of uncertainty for the diagnostic. In order to allow a systematic treatment of the uncertainties, a completely new, general Bayesian forward model of a calibration unit with rotating mirror was developed and tested within the Bayesian modeling framework Minerva.

The calibrated data then allow to obtain a radiation temperature spectrum. The actual desiderata, i. e. the sought-after quantities, however, are the electron temperatures on the effective plasma radius. Traditionally, the emission region is approximated by the cold resonance location. However, this method is inaccurate if, for example, relevant plasma pressure is reached that leads to a modification of the magnetic field along the line of sight. For the more precise determination of the emission region and the underlying electron temperature profiles, forward modelling must therefore be carried out taking into account the radiation transport along the line of sight. Furthermore, the ray should be determined by raytracing, since ray deflection via the plasma parameter-dependent refractive index can have a serious influence on the model predic-

tions, especially at higher densities. Both is achieved by incorporating the tracing visualized (TRAVIS) code into Minerva, in which the forward model of the electron cyclotron emission (ECE) for W7-X is written. The model includes a prediction of line-integrated electron density via interferometry. One of the advantages of this completely new model is that it is relatively general and should allow easy transferability to other machines, as well as compatibility with the »plug'n'play« neural network generator currently in development.

As examples for applications, the model is used to obtain information about the absolute values of the electron density profile during low and high density plasma discharges, which is a good addition to the already existing possibilities of density measurements by, for example, the Thomson scattering diagnostic. Finally, the ECE data is used on a simple Bayesian heatwave analysis model in an attempt to obtain the electron heat diffusivity.

Zusammenfassung

Die Elektronen-Zyklotron-Emissionsspektroskopie (ECE) ist eine Standarddiagnostik am optimierten Stellarator Wendelstein 7-X (W7-X), die zeitlich hochaufgelöste Messdaten aufnehmen kann. Ihre räumliche Zuordnung über die Schwarzkörperemission einer Plasmaschicht hängt von der optischen Tiefe und damit Plasmaparametern sowie dem Magnetfeld entlang der Beobachtungsrichtung ab. Die Messungen mit einem Vielkanalradiometer enthalten einen großen Informationsanteil über das Elektronentemperaturprofil, auch ist die Diagnostik eine sehr empfindliche Messmöglichlichkeit des Magnetohydrodynamikequilibriums an W7-X.

Zu Beginn wurde die Diagnostik neu in Betrieb genommen und absolut kalibriert. Am W7-X wird dies durch eine zur Plasmamessung identisch aufgebaute Optik erreicht, die mittels einem rotierenden Spiegel abwechselnd Raumtemperatur und Flüssigstickstofftemperatur eines Mikrowellenschwarzkörperstrahlers misst. Die mit der Temperaturveränderung einhergehende Signaldifferenz erlaubt dann die Bestimmung der Kalibrierfaktoren, deren Genauigkeit die bedeutendste Fehlerquelle der Diagnostik darstellt. Um eine systematische Behandlung der Unsicherheiten zu erlauben, wurde ein komplett neues, generelles Bayessches Vorwärtsmodell einer Kalibriereinheit mit rotierendem Spiegel entwickelt und im Rahmen des Bayesschen Modellierungsframeworks Minerva getestet.

Die kalibrierten Daten erlauben dann die Messung eines Strahlungstemperaturspektrums. Der eigentlich gewünschte Wert ist jedoch die Elektronentemperatur in Abhängigkeit vom effektiven Plasmaradius. Traditionell erfolgt die Approximation des Emissionsbereiches über den Ort der kalten Resonanz, allerdings ist diese Methode ungenau, wenn beispielsweise relevante Werte des Plasmadrucks erreicht werden, was zu einer Modifikation des Magnetfeldes entlang der Sichtlinie führt. Für die genauere Bestimmung des Emissionsbereiches und der zugrundeliegenden Elektronentemperaturprofile ist darum eine Vorwärtsmodellierung unter Berücksichtigung des Strahlungstransportes entlang

der Sichtlinie durchzuführen. Der Strahlverlauf im Plasma sollte des Weiteren mittels Raytracing bestimmt werden, da die Strahlablenkung über den plasmaparameterabhängigen Brechungsindex insbesondere bei höheren Dichten die Modellvorhersagen gravierend beeinflussen kann. Dies wird durch einbinden des Strahlungstransport-Raytracing-Codes tracing visualized (TRAVIS) in Minerva erreicht, in welchem auch das Vorwärtsmodell der ECE für W7-X geschrieben ist. Das Modell beinhaltet des Weiteren eine Vorhersage der linienintegrierten Elektronendichte via Interferometrie. Einer der Vorteile dieses komplett neuen Modells ist, dass es relativ allgemein gehalten ist, und somit eine leichte Übertragbarkeit auf andere Maschinen erlauben sollte, sowie die Kompatibilität mit dem zurzeit in Arbeit befindlichen »Plug'n'Play«-Neuronale-Netze-Generator.

Die Diagnostik wird beispielsweise angewandt um zu versuchen bei Niedrig- und Hochdichteplasmaentladungen Informationen über die absoluten Werte des Elektronendichteprofils zu erlangen, was eine gute Ergänzung zu den bereits existierenden Möglichkeiten der Dichtemessungen durch beispielsweise die Thomsonstreuungsdiagnostik darstellen würde. Schließlich werden die ECE-Daten verwendet um mittels eines simplen Bayesschen Wärmewellenmodells auf die Elektronenwärmediffusivität zu schließen.

Contents

Contents

1

Introduction

The primary electricity production has grown on average 3.6 % per year between 1985 and 2011 [1]. Conservative estimates expect a further threefold increase between 2000 and 2100 [2]. As mankinds hunger for energy keeps growing, climate change poses serious threats to the stability of societies and the environment [3]. Using renewable energy sources helps alleviating CO_2 emissions, but has drawbacks, for example with respect to the thus required energy storage capabilities [4].

Nuclear fusion is a new primary, clean energy source under investigation. It is expected to contribute to the energy mix in the second half of the century [5]. The advantages of nuclear fusion include large amounts of fuel, lack of both CO_2 emissions and long-lived radioactive waste [6].

However, for fusion to happen, a deuterium-tritium plasma with temperatures exceeding 100 million kelvin needs to be confined contactlessly in a vacuum chamber with a confinement time on the order of seconds. The two main magnetic confinement concepts are the tokamak [7, 8] and the stellarator [9]. Both have a toroidal magnetic field, but different ways to generate an additional poloidal magnetic field. Tokamaks often use the ohmic current induced by a transformer for this, while stellarators rely on external coils. The superposition of the magnetic fields leads to twisted magnetic field lines. The ratio of poloidal to toroidal turns is called rotational transform, ι. To achieve sufficient confinement, closed and nested magnetic flux surfaces with a finite ι are necessary. Flux surfaces are defined as a surface for which $\vec{B} \cdot \vec{n} = 0$ with $\vec{B}$ the magnetic field and $\vec{n}$ the normal to the surface holds. The force balance equation,

$$\vec{\nabla} p = \vec{j} \times \vec{B},\qquad(1.1)$$

allows to get the static, ideal magnetohydrodynamic (MHD) equilibrium with p the pressure and $\vec{j}$ the plasma current. As $\vec{\nabla}p \perp \vec{B}$, the pressure is constant along a field line. Hence, having nested flux surfaces made up of field lines is important for analyses of magnetic confinement plasmas, as they allow to reduce the dimensionality that has to be handled within the flux surfaces from 3D to 1D. In a tokamak the poloidal magnetic field is generated by the plasma current, which allows to keep the plasma toroidally symmetric. Tokamaks suffer from disruptions [10] and plasma pressure limiting neoclassical tearing modes that degrade the plasma energy and can initiate current driven instabilities [11, 12]. Further, it is difficult to get tokamaks with steady state capability due to their reliance on the plasma current to get the desired shape of the magnetic cage. Recently, notable progress in non-inductive operation has been achieved at ASDEX Upgrade (AUG) [13]. The magnetic cage of stellarators is in contrast provided entirely by external coils, thus easing steady state operation due to easier control of the current in external coils. Roughly similar construction costs are anticipated for tokamaks and stellarators on comparable reactor scales [14].

Table 1.1.: Typical (plasma) parameter values at Wendelstein 7-X.

parameter	typical value
B_0	2.5 T
$T_{e,0}$	$\approx 4\,\text{keV}$
$T_{i,0}$	$\approx 1.5\,\text{keV}$
$\bar{n}_e$	$\approx 5 \times 10^{19}\,\text{m}^{-3}$
τ_E	$\approx 150\,\text{ms}$
t_{puls}	$\approx 10\,\text{s}$

For stellarators the energy confinement time τ_E, representative of the quality of the thermal insulation, can be described with scaling laws, i.e. the international stellarator scaling 1995 (ISS95) [15], or more recently, the international stellarator scaling 2004 (ISS04) [16]. ISS04 predicts

$$\tau_E \propto \bar{n}_e^{0.54}, \tag{1.2}$$

with $\bar{n}_e$ the line averaged electron density. Thus, high densities are beneficial, as they increase the energy confinement time. Neoclassical transport, i.e. transport considering geometrical effects neglected in classical transport theory, is detrimental to good confinement. Consequently, neoclassical optimization [17] aims to lessen the impact of neoclassical transport [18]. It necessitates high densities as well [19]. Furthermore, high edge densities are mandatory for safe divertor operation as they help to alleviate the heat load the divertor has to withstand [20, 21]. The densities achievable in stellarators are limited by heating power and radiation losses [22–24], in contrast to tokamaks, which have the so-called Greenwald limit [25] as the upper density limit.

The most modern, optimized stellarator is Wendelstein 7-X (W7-X) [24, 26–29], cf. figure 1.1 and table 1.1. W7-X makes use of modular coils, in contrast to »classical« stellarators, reducing stress between the coils. This thesis focusses on the analysis at W7-X, although the findings should apply to tokamaks as well. At W7-X, the backbone of the plasma heating is done via electron cyclotron resonance heating (ECRH) [30] that allows practical heating up to electron densities of about $10^{20}\,\mathrm{m}^{-3}$ in extraordinary mode (X mode), or about $2 \times 10^{20}\,\mathrm{m}^{-3}$ in ordinary mode (O mode) for up to 30 min, which is a necessity to show the intended steady state capability. In principle, neutral beam injection (NBI) is also available and has no upper density limit, but in contrast to the ECRH this system, at least as currently implemented in W7-X, is not steady state capable. Additionally, ECRH offers the option to drive plasma currents (i. e. the so-called electron cyclotron current drive (ECCD) [31–33]) and shape the profiles of plasma parameters, which can be an important aspect in controlling the plasma [34]. For the first time plasma heating using only the second harmonic O mode has been achieved with the ECRH system at W7-X [35], thus opening the operational window to higher electron densities.

The complementary physical process to ECRH is electron cyclotron emission (ECE), which can be detected via microwave spectroscopy and can be used to obtain electron temperature profiles with high temporal resolution [36, 37]. In a first, simple approach, the measured ECE intensity can be considered to stem from a blackbody, which is a reasonable first estimation for typical plasma parameters at W7-X (cf. table 1.1) given a carefully chosen line of sight allowing to obtain the electron temperature. The location where the plasma emission originates from can be approximated by the so-called cold resonance such that radially resolved electron temperature profiles can be obtained. This approximation neglects relativistic effects, consequently the radiation originates solely from electrons emitting at their cyclotron frequency. More details about ECE and the system at W7-X can be found in chapter 2. Bayesian modelling, summarized in chapter 3, is used to optimally combine information from different measurements, and to properly formalize the physics models. Thus, a model for the absolute intensity calibration of the ECE has been developed and deployed, cf. chapter 4. However, the calibrated intensities cannot always be equated with the electron temperature directly as there are a number of effects that may complicate the interpretation. For instance, the locations from which the emission originates depend on the radiation transport, potentially deviating notably from the cold resonance. In addition, the electron energy distribution function may
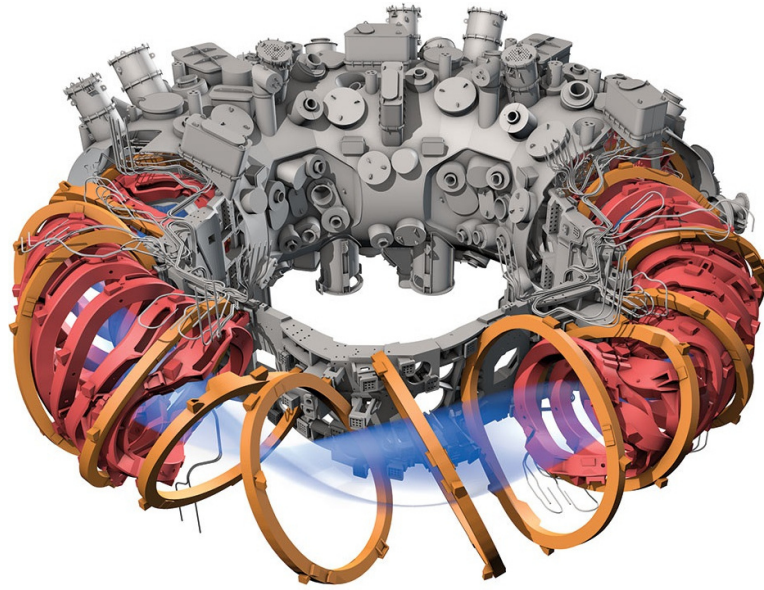
Figure 1.1.: Schematic structure of W7-X. The half transparent blue structure indicates the plasma (corresponding to $\approx 30\,\mathrm{m}^3$), red and orange the modular and planar coils and grey the outer vessel including the ports. © IPP

contain a sufficient number of electrons in the high energy tail that contribute to spectral intensities at significantly lower frequencies and thus hamper the simple cold resonance approximation. These electrons emit at lower frequencies than naïvely to be expected due to their relativistic mass increase. Especially in plasmas with lower electron densities this leads to the so-called »shinethrough« of the emission that stems from hot core electrons, effectively preventing ECE measurements to be interpreted as electron temperatures. Furthermore, at high electron densities refraction of the line of sight starts to play a role. Above a certain electron density, the ECE cutoff conditions are fulfilled for the probed frequencies, preventing the microwave emission and no longer contributing directly to the measured intensity. For the interpretation of the ECE it is therefore preferable to not rely on the blackbody cold resonance ansatz due to the aforementioned effects. This is especially the case when approaching operational parameters at which the classical approach starts to break down. At W7-X this is frequently the case, as it aims at high electron densities, near or above the cutoff densities of the second harmonic X mode. Consequently, a model that

combines the full radiation transport calculations for the ECE system with the line integrated electron density from a single channel dispersion interferometry system, particularly for W7-X, is introduced in chapter 5. This Bayesian model in turn allows to push the model to its high density limits as detailed in section 6.1. The full potential of ECE as a diagnostic makes use of its high temporal resolution, allowing for measurements of dynamic plasma phenomena. As an example, ECE can be used to infer information about the electron heat transport in a plasma, for example via heat pulse modulation experiments [38] or via transfer entropy analysis [39]. Examples of this are highlighted in section 6.3. Finally, the conclusions in chapter 7 wrap the thesis up.

2

Electron cyclotron emission spectroscopy

One of the advantages of ECE spectroscopy is that the functionality of the optical equipment needed near the torus, i. e. mirrors, horns and waveguides, is not sensitive to neutron radiation. The analysis of ECE radiation is a standard tool [36, 40] to obtain with a passive measurement, i. e. a measurement that does not perturb the plasma, information about the electron temperature at a high sampling rate. Thus, this chapter will first explain the general principles at work, where the first part tries to give a more intuitive understanding and the second part briefly goes through the relevant equations describing ECE quantitatively. Then, the tracing visualized (TRAVIS) code used to calculate the quantitative equations, its advantages and limitations are shortly explained. Finally, the ECE at W7-X, especially the experimental setup and the calibration unit, are introduced. Parts of this chapter have been published in [41].

2.1. General principle

This section aims to give an overview of the governing principles relevant for ECE. First, a more qualitative description of the physics is given in section 2.1.1. Then, section 2.1.2 describes the physical processes quantitatively.

2.1.1. Qualitative description

As typical plasma experiments aiming to pave the way to a fusion reactor are operated with strong magnetic fields, the charged particles making up the plasma

are subject to the Lorentz force, therefore leading to gyrations around the magnetic field lines. For the radius of the gyration one finds the so-called Larmor radius,

$$r_{\mathrm{L}} = \frac{m_{\mathrm{e},0} v_\perp}{eB} = \frac{v_\perp}{\omega_{\mathrm{c},0}}. \tag{2.1}$$

Here, $m_{\mathrm{e},0}$ is the electron rest mass, v is the velocity of the electron, B the magnetic field, e the electron charge and $\omega_{\mathrm{c},0}$ the cyclotron frequency. Due to the finite Larmor-radius effects in magnetized plasmas the emission near the harmonics of the cyclotron frequency can be calculated via

$$f_{\ell,0} = \ell \frac{\omega_{\mathrm{c},0}}{2\pi} = \ell \frac{eB}{2\pi m_{\mathrm{e},0}} \tag{2.2}$$

with ℓ corresponding to the harmonic.

According to equation (2.2) the emission frequency is tied to a magnetic field. The magnetic field should be strictly monotonically increasing along the line of sight of the ECE (looking towards the emission) to allow an unambiguous interpretation[1]. Therefore, the intensity collected at a higher frequency will be emitted at a greater distance from the antenna. The presence of a plasma modifies the magnetic field strength locally, e. g. by the diamagnetic effect and the Shafranov shift [42], which can be described e. g. via a variational moments equilibrium code (VMEC) equilibrium. Hence, a VMEC equilibrium provides a relationship between the magnetic field strength in real space and a normalized effective radius coordinate system describing the plasma within the last closed flux surface. The effective radius coordinate system is important to ease analysis by lowering the dimensionality.

Thus, solving equation (2.2) for B ties an ECE channel frequency via VMEC to an effective radius ρ. This is the so-called cold resonance,

$$\rho_{\mathrm{CR}}(f) = g_{\mathrm{VMEC}}\left(\frac{2\pi f_{\ell,0} m_{\mathrm{e},0}}{e\ell}\right) \tag{2.3}$$

In principle this is not tied to VMEC as any method providing a similar mapping can be used.

---

[1] While the W7-X ECE system's low field side antenna fulfills this condition, the antenna on the high field side (HFS) does not. However, as the HFS antenna was not used anywhere throughout the thesis, the description henceforth deals only with the case of a monotonically increasing magnetic field strength along the line of sight.
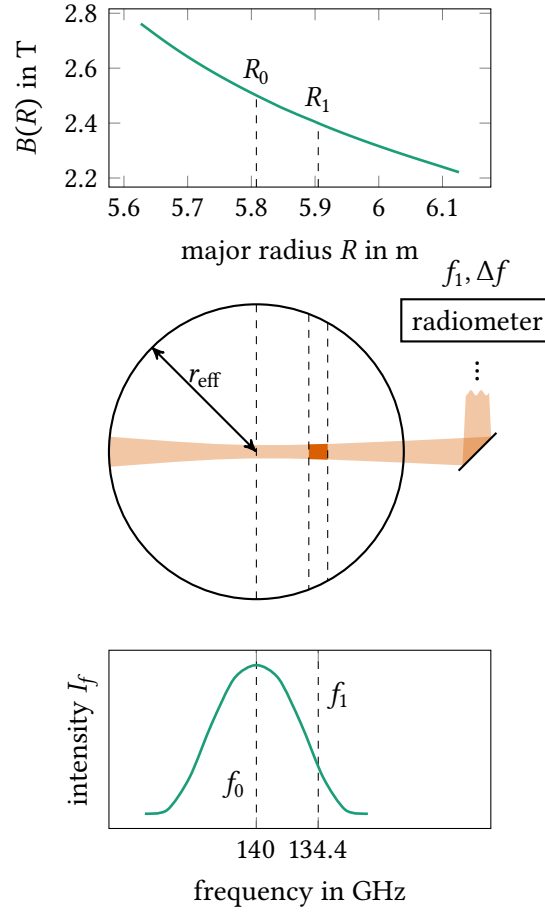
Figure 2.1.: Illustration of the principle of ECE. The magnetic field corresponds to the equilibrium used in chapter 5 and the frequencies correspond to real values of ECE channels at W7-X, while the shape of the intensity curve is purely schematic. Based on [36].

Figure 2.1 shows the general principle that an ECE spectroscopy system in a plasma fusion experiment makes use of. The minor effective radius is denoted r_{eff} (approximately 0.52 cm in W7-X), the probing ECE channels center frequency is called f_1 with a bandwidth of Δf broad enough to collect meaningful intensities. For this example the chosen frequency corresponds to channel 9. The plasma center is roughly at 140 GHz, denoted by f_0. In the top part of the picture one can see the magnetic field strength over the major radius of W7-X

along the ECE line of sight. The plasma center is roughly at $R_0 \approx 5.81\,\mathrm{m}$. Further-more, channel 9 is chosen exemplarily on the gradient region of the profile. In the top part of the figure channel 9 corresponds to $R_1 \approx 5.91\,\mathrm{m}$. The middle part of the figure shows the plasma center and channel 9 indicated by dashed lines in an idealized cut through the torus. Channel 9 corresponding to frequency f_1 has a finite bandwidth Δf, thus the radiation collected by the radiometer stems from the orange region, provided the relativistic effects in a hot plasma are neglected. The lower part of the figure shows schematically the intensit-ies as measured by the radiometer. For radiometers like the one at W7-X, the intensity is, in a first approximation for sufficiently dense and hot plasmas, pro-portional to the electron temperature. This assumes that the plasma acts like a blackbody, in which case the Rayleigh-Jeans approximation is valid. Thus, a calibration with other blackbody sources of known temperatures is possible. This calibration (cf. chapter 4) allows to transform the measured intensity into a so-called radiation temperature which is the temperature a blackbody would have to have to produce the measured intensities. Much of the following deals with deviations of the radiation temperature from the ideal blackbody case (cf. chapter 5).

Figure 2.2 shows the electron cyclotron frequencies for the first three har-monics. Due to the large aspect ratio of W7-X, which is approximately 10, the different harmonics are well separated. This allows an easier interpretation of the ECE spectrum, as the radiation collected by a frequency channel stems from one place in the plasma (assuming sufficient optical thickness and neglect-ing Doppler and relativistic effects) and is not a mixture of contributions from different harmonics.

As one can see from equation (2.2) the relativistic mass gain of the electrons is not taken into account. However, in a typical plasma experiment aiming at nuclear fusion electrons in the high energy tail of the electron energy distri-bution function reach velocities at which the relativistic mass gain is notable. Furthermore, depending on the angle, Doppler shifts can play a relevant role for the frequencies observed in the laboratory coordinate system. These effects can be taken into account by

$$f_\ell = f_{\ell,0} \frac{\sqrt{1 - (v/c)^2}}{1 - (v_\parallel/c)\cos(\theta)}, \tag{2.4}$$

with $f_{\ell,0}$ the nonrelativistic electron cyclotron frequency (cf. equation (2.2)), v the electron velocity, $v_\parallel$ the electron velocity along the magnetic field lines, c
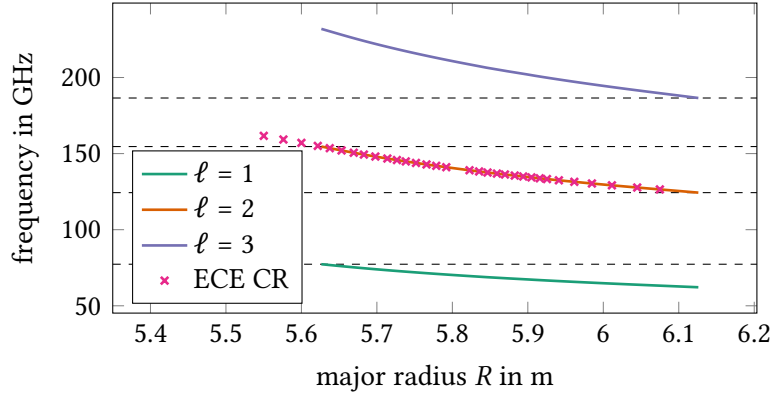
Figure 2.2.: First, second and third harmonics of electron cyclotron frequencies (cf. equation (2.2)) along the line of sight of the ECE at W7-X. The cold resonance (CR) positions are depicted with purple crosses. One can see that the harmonics are well separated. Hence there is no harmonic overlap, thus this poses no problem for the localisation of the ECE channels. The equilibrium used is the same as the one used for the analysis in chapter 5.

the speed of light and θ the angle of the observer relative to the magnetic field lines. The relativistic mass increase, represented by the term in the numerator, leads to a red shift. This is schematically shown in figure 2.3 (red line). In this example, radiation emitted at the major radius of 5.77 m is (partially) redshifted, such that it appears to be a contribution from around 5.9 m (corresponding to f_1). This complicates the interpretation of the obtained spectrum and highlights the need for a way to take relativistic downshifts into account. For the Doppler shift, corresponding to the term in the denominator, one can see that it can be minimized by choosing a sightline perpendicular to the magnetic field lines, which is well satisfied for the low field side antenna of the ECE at W7-X.

The given explanations consider single electrons. Taking into account that the electrons at each location exhibit a distribution of velocities leads to a relatively broad range of major radii where the electron population includes electrons that emit radiation with the same frequency, as shown by the marine curve in figure 2.3. However, much of the radiation is reabsorbed by resonant electrons along the line of sight, effectively narrowing the broad spatial emissivity profile. For an example, see the curve filled with orange in figure 2.3. The propagation of radiation through the plasma is described by radiation transport.
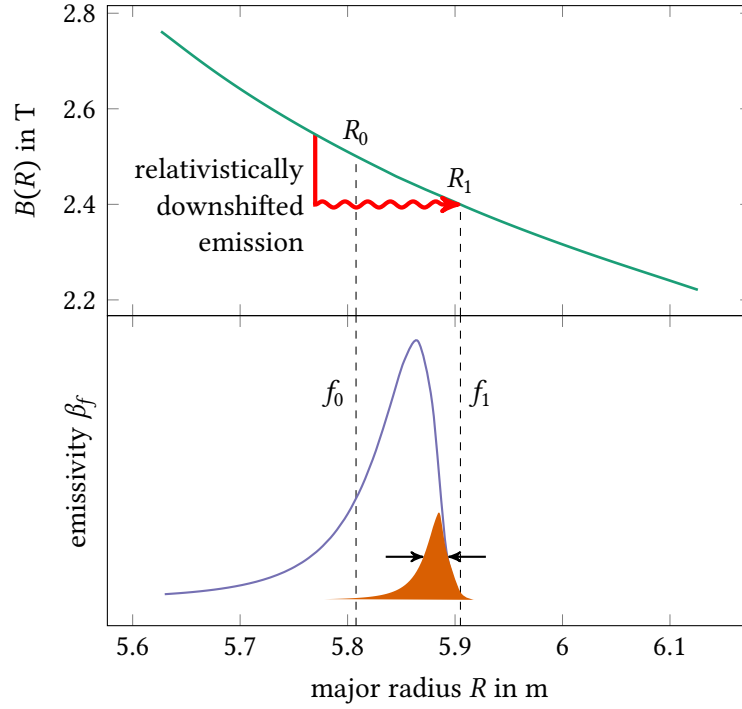
Figure 2.3.: Illustration of the effect of radiation transport on the ECE. The magnetic field and major radii correspond to real values of W7-X, while the shape of the emissivity curve is purely schematic.

It is important to note that the description of the wave emission of gyrating electrons in a magnetized plasma within the »cold« plasma approximation allows for two solutions for lines of sight perpendicular to the magnetic field; one solution has the electric field of the wave parallel to the magnetic field, yielding the so-called O mode, and the other solution has the electric field perpendicular to the magnetic field, yielding the so-called X mode. Consequently, their cutoff densities in a plasma appear at different frequencies. The O mode and X mode have their cutoff densities at

$$n_{\mathrm{e,O}}^{\mathrm{cut}} = \frac{(2\pi f_{\ell,0})^2 \varepsilon_0 m_{\mathrm{e,0}}}{e^2} \overset{\mathrm{W7X}}{\approx} 2.4 \times 10^{20}\,\mathrm{m}^{-3} \text{ for } f_{\mathrm{obs}} = 140\,\mathrm{GHz}, \tag{2.5a}$$

$$n_{\mathrm{e,X}}^{\mathrm{cut}} = n_{\mathrm{e,O}}\left(1 - \frac{f_{\mathrm{c,0}}}{f_{\mathrm{obs}}}\right) \overset{\mathrm{W7X}}{\approx} 1.2 \times 10^{20}\,\mathrm{m}^{-3} \text{ for } f_{\mathrm{obs}} = 140\,\mathrm{GHz} \text{ at } B = 2.52\,\mathrm{T},$$

$$\tag{2.5b}$$

with ε_0 the vacuum permittivity, f_{obs} the observation frequency and the approximate values specific for the standard ECE at W7-X. The optical thickness τ, a measure for the opacity of the plasma for the wave (cf. section 2.1.2 and equation (2.23)), is different for the two modes. It depends largely on the electron temperature, density and the gradient length of the magnetic field. This means that the plasma can be optically thick for the X mode, so one pass through the plasma is enough to obtain radiation temperatures reasonably close to the electron temperatures, while the O mode may still see an optically thin plasma. Therefore, an appropriate description of the O mode would require to take reflections at the vessel wall into account. A model assuming infinite reflections between vessel walls, acting as mirrors perpendicular to the wave propagation direction, is an often used approximation for many ECE systems. According to the infinite reflection model the intensity is [40],

$$I = I_0 \frac{1 - \exp(-\tau)}{1 - R_w \exp(-\tau)}, \tag{2.6}$$

with I_0 the intensity one would expect from an optically thick plasma, τ the optical thickness, R_w the wall reflection coefficent and I the intensity corrected for optical thickness effects.

To relax the assumptions mentioned above required for the infinite reflection model, which are clearly violated in W7-X due to its complex 3D geometry of the plasma vessel, a complete 3D ray tracing is beneficial. For practical usage, a quantitative description of ECE spectroscopy is required, thus the necessary formulæ are outlined in the next subsection.

2.1.2. Quantitative description

A quantitative description of the major physical processes leading to ECE spectroscopy measurements is crucial to obtain detailed and reliable electron temperature profiles. This subsection follows [40, 43] and uses the notation of [43].

The radiation transport along a path s through the plasma is described by

$$N_r^2 \frac{\mathrm{d}}{\mathrm{d}s}\left(\frac{I_\omega}{N_r^2}\right) = \beta_\omega - \alpha_\omega I_\omega, \tag{2.7}$$

with N_r the ray refractive index, I_ω the wavelength dependent intensity, β_ω the emissivity and α_ω the cyclotron absorption coefficient. Each of these variables

13

will be explained subsequently. New paragraphs indicate where the next of these variables is defined due to the amount of space required for each explanation.

Ray refractive index   First, the ray refractive index is given by

$$N_r^2 = N^2 \left| \cos(\vartheta) \frac{\mathrm{d}\Omega}{\mathrm{d}\Omega_k} \right|^{-1} \tag{2.8}$$

with N the refractive index. Furthermore,

$$\cos(\vartheta) = (\mathbf{k} \cdot \mathbf{F})/(kF), \tag{2.9}$$

the ray solid angle is $\mathrm{d}\Omega$ and the wave vector solid angle is $\mathrm{d}\Omega_k$ [44]. The wave vector is specified by $\mathbf{k}$ and the dimensionless normalized power flux density by $\mathbf{F}(\omega, \mathbf{k})$.

Emissivity   Second, for the emissivity one finds

$$\beta_\omega = N_r^2 \frac{m_e \omega^3}{8\pi^2 c^2 |\mathbf{F}(\omega, \mathbf{k})|} \frac{\omega_{p,e}^2}{\omega^2} \sum_\ell \int \mathrm{d}u_\parallel [\mathcal{D}_{\mathrm{ql}} f_e]_{\gamma = \gamma_{\mathrm{res}}}, \tag{2.10}$$

with the electron plasma frequency given by

$$\omega_{p,e} = \sqrt{\frac{n_e e^2}{\varepsilon_0 m_e}}, \tag{2.11}$$

ℓ the harmonics number, $u_\parallel$ the momentum per unit mass along the ray, the normalized quasilinear diffusion coefficient

$$\mathcal{D}_{\mathrm{ql}} = (u_\perp/c)^2 |\mathbf{\Pi}_\ell|^2, \tag{2.12}$$

which in turn contains the complex polarization vector

$$\mathbf{\Pi}_\ell = \mathbf{e}_- \mathcal{J}_{\ell-1}(k_\perp r_{\mathrm{L}}) + \mathbf{e}_+ \mathcal{J}_{\ell+1}(k_\perp r_{\mathrm{L}}) + \mathbf{e}_\parallel u_\parallel \mathcal{J}_\ell(k_\perp r_{\mathrm{L}})/u_\perp \tag{2.13}$$

with $\mathcal{J}$ Bessel functions of the first kind,

$$\mathbf{e}_\pm = \mathbf{e}_x \pm \mathrm{i}\mathbf{e}_y, \tag{2.14}$$

as well as

$$k_\perp r_\text{L} = \frac{N_\perp u_\perp \omega}{\omega_{\text{c},0} c} \tag{2.15}$$

with the Larmor radius (cf. equation (2.1)), $u_{\perp/\parallel} = \gamma v_{\perp/\parallel}$ and $N_\perp$ the refractive index perpendicular to the magnetic field. For the electron velocity distribution function f_e often a normalized relativistic Maxwellian is assumed,

$$f_\text{e}(u) = \left[\sqrt{\frac{\pi}{2\mu}} \frac{\exp(-\mu)}{\mathcal{K}_2(\mu)}\right] \frac{\mu^{3/2}}{\sqrt{2\pi}} \exp\left(-\mu(\gamma - 1)\right), \tag{2.16}$$

wherein

$$\mu = m_\text{e} c^2 / T_\text{e}, \tag{2.17}$$

with T_e in units of keV. The Bessel function of the second kind is denoted by $\mathcal{K}$. As the integral of equation (2.10) is performed along the resonance lines, the integrand is only evaluated where $\gamma = \gamma_\text{res}$ holds. Therein,

$$\gamma = \sqrt{1 + (u/c)^2}, \tag{2.18}$$

$$\gamma_\text{res} = (\omega_{\ell,0}/\omega) + N_\parallel u_\parallel / c. \tag{2.19}$$

Note that the integral in equations (2.10) and (2.20) is just over $u_\parallel$, thus the starting and end points of the integral can be determined by solving $\gamma = \gamma_\text{res}$ for $u_\perp = 0$.

Cyclotron absorption coefficient   Third, the cyclotron absorption coefficient can be calculated via

$$\alpha_\omega = -\frac{\pi\omega}{c^2 |\mathbf{F}(\omega, \mathbf{k})|} \frac{\omega_{\text{p,e}}^2}{\omega^2} \sum_\ell \int \mathrm{d}u_\parallel [\mathcal{D}_\text{ql}\hat{\mathcal{L}}(f_\text{e})]_{\gamma = \gamma_\text{res}}. \tag{2.20}$$

In there, the quasilinear differential operator is defined by

$$\hat{\mathcal{L}} = \partial_\gamma + cN_\parallel \partial_{u_\parallel}. \tag{2.21}$$

One can find that for a Maxwellian given by equation (2.16) α_ω and β_ω satisfy Kirchhoff's blackbody radiation law,

$$\frac{\beta_\omega}{\alpha_\omega} = N_r^2 \frac{\omega^2}{8\pi^3 c^2} T_\text{e}. \tag{2.22}$$

The cyclotron absorption coefficient allows the calculation of the optical thickness along the ray,

$$\tau = \int \alpha_\omega \, ds. \tag{2.23}$$

Spectral intensity   Forth, solving equation (2.7) at the point where the ray leaves the plasma on its way to the antenna, b, leads to the following definition of the spectral intensity at the microwave receiver optics

$$I_\omega(b) = I_\omega^{\mathrm{inc}} \exp\big(\tau_\omega(a) - \tau_\omega(b)\big) + \int_a^b ds' \frac{\beta_\omega}{N_r^2} \exp\big(-\tau_\omega(b) + \tau_\omega(s')\big), \tag{2.24}$$

where I_ω^{inc} is the spectral intensity at the point at which the ray ending in b entered the plasma. This point is described by a. It is clear that for optically thick plasmas I_ω^{inc} is neglectable as the exponential term will damp its contributions strongly. Note that $N_r(s = a, b) = 1$ needs to be fulfilled as a boundary condition, effectively requiring the electron density and temperature to vanish at the plasma edge.

Finally, one can predict the radiation temperature to be expected from a plasma with given electron temperature and density with

$$T_{\mathrm{rad}}(\omega) = \frac{8\pi^3 c^2}{\omega^2} \sum_r w_r I_{\omega,r}(b) \overset{\tau \gg 1}{\approx} T_{\mathrm{e}}, \tag{2.25}$$

wherein the sum collects contributions from multiple rays r (only required if the beam has been split in multiple rays beforehand) reaching the antenna, each weighted by w_r. Given a large optical thickness τ, T_{e} can be very well approximated by T_{rad}, which is the basis for the »classical« ECE analysis. For practical application, Marushchenko and Turkin developed an optimized code for these calculations which is shortly explained in the next subsection.

2.2. Tracing visualized code

TRAVIS is a Fortran code that is used for calculations for the ECE and ECRH systems, as well as for reflectometer systems, mostly at W7-X, but calculations for the International Thermonuclear Experimental Reactor (ITER) have also been

done. This section is based heavily on [43] and summarizes the assumptions that are relevant for the ECE mode of TRAVIS, as well as the (dis)advantages.

TRAVIS contains a ray tracer that takes the electron density and temperatures, as well as the magnetic equilibrium provided for example via VMEC or EFIT into account and implements the theoretical calculations from section 2.1.2. To account for the finite width of the antenna pattern of the ECE diagnostic TRAVIS can automatically create multiple appropriately weighted rays located at concentric circles around the central ray. However, each ray is an independent calculation, such that the computational cost of such an approach can be prohibitive for calculations that require a large number of iterations, like for example a Markov chain Monte Carlo (MCMC) (cf. section 3.3). For such cases using only the central ray offers an approximation that yields good results for beams which do not diverge to a significant extent and the plasma density is far enough from the cutoff density.

In a second step, the radiation transport equation (cf. equation (2.7)) is then solved in reverse direction along the calculated rays such that a radiation temperature for each frequency channel is predicted. Do note that TRAVIS does its calculations either for the X mode or the O mode. This implies that two separate calls to TRAVIS are required to predict both modes. While TRAVIS can handle wall reflections in full 3D, which is required for stellarators due to their complex geometry, transported radiation that would reenter the plasma in the wrong mode is discarded. Computationally, it is expensive to search for the intersections of rays with the wall polygons. For diagnostics like the ECE at W7-X this poses no problem if electron densities above approximately 10^{19} m^{-3} and temperatures above about $1\,$keV are achieved, as the X mode is then optically thick enough that a second plasma pass and thus reflections are not required.

The TRAVIS code is provided in the general W7-X service-oriented architecture (SOA) (cf. [45, 46]) as a webservice to decrease maintenance and improve accessibilty for the W7-X team.

2.3. ECE diagnostic at W7-X

At W7-X, an ECE heterodyne radiometer[2] is the main diagnostic for obtaining electron temperature core profiles with a continuous high temporal resolution

---

[2]That is, the input signal is mixed with the signal of a local oscillator, leading to a signal at the sum and the difference of the two signals.

of up to 2 MHz. In contrast, the other main temperature diagnostics, Thomson scattering [47] and X-ray imaging crystal spectroscopy [48] have sampling rates of up to 30 Hz respectively 500 Hz, although the Thomson scattering system can reach sampling rates of up to 10 kHz in »burst« mode. However, the »burst« mode can only be used for 1.2 ms with a repetition rate of 5 Hz [49] and thus is not capable of providing data with high temporal resolution continuously. Thus, the ECE is crucial in understanding e. g. fast plasma phenomena at W7-X.

2.3.1. Setup

The setup of the standard ECE system at W7-X [37] can be found in figure 2.4. It measures at the outboard, low field side of the torus[3]. As W7-X has a five-fold symmetry, there are 10 toroidal angles[4] at which the whole line of sight of the ECE could have been truly perpendicular to the magnetic flux surfaces. However, an ECE measurement near the plasma center would be completely masked by strong non-absorbed ECRH radiation, as the ECRH heats the plasma in these symmetry planes. The ECE line of sight was chosen to have a small toroidal component, while still being near one of the 10 toroidal angles mentioned above[5]. As a consequence, the ECE does measure a Doppler shift, and refraction does play a role for higher densities (see figure 6.3). However, the sightline has been optimized to reduce the influence of both these effects as much as possible [50]. Note that the line of sight in figure 2.4 does not appear to be perpendicular to the flux surfaces, which is an artifact from the projection into the (R, z) plane in polar coordinates. The VMEC calculated magnetic flux surfaces for a high mirror configuration[6] for operation phase (OP) 1.2b are shown for no plasma beta. A wideband optical Gauss telescope system (i. e. with a wavelength independent waist position [36]) with an $1/e^2$ beam intensity width of approximately 20 mm is used. The variation of the width over the sightline in the plasma (±2 mm) can be considered small.

---

[3]The exact toroidal angle is $\phi \approx 223°$.

[4]To be precise: There are five »triangular« planes and five »bean shaped« planes, i. e. the magnetic flux surfaces take on the shape of a bean, similar to the magnetic flux surfaces shown in figure 2.4, but symmetric around a horizontal cut at $z = 0$. The triangular planes are not well suited for ECE measurements from the low field side as the magnetic field is not monotonously increasing along the potential line of sight.

[5]The toroidal angle $\phi = 216°$ corresponds to the bean shaped plane due to the fivefold symmetry of W7-X ($\phi = 0°, 72°, 144°, 216°, ...$).
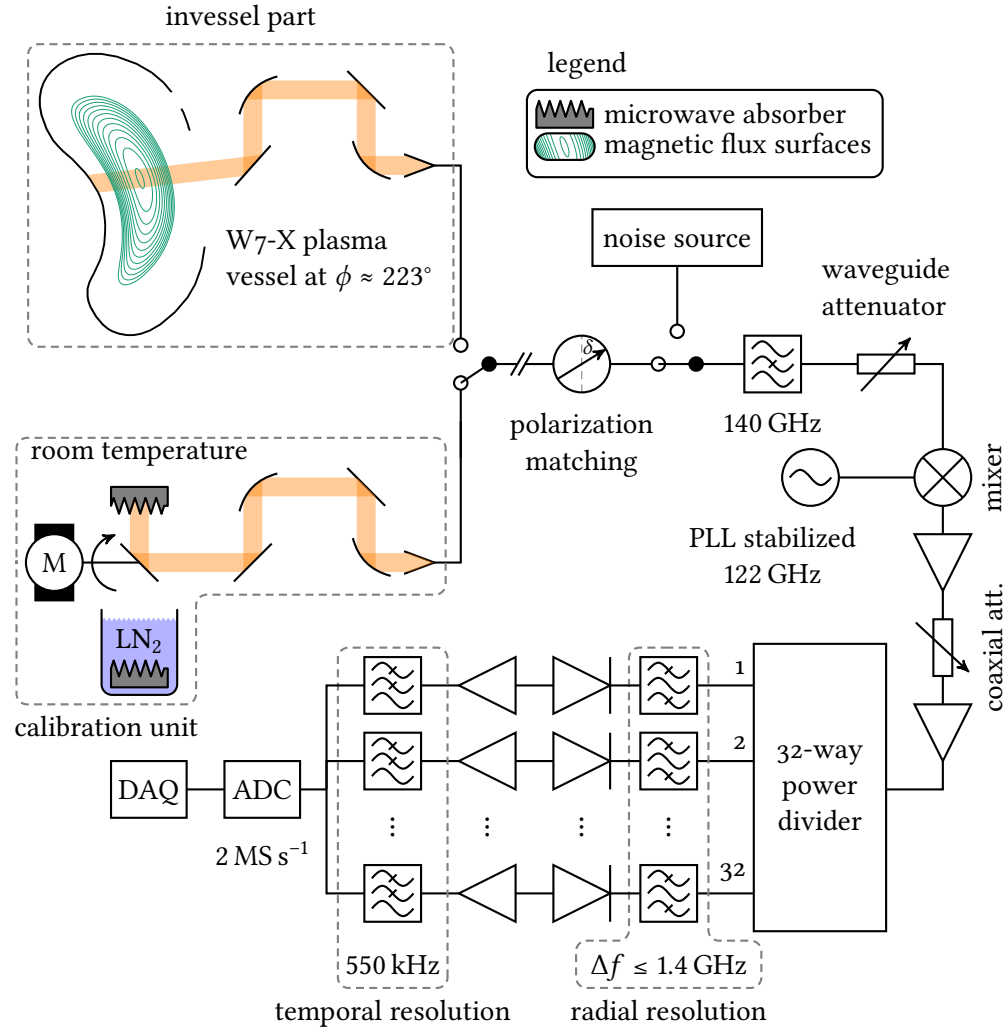
[6]KJM, see w7x_ref_66.

Figure 2.4.: Setup of the standard ECE system at W7-X. Black lines around the shown magnetic flux surfaces indicate accurately the vessel cross section. The magnetic flux surfaces correspond to the VMEC equilibrium used for the analysis in chapter 5. Also, the angle of the line of sight through the plasma is accurate.

To calibrate the system, exactly the same optical system including the vacuum windows is used below the W7-X experiment in the calibration unit, which is explained in more detail in section 2.3.2. The only difference between the two optical systems is the small tilt of the last plane in-vessel mirror to achieve a sightline perpendicular to the magnetic flux surfaces to suppress Doppler shift contributions. Instead of the inner plasma vessel, the calibration unit has its sightline leading to a rotating mirror, which is surrounded by a highly microwave absorbent foam (ECCOSORB®) at room and liquid nitrogen temperature, see also section 2.3.2 and figure 2.6. A wire grid separates the incoming radiation in X mode and O mode, for which two separate but similar transmission lines exist. An oversized waveguide of approximately 23 m length, including two tapers, 11 mitre-bends and one polarization tuner transmits the radiation from either the calibration unit or the plasma to the detection system outside the torus hall, allowing easy access during operation. The overall loss of this transmission line is 13.3 dB. The polarization tuner allows to adopt the mode to the radiometer input. Cross polarization coupling of the transmission line is on the order of 1 % to 2 % [37] as has been measured with a polarized signal source. A calibrated noise source can be selected by a waveguide switch instead of the transmission line, which allows to calibrate with a higher signal to noise ratio at the expense of not taking the influence of components in front of the noise source switch into account. A Bragg reflection notch filter with at least 55 dB damping within (140.0 ± 0.5) GHz and approximately 5.3 dB insertion loss outside was used to block non-absorbed 140 GHz ECRH radiation [51]. Afterwards, the signal is down-converted to 4 GHz to 40 GHz via a broadband low noise mixer fed by a phase locked loop (PLL) stabilized local oscillator (LO) at 122.06 GHz [52]. Using several staged amplifiers and power dividers the signal is split into 32 channels, which subsequently are band-pass filtered with center frequencies between 4.4 GHz and 39.6 GHz and bandwidths of 0.25 GHz $\leq \Delta f \leq$ 1.4 GHz, with a resolution in real space between 0.5 cm and 1.5 cm [50]. Then, the signal passes through the detection diodes and fixed postdetection amplifiers. Highly linear amplifiers with a variable gain and an adjustable DC offset allow to choose a reasonable signal amplitude for each plasma discharge, therefore making maximum use of the analog digital converter (ADC) range. Subsequently, a low-pass filter with a 3 dB point at 550 kHz determines the maximum temporal resolution and ensures that no aliasing occurs. Finally, a 16 bit ADC with a sampling rate of up to 2 MS s^{-1} is used, before the data are stored in the central W7-X database by the data acquisition (DAQ).
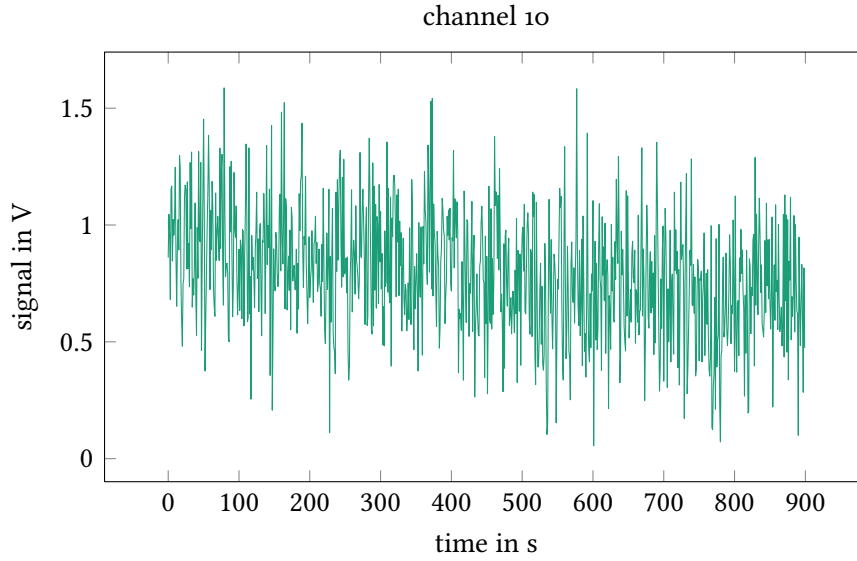
Figure 2.5.: Example of the drifts of an ECE channel during a few minutes. Channel 10 has a signal difference between the two reference temperatures of around 4.7 mV. Only every 1000[th] point is shown.

2.3.2. Calibration unit

Many ECE radiometers that are absolutely calibrated use a rotating blade to switch between two reference temperatures [53, 54]. Another method to switch between the reference temperatures is given by a rotating mirror, as described by Hartfuß et al.[36], and is sketched in figure 2.6. The advantage of these two methods in contrast to just recording data at one reference temperature for several minutes and then at another reference temperature for several minutes (both without using a chopper) is the much lower measurement signal drift sensitivity, which, if the drifts are not strictly linear, would change the ratio of the bit signal corresponding to the reference temperatures over time. The rotating mirror and the subsequently applied conditional averaging act as a bandpass which suppresses drifts on timescales larger than a rotation period, as the drifts correspond to a low frequency contribution. As drifts on timescales above 10 s are not negligible at W7-X (see figure 2.5) despite having the electronics in a temperature controlled rack, the rotating mirror method has been chosen for W7-X as the temperature control allows only for temperature stability of the

electronics on the order of 1 K to 2 K. The advantage of the rotating mirror over a rotating chopper lies in the better symmetry of the intermediate temperatures that are measured when radiation from multiple radiation sources at different temperatures is collected. Drifts on the magnitude observed here pose a considerable problem for long term plasma operation. Either further measures to suppress drifts have to be taken, or a regularly repeated offset determination within a discharge has to be performed. For W7-X, it is planned to repeatedly close the shutter in the planned 30 min plasma discharges to correct the offset.
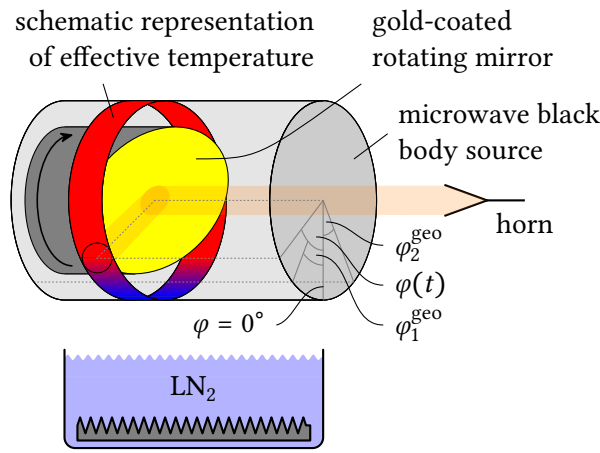
The calibration unit of the ECE at W7-X contains a gold-coated brass mirror rotating steadily with approximately 3.6 Hz, see figure 2.6 for a schematic illustration. Arranged cylindrically around the mirror, a microwave absorber guarantees a black body emitter at room temperature T_{RT}, which is kept at (294.45 ± 3.50) K in the torus hall. However, a small part at the lower side of this cylinder is cut to allow the observation of a stainless steel container

Figure 2.6.: Schematic drawing of the ECE calibration unit at W7-X.

thermally insulated by styrofoam. The stainless steel container's inner wall is lined with a pyramidical microwave absorber. The optical system between the rotating mirror part of the calibration unit and the microwave antenna, which is identical to the invessel optical system, is not shown. Light orange indicates the line of sight of the microwave antenna. The microwave antenna is characterized by a Gaussian beam. The beam is reflected at the gold-coated mirror and finally »sees« either room or liquid nitrogen temperature, the latter being produced by a liquid nitrogen tank underneath the rotating mirror. The cold reference temperature is not directly given by the temperature of the liquid nitrogen, $T_{LN_2} = (77.2 \pm 0.5)$ K, as water vapour (assumed to be at $T_{H_2O} = (280 \pm 10)$ K with an uniform emissivity of $0.01 < \varepsilon_{H_2O} < 0.03$) accu-

mulates above the liquid nitrogen reservoir. Moreover, the temperature of the mirror needs to be taken into account (corresponding to room temperature described above, with an emissivity $0.01 < \varepsilon_{\mathrm{mirror}} < 0.03$). This leads to an effective temperature difference between the hot and the cold source of about 208 K and an associated uncertainty on the order of 4 K that is used in further calculations. Due to the finite size of the beam, the effectively measured temperatures are smeared out at the hot/cold edges, as different parts of the beam »see« different temperatures. Details about the effective temperature estimation are given in section 4.2.1.

As stated beforehand, the analysis of the calibration process allows to obtain radiation temperature spectra (cf. section 2.1). W7-X is planned to have most of its standard analyses done in a Bayesian way in the long term. Therefore, the next chapter introduces the mathematical and philosophical concepts of Bayesian analysis.

3
Bayesian modelling

The actual science of logic is conversant at present only with things either certain, impossible, or entirely doubtful, none of which (fortunately) we have to reason on. Therefore the true logic for this world is the calculus of probabilities, which takes account of the magnitude of the probability which is, or ought to be, in a reasonable man's mind.

(James C. Maxwell)

Most analyses in science want to test how well a model describes some observation. For physics that typically means that a forward model (or a set thereof) predicting the observation needs to be constructed. Using Bayes' formula [55, 56],

$$\underbrace{P(F|D)}_{\text{posterior}} = \frac{\overbrace{P(D|F)}^{\text{likelihood}}\ \overbrace{P(F)}^{\text{prior}}}{\underbrace{P(D)}_{\text{evidence}}}, \tag{3.1}$$

originating from the product rule of probability theory allows to infer the desiderata (i. e. the sought-after quantities). In equation (3.1) $P(\cdot)$ represents a probability, F the free parameters to be inferred and D the data. A priori knowledge about the free parameters is encoded in the *prior* probability distributions. The probability of the measured data given the free parameters is called the *likelihood*. The normalization factor in the denominator, often called *evidence*, is important for model comparison. While calculable via marginalisation,

$$P(D) = \int P(D|F)P(F)\,\mathrm{d}F, \tag{3.2}$$

this integral is in general difficult to compute. However, there are methods available that can reconstruct the posterior without the explicit need to calculate the evidence, for example with the MCMC algorithm described in section 3.3. The *posterior* is the term we want to calculate, as it yields the probability distribution of the sought-after free parameters given the data. As the evidence does not change the shape of the posterior for a given model, one can neglect that term if one is not interested in comparing models explicitly. For comparing models, however, the evidence is required, and it can be shown that this term penalizes complexity (that is, applying Occam's razor), as the probability mass of more complex models is spread over a larger hypervolume. If no forward model is used the difficulty to obtain consistent results if multiple observations are available is severly increased. An example from plasma physics: Predicting the measurements of X-ray imaging crystal spectroscopy and ECE from a common T_e profile is straightforward, while a routine that calculates T_e separately for both diagnostics will hardly get consistent T_e profiles as a result.

The philosophy underlying the use of Bayesian probability theory in this thesis (and in general in Minerva, see section 3.4) does refute the common notion of *measurement error* in favour of putting the uncertainties on the *model*. Thus, observations in the models in this thesis are assumed to be *exact*. This has no directly obvious consequences, as the uncertainty distributions are usually assumed to be symmetric. However, transfering the uncertainties from the observations to the predictions allows to treat the uncertainties as any other part of the model, practically acting as a »uncertainty submodel«. In particular, one can multiply the predicted uncertainties with a free parameter which can be optimized by maximizing the probability density[1].

A specific and very powerful technique used throughout the thesis are so-called *virtual observations*. They are useful for imposing certain constraints on the prior and work in the following way: First, the equation describing the prior constraint is reorganized such that the terms that can be calculated from the forward model are on the left hand side. The result that one would expect to get if the equation is exactly fulfilled corresponds then to the right hand side of the equation. However, now an uncertainty on the prediction is introduced. Thus, these prior constraints act like constraints imposed by observations from

---

[1] Note that this works because larger uncertainties correspond to a higher complexity, as more models can predict the observation well. This corresponds to a spreading of the probability mass over a larger hypervolume mentioned above.

some diagnostic. However, as they do not relate to data obtained by a real world measurement device, they are called »virtual«. Note that this implies that whatever equation is implemented this way is thus only *probabilistically* enforced.

An important part of a typical plasma model is the way the plasma profiles, for example the electron density profile, is described. Thus the next section will specify how the plasma profiles are described in this thesis.

3.1. Gaussian processes

One of the typical problems with descriptions of physics problems is that the exact analytic formula to represent, for example, an electron temperature profile is not known. One can choose a simple parameterization, for example the classical approach for the electron temperature profile,

$$T_e(\rho) = T_{e,0}(1 - \rho^\beta), \tag{3.3}$$

wherein ρ is the effective radius, β a parameter determining the shape and $T_{e,0}$ is the electron temperature at the core. However, the choice of a specific parameterization can put severe constraints on the posterior that might hinder proper physical interpretation, here for example of the profile shape. With these parameterized methods, one can optimize their corresponding parameters. Exemplarily, choosing equation (3.3) for the electron temperature profile implicitly puts heavy assumptions on the *transport* processes in the plasma, as the parameterization influences the gradients that are themselves a result of transport occuring in the plasma. Therefore, it would be good to have a *nonparametric* way to constrain the posterior less.

A nice method to achieve this is given by Gaussian processes [57]. In contrast to the parameterized methods one does *not* optimize in the parameter space directly, but rather in a *function* space determined by *hyperparameters*. As a consequence, a far greater freedom in function shapes is allowed. A Gaussian process is defined as

$$f(\mathbf{x}) \sim \mathcal{GP}\big(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\big), \tag{3.4}$$

wherein $\mathbf{x}$ corresponds to the finite number of support points[2], $m(\mathbf{x})$ to the

---

[2]In the case of an electron temperature profile these are e. g. the effective radius positions at which the Gaussian process is evaluated.

mean function and $k(\mathbf{x}, \mathbf{x}')$ to the covariance function[3], specifying the covariance of random variables. Thus, a stochastic process is a Gaussian process if the values $f(\mathbf{x})$ are elements of a multivariate normally distributed vector. The hyperparameters only determine the mean and covariance of the multivariate normal, not the realisation the function takes at a specific point, which in a practical setting has to be determined as well. The covariance function corresponds to the so-called kernel. An often used kernel is the squared exponential kernel [58]. However, such a stationary (i. e. invariant under translations) kernel has the disadvantage that the covariance, that is the assumed a priori smoothness, is the same for all points. Thus, a stationary kernel will struggle to describe, for example, a electron density profile that is flat throughout the core region but contains notable gradients towards the plasma edge, as different smoothnesses would be necessary. A solution to this is to use a kernel that depends on the position. Consequently, the Gaussian processes used throughout this thesis use the nonstationary kernel by Paciorek [59–61][4] and a correlation length defined by a hyperbolic tangent length scale function [62],

$$
\ell(x) = \frac{\ell_1 + \ell_2}{2} - \frac{\ell_1 - \ell_2}{2} \tanh\left(\frac{x - x_0}{x_w}\right), \tag{3.5}
$$

in which ℓ_1 is the core saturation value of the correlation length, ℓ_1 is the saturation value of the correlation length at the edge, x_0 the point at which the two length scales coalesce and x_w the characteristic width of the coalescent region. This allows, for example, to have a flat electron density profile in the core while still getting a steep gradient at the edge required to adequately describe observed plasma profiles. Even though a hyperbolic tangent length scale function is used is the profile shape *not* locked into a hyperbolic tangent.

An example to foster intuition on how the points correlation in a Gaussian process influences the shape is depicted in figure 3.1 with samples drawn from a Gaussian process prior. For example, the points 1. and 2. on one of the prior samples show similar T_e values, indicating that the points are, probably, highly correlated. The T_e values between point 1. and 3. show a larger T_e difference, indicating that the correlation between the points is lower as compared to

---

[3]Here, $\mathbf{x}'$ is used to indicate that the covariance function describes the correlations of each component of $\mathbf{x}$ with each other component of $\mathbf{x}$.

[4]Note that Gibbs derived almost the same kernel in his PhD thesis in 1997, only differing by a factor of two in the exponential. The generalized kernel by Paciorek is based on Higdon et al. and was used here.
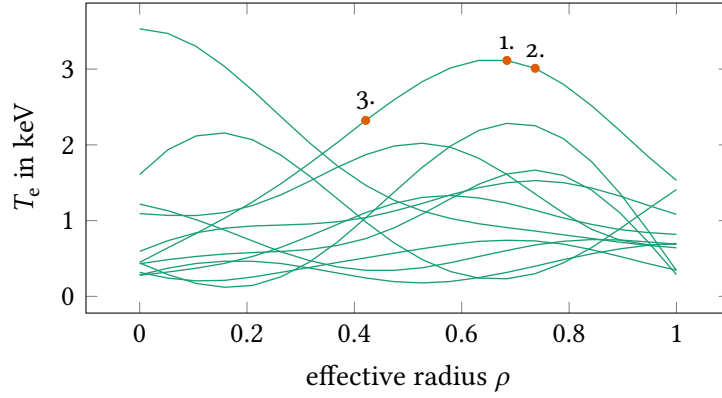
Figure 3.1.: Shown are some examples for profile shapes that can be obtained by sampling from a Gaussian process prior with fixed hyperparameters (cf. the cyan T_e node in figure 5.1). Corresponds to the T_e Gaussian process used for the profile inversion in chapter 5.

points 1. and 2. Be aware however that the samples are drawn probabilistically, therefore drawing conclusions from a single curve has to be taken cum grano salis.

With the knowledge about Gaussian processes the models used throughout this thesis can be constructed. In what follows the techniques that will be used to analyse the models will be introduced.

3.2. Maximum a posteriori

A typical problem in model analysis is the determination of the parameter space point with the highest probability, that is the maximum a posteriori (MAP),

$$F_{\text{MAP}} = \arg\max_F [P(F|D)]. \tag{3.6}$$

Note that usually not the probability density directly, but rather the so-called *joint* defined by $P(D|F)P(F)$, which is proportional to the probability density if no hyperparameters are changed, is calculated. There are various methods available to find that spot, for example gradient based search algorithms [63]. Another class of algorithms use patterns, amongst which the Hooke and Jeeves algorithm [64] is well established. This particular algorithm was used extens-
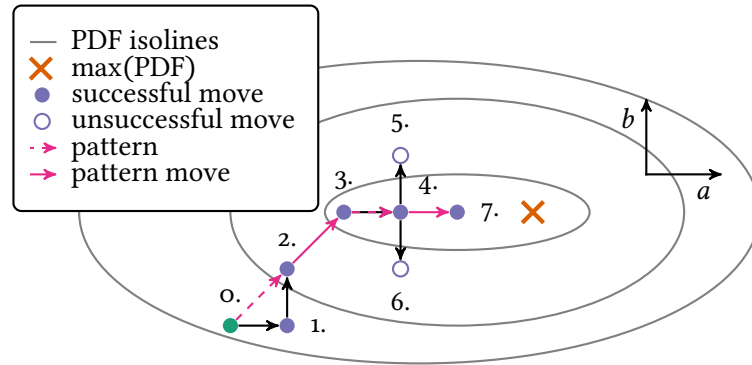
Figure 3.2.: Example of the modus operandi of an maximum a posteriori (MAP) in 2D based on the Hooke and Jeeves pattern search algorithm [64].

ively throughout this thesis. The way the algorithm works is shown schematic-ally for two parameters a and b in figure 3.2. The starting point of the analysis is at 0. Then, the algorithm takes a step along the axis of parameter a and checks whether this increases the joint probability of the model. If this is the case the reference point moves there. As it is indeed the case in the shown example, 1. is made the new reference point. Subsequently, the algorithm moves along the second parameter axis, b. Again the probability density increases, thus the pos-ition is moved to 2. Moves along all dimensions have hence been undertaken. The vector that combines the successful moves of this iteration corresponds to the found pattern. Applying the pattern to 2. moves the point to 3. Then the next iteration begins, probing a new position along the a axis which is, here, again improving the joint probability, thus moving to 4. Probing a new position along the b axis would move to 5., but the joint probability is lower there, thus the point is discarded. Consecutively, the step along the b axis is attempted in the other direction which would lead to 6. As the joint probability is again lower at this point it is rejected as well. Thus, the pattern found for this iteration leads to a move to point 7. Continuing these iterations[5] will yield a parameter space position nigh the maximum of the probability densitiy function (PDF). One criteria that can be used to determine when to end the chain is to define a maximum number of iterations in which the reference point does not move, and abort if this number is exceeded. It is worthy to note that the Hooke and

---

[5]This includes, for the analyses performed in this thesis, a halving of the stepsize in each iteration in which no move was found, until the stepsize falls below 10^{-5}.
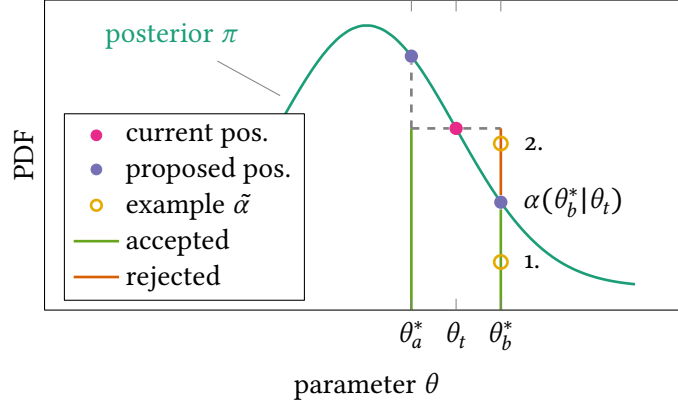
Figure 3.3.: Example of the modus operandi of an MCMC based on the Metropolis-Hastings algorithm.

Jeeves algorithm might get stuck in a local maximum. Even if not stuck in a local maximum, it is desirable to find not only the maximum of the posterior, but to reconstruct the *shape* of the posterior. If it is available, potential nontrivial dependencies between the parameters can be found. Also, one can see from the shape how stable the found MAP is. An MAP estimation yields a point estimate and is thus not able to reconstruct the posterior. A technique with which the shape of the posterior can be found is introduced in the next section.

3.3. Markov chain Monte Carlo

Markov chain Monte Carlo (MCMC) is a useful method to obtain the posterior shape. A Markov chain allows to make a probabilistic statement about the next step in the chain, wherein the knowledge of the full history does not yield an advantage over knowing only the current state, i. e. the chain has no memory. The goal of an MCMC analysis is often to sample from the posterior, which corresponds to the stationary distribution of the Markov chain. As an example, the Metropolis-Hastings algorithm [65, 66] is described, corresponding to the algorithm used throughout the thesis. Figure 3.3 serves as a visualization of some of the steps. First, a proposal distribution[6] q is used to obtain a candidate

---

[6]The proposal distribution is the distribution from which a sample determines where the next proposed parameter position will be.

with a new parameter space position, θ_{t+1}^*, from the current parameter space position, θ_t. This allows the computation of the acceptance probability of the new parameter space position, given by

$$\alpha(\theta_{t+1}^*|\theta_t) = \min\left(\frac{\pi(\theta_{t+1}^*)q(\theta_t|\theta_{t+1}^*)}{\pi(\theta_t)q(\theta_{t+1}^*|\theta_t)}, 1\right), \tag{3.7}$$

with $\pi(\cdot)$ the (unnormalized) posterior probability of the given position. Note that this implies that the evidence $P(D)$ (cf. equation (3.1)) does not need to be calculated explicitly, as only the ratio $\pi(\theta_{t+1}^*)/\pi(\theta_t)$ is needed – only the relative change of »height« is required. In figure 3.3 two examples are shown; one point jumps to a position with a higher posterior probability, θ_a^*, and another jumps to a position with lower posterior probability, θ_b^*. In the case of the former, the acceptance probability is 100 %, in the latter it is a bit over 50 %. Subsequently, a random number from a uniform distribution between 0 and 1 is drawn, here denoted by $\tilde{\alpha}$. This value is then compared to the acceptance probability, if $\tilde{\alpha} \leq \alpha$ the new position is accepted, else it is rejected and the current position is kept. To help to get an intuition for the modus operandi of an MCMC, example random numbers are drawn to compare the outcome for the examples in figure 3.3. The proposed jump to θ_a^* is accepted for any value of $\tilde{\alpha}$, as the posterior probability of θ_a^* is larger than the one at θ_t. For the other proposed point at θ_b^*, 1. corresponds again to an $\tilde{\alpha}$ below the posterior curve and thus θ_b^* would be accepted, while for 2. the posterior curve is below $\tilde{\alpha}$ and the jump is thus rejected. A notable simplification of the algorithm is achieved if the proposal distribution is symmetric as the q terms in equation (3.7) cancel then. Many proposal distributions do satisfy this criterion, but for example a proposal distribution with a fixed position generally does not. In this case, the q terms remove the influence of the proposal distribution, which otherwise would lead to an over- or underweighing of certain parameter space points of the posterior. The efficiency of the MCMC depends on the shape of the proposal distribution, although in principle any proposal distribution will work. Thus, adaptive changes of the proposal distribution are a way to automatically improve the efficiency of the MCMC [67]. Noteworthily, adapting the adaptation method itself might help in further speeding up convergence [68], which however has not been used in this thesis. However, to ensure that the posterior remains the stationary distribution that is converged to, the adaptive changes to the proposal distribution need to diminish.

A burn-in is the practice of running a specific number of MCMC iterations at the beginning of an MCMC and use the resulting point as a new start point, while discarding the results from the burn-in samples. Strictly speaking, a burn-in is not necessary given the MCMC is run long enough. However, it can be a way to get to a reasonably likely starting point, as one does not want to start in the tails of the probability distribution due to the limited number of samples that will be drawn during the MCMC. Do note that the MCMC in use in this thesis are essentially black box MCMCs, thus the only way to ensure convergence instead of just pseudo-convergence is by using sufficiently long runs [69]. In general, the optimal acceptance rate depends on the exact problem, but for a wide variety of models 0.234 is optimal[7].

As MCMC samples are typically autocorrelated, a thinning is done, respectively only every third point is kept within the MCMC runs of this thesis, which reduces the autocorrelation. The posterior can be obtained from the samples for example by using a kernel density estimate [72, 73] on the MCMC samples.

3.4. Minerva

Minerva [74] is a Java based, modular, general Bayesian modelling framework that can handle arbitrarily complex models and is used among several other large fusion experiments around the world, for example at the Joint European Torus (JET) [75] and the Mega-Ampere Spherical Tokamak (MAST) [76] and is the main inference framework at W7-X. It is designed to work for *all* kind of inference problems and is not limited to physics or fusion problems. In contrast to many other frameworks the methods used to infer the desiderata are kept completely orthogonal to the modelling part based on graphical models. This means that the program parts handling the inference are decoupled from the program parts describing the models.

Graphical models are a powerful tool to describe the conditional dependency structure of a probabilistic model [77]. A Minerva graphical model is a Bayesian

---

[7]Roberts et al. [70] showed that running an MCMC with an acceptance rate of 0.234 is optimal, if one uses a multivariate Gaussian proposal distribution and a high-dimensional, independent and identically distributed target (posterior) distribution. However, if the convergence rates of some of the posterior dimensions deviate significantly when compared to other dimensions, 0.234 is usually no longer optimal and can be much smaller [71]. This applies to most of the MCMCs ran during the preparation of this thesis.

network, more precisely, a directed acyclic graph. Examples of Minerva graphs can be found in figures 4.1 and 5.1. It consists of deterministic and probabilistic nodes, which are connected via edges. The graphical model encodes the whole joint probability of the free parameters and the data. The joint probability can be calculated by adding up the logarithms of the probabilities of all probabilistic nodes conditioned on their parent, i. e. the nodes describing the priors (colored cyan in the aforementioned figures) and the nodes describing the observations and predictions (colored grey in the aforementioned figures). Based on these models, sample predictions that may be used for example for neural net training sets [78] can be generated, as well as inversions based on different techniques, such as MAP and MCMC methods. Scientific traceability of the models can be ensured easily by serializing the model structure to an `.xml` file.

Some basic concepts used for the implementation at W7-X that help to ease the handling of multiple models and their analysis are described henceforth. As these concepts are relatively generic they probably can be used at other experiments as well.

3.4.1. Datasources

A Minerva datasource is a node in a graphical model. Datasources should be the only experiment specific nodes in the whole graphical model to allow an easy adaptation of the model to another experiment by exchanging the datasources. More specifically, a datasource loads the data from an immutable database and does some preprocessing, if necessary. Furthermore, it provides metadata like lines of sight or frequencies. There have been many attempts to provide metadata in a generalized manner. The problem with that lies in the plethora of *types* of metadata that is required for different usecases. An example from W7-X: The ECE datasource contains the metadata for its calibration unit (cf. section 2.3.2), e. g. its diameter. However, not all diagnostics, let alone all datasources require »calibration unit« type metadata. This makes it difficult to generalize without burdening datasources that have no need for this type of metadata in the first place. Thus, all the metadata provided by datasources is handcrafted and specifically tailored for each datasource. This is time consuming, but still the most pragmatic way the author knows to ensure a satisfying user experience. An example of how a datasource is used in the construction of a graphical model can be seen in code 3.1.

Code 3.1: Example code from the plasma profile inversion model, see chapter 5. Code corresponds to [79] with small adjustments to improve readability.

```
239  // the graphical model
240  GraphicalModel g = new GraphicalModel("EceInterfModel");

242  // this node handles the time to be analyzed
243  W7xProgramTimeToNanos nanosecond = new W7xProgramTimeToNanos(g
         , "time");
244  nanosecond.setTime(201808023016002, 4.45);
                              ⋮
276  // provides ECE data and metadata
277  W7xEceQmeDataSource ds = new W7xEceQmeDataSource(g, "ds");
278  ds.setNanosecond(nanosecond);
                              ⋮
640  // connection to the TRAVIS webservice
641  Travis travis = new Travis(g, "travis");
                              ⋮
644  travis.setFrequencies(new Source(ds, "getFrequenciesGHz"));
```

In line 240 the model to be constructed is instantiated, lines 243 and 244 instantiate the node describing the discharge time and set the time to be analyzed. Exemplarily, the second scenario of the 16[th] experiment on the 23[rd] August 2018, 4.45 s after the heating was switched on, is selected. Line 277 instantiates an ECE datasource for W7-X. Subsequently, line 278 connects the W7-X time node to the ECE datasource node. Similarly, line 641 instantiates a TRAVIS node, while line 644 connects the ECE datasource node to the TRAVIS webservice wrapped in the aforementioned node. This is done to ensure that TRAVIS utilizes the frequencies, line of sight and further diagnostic specific settings used during experiments for its calculations. Note that the TRAVIS node is not W7-X specific: The model could be easily ported to e. g. JET, if the ECE datasource and the node indicating the time to be analyzed get substituted with their corresponding counterparts.

In the course of this thesis datasources for the ECRH, ECE, X-ray imaging crystal spectroscopy, snifferprobe, single channel dispersion interferometry, NBI and Langmuir probe systems have been written for W7-X and made available within the Minerva framework in preparation of a combined Bayesian analysis on a large number of diagnostics.

3.4.2. Evaluation strategies

The evaluation strategy in Minerva is generally the following: First get reasonable estimates for the starting parameters via the datasource. Consecutively, the analysis method has to be selected. If the model is linear, the typical analysis method chosen is the linear Gaussian inversion explained in more detail below. For nonlinear problems, the standard procedure continues with an MAP inversion[8], which, if not trapped in a local maximum will generally yield a result where relevant amounts of probability mass can be found. Consequently, the MAP result is used as a starting point for the MCMC inversion with a sufficient burn-in and an adaptive Metropolis adapter [67] that can be deactivated once the MCMC chain is stable.

The linear Gaussian inversion [80–82] (for the notation and full derivation see [82]) is a very powerful technique to determine the full posterior distribution, which can be applied if the model is linear and both the likelihood and the priors can be expressed as multivariate normal distributions,

$$P(\mathbf{D}|\boldsymbol{\mu}) = \mathcal{N}(\mathbf{D}; \boldsymbol{M}\boldsymbol{\mu} + \mathbf{C}, \boldsymbol{\sigma}_{\mathrm{D}}) \tag{3.8}$$

$$P(\boldsymbol{\mu}) = \mathcal{N}(\boldsymbol{\mu}; \boldsymbol{\mu}_{\mathrm{p}}, \boldsymbol{\sigma}_{\mathrm{p}}). \tag{3.9}$$

with $\mathbf{D}$ the data vector, $\boldsymbol{\mu}$ the free parameter vector, $\boldsymbol{\sigma}_{\mathrm{D}}$ some known variance, the mean of the likelihood a linear combination given by $\boldsymbol{M}\boldsymbol{\mu} + \mathbf{C}$ wherein $\boldsymbol{M}$ is the response matrix and $\mathbf{C}$ the constants vector. Similarly, the prior is also given by a multivariate Gaussian distribution, with mean $\boldsymbol{\mu}_{\mathrm{p}}$ and variance $\boldsymbol{\sigma}_{\mathrm{p}}$. Note that »;« stands for »parameterized by«. This allows to determine the posterior exactly, which is also given by a multivariate normal with a mean and covariance of

$$\boldsymbol{\sigma} = (\boldsymbol{M}^{\mathrm{T}}\boldsymbol{\sigma}_{\mathrm{D}}^{-1}\boldsymbol{M} + \boldsymbol{\sigma}_{\mathrm{p}}^{-1})^{-1}, \tag{3.10}$$

$$\boldsymbol{\mu}_0 = \boldsymbol{\sigma}[\boldsymbol{M}^{\mathrm{T}}\boldsymbol{\sigma}_{\mathrm{D}}^{-1}(\mathbf{D} - \mathbf{C}) + \boldsymbol{\sigma}_{\mathrm{p}}^{-1}\boldsymbol{\mu}_{\mathrm{p}}] \tag{3.11}$$

Note that throughout this thesis matrices are indicated by bold italic symbols, whereas vectors are indicated by bold upright symbols.

A particularly powerful approach here lies in combining the linear Gaussian inversion with the Hooke and Jeeves pattern search algorithm (cf. section 3.2)

---

[8]For example using the Hooke and Jeeves pattern search algorithm mentioned above, optimizing with respect to the joint probability density.

allowing to optimize with respect to the *posterior probability*, not just with respect to the product from likelihood and prior. This approach is used for example in the calculation of the initial guesses for the electron temperature and density profiles (see section 5.1).

4

ECE diagnostic calibration[1]

A common way to absolutely calibrate the radiometers is the hot/cold calibration method, which uses two reference temperatures to determine the sensitivities respectively the calibration factors [36], although often only a relative calibration with respect to a Michelson interferometer or a Thomson scattering diagnostic is done [53]. As described in section 2.3.2 the reference temperatures at W7-X are given by a microwave absorber at room temperature and a cold source at liquid nitrogen temperature. In principle, a hot source could also be used, as planned for ITER [83]. The signal chain from the horn to the DAQ system contains many components. For a complete modelling of the ECE, one would like to have individual uncertainties for each component, but the large number of components involved effectively prohibit this approach. This makes it challenging to get appropriate uncertainties. Realistic uncertainties are especially important for modelling of multiple diagnostics, as too small or too big uncertainties will artificially bias the result. Good estimates for the overall uncertainties of the calibration are obtained by inverting the forward model of the calibration process described in this chapter (cf. figure 4.1 for the corresponding graphical model). Easy automation of the whole calibration procedure is possible as the implementation is capable of finding reasonable defaults where necessary.

This chapter will begin by giving an overview of the available models, and their limitations in section 4.1. Thereafter, the formulæ constituting the models are layed out in section 4.2, first for the specific case at W7-X, then generalised for arbitrary calibration units. Then, the results are discussed in section 4.3 and conclusions are drawn in section 4.4.

---

[1]Most of this chapter has been published in [41]; the format has been adapted to the thesis.

4.1. Minerva implementation

To obtain the physical quantities that are of interest for the evaluation of an ECE calibration, namely the calibration factor of each individual channel (that is the inverse sensitivity) and the effective beam width, the calibration procedure itself is modeled. To evaluate multiple ECE channels in a consistent way, a forward model predicting the different channel sensitivities has to be used. See also section 3.4 for the required Minerva background.

For comparison of different modeling approaches, different descriptions of the calibration process were implemented. The following properties of the model have been combined, creating a total of eight structurally different models: i) Evaluation of individual channels respectively multiple channels simultaneously (indicated from here on with the keywords »single« respectively »multi«), ii) with and without the use of a channel specific scaling factor α_i to scale the variance of the observed conditionally averaged signal (cf. the footnote in chapter 3) and iii) by having a channel specific beam width or a beam width fixed for all channels by a scaling following

$$w(f) = w(140\,\text{GHz})\sqrt{140\,\text{GHz}/f}, \tag{4.1}$$

as expected from broadband Gauss telescope optics. Henceforth, models using a channel specific beam width are denoted »shared«, while the models using channel specific beam widths are denoted »individual«.

In the case of the datasource (cf. section 3.4.1) used for the evaluations presented here, the preprocessing consists of the following steps: First, the chopper channel data are used to automatically identify the individual mirror rotation periods. Subsequently, the mean of each period is subtracted to avoid an overestimation of the uncertainties. This also removes the explicit dependence on the absolute values of the reference temperatures (cf. section 4.2.1 and appendix A). Then, each period is binned to avoid influences by variations of the motor speed. Finally, a coherent averaging is done to get a signal with a better signal to noise ratio.

Due to the small signal to noise ratio, many rotations of the rotating mirror need to be acquired to be able to extract enough information for meaningful results. The preprocessing is necessary as working with the full data is unpractical due to the sheer volume of raw data. In the case studied in this chapter, this decreases the number of effective temperatures to predict from about 10^8 to

less than 10^4 for the full combined model. An example of such a preprocessed signal for a channel with good sensitivity can be seen in figure 4.4.

The advantage of evaluating all channels simultaneously is the consistency gained for channel independent parameters, namely of the angles φ_1^{geo} and φ_2^{geo} that determine where the radiometer starts respectively stops to collect radiation mainly from the hot source (cf. section 2.3.2), and, depending on the model, the beam width. The disadvantage that comes with evaluating multiple channels simultaneously lies in the *curse of dimensionality*, i. e. the evaluation time increases notably. This »curse« can be understood in the following way: Consider that the posterior is normalized, i. e. that the integral over its PDF is exactly 1. Consequently, as its dimensionality increases the probability density at any given parameter point shrinks. Therefore the information contributed by one point to an accurate reconstruction of the posterior, e. g. via an MCMC, diminishes accordingly, effectively requiring much more evaluations of the PDF to achieve the same level of accuracy compared to a lower dimensional case.

Exemplarily, the graphical model for the multi-channel case with variance scaling factors α_i and channel specific beam widths w_i is shown in a simplified way in figure 4.1. The user only has to set which channel contains the chopper signal and the calibration segments (i. e. the time windows that contain the data), as the datasource then automatically fetches the radiometer settings and the raw data from the W7-X archive database (ArchiveDB) and preprocesses the data, including the merging of the calibration segments, if multiple segments are given. T_{LN_2} represents the liquid nitrogen temperature, $T_{\text{H}_2\text{O}}$ the temperature of the water vapour that forms above the liquid nitrogen with $\varepsilon_{\text{H}_2\text{O}}$ the emissivity of the water vapour, T_{RT} the room temperature and $\varepsilon_{\text{Mirror}}$ the emissivity of the mirror. ΔT is the effective temperature difference between the hot and the cold source. As stated above, φ_1^{geo} and φ_2^{geo} denote the mirror angles between which the radiation collected by the radiometer originates predominantly from the hot source, cf. figure 2.6. Together with the Gaussian beam width of the microwave antenna characteristic w_i, the number of bins and the effective hot and cold temperatures the predicted effective temperatures $T_{\text{eff},i}(\varphi)$ can be calculated. The calibration sensitivity $\tilde{\eta}_i$ as calculated from the bit change Δs_i and the temperature difference ΔT_{eff} can be multiplied with $T_{\text{eff},i}(\varphi)$ to obtain the prediction $f_i^{\text{pred}}(\varphi)$ of the measurement signal for channel i. For the uncertainty of the prediction the variance as estimated from the conditional averaging, scaled by a factor α_i, is used. This variance scaling factor allows to find the uncertainties matching the predictive capabilities of the model, as too small uncertainties are
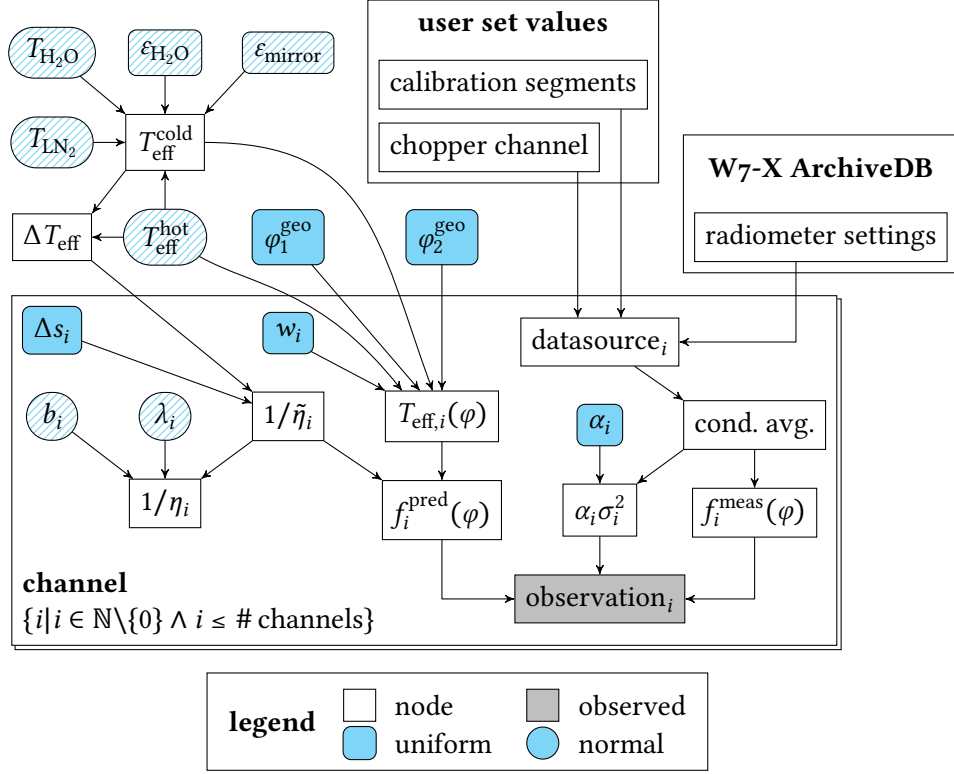
Figure 4.1.: The simplified Minerva graph showing the model dependencies. Free parameters have a cyan background and their shape, as specified in the legend, corresponds to the prior distribution used. Note that the nodes with a hatched background sample completely randomly from their distribution, thus propagating the corresponding uncertainties without increasing dimensionality by parameters which do not influence the prediction, cf. appendix A.

penalized by increasing the relative distance between prediction and observation, while too large uncertainties correspond to a higher complexity and are thus penalized by the implicit penalty for complexity in Bayes' formula.

The conditionally averaged and binned signal $f_i^{\mathrm{meas}}(\varphi)$ in bit constitutes the observation. In combination with the bit to Volt conversion factor b_i, and the factor taking the measured differences between the invessel and calibration optics into account, λ_i, the calibration factor $1/\eta_i$ can be calculated.

The single channel evaluation will be compared with the multi-channel evaluation in section 4.3.1. The evaluation strategy is described in section 3.4.2. This

allows for a realistic estimation of the uncertainties of the drop of the measured signal associated with the decrease in radiation temperature as produced by the switch from room temperature to the cold temperature. It has to be noted that the measured signal stored in the database is the raw bit signal, which is also what the model predicts. The trace of the logarithm of the joint (i. e. prior times likelihood) probability density, which is a measure for how probable a specific realisation is, of an MCMC of the full model, and example plots of each the calibration and scaling factor traces are shown in figure 4.2. All MCMCs have been run with a burn-in of at least 10^6 iterations, until the MCMC traces of the free parameters did not show notable drifts of the running average of the logarithm of the joint probability density. For all MCMCs ran so far, convergence was observed after at most 8×10^5 MCMC iterations, the single channel evaluations typically needed less than 10^4 iterations. This evaluation is fully automatized and writes the results to the central W7-X database, making automatic calibration overnight feasible. However, in practice the lack of liquid nitrogen availability in the torus hall prevents regular overnight calibrations.

4.1.1. Limitations

Practical limitations are given for example by the number of bins as well as the number of channels. If the sampling rate notably exceeds 1 kHz the evaluation slows down considerably as well, as the data that have to be loaded for the pre-processing increases accordingly. A practical number of bins can be determined automatically, which gives a value close to the average number of data points per mirror rotation. As there are computationally expensive steps involved for each bin, increasing the number of bins also leads to an increase of required computation time. The full 98D model (made up of φ_1^{geo}, φ_2^{geo} and 32 channels, each with free parameters for the beam width, variance scaling and bit dip) includes all 32 channels and the scaling of the variance, and takes roughly 230 hours with a Intel® Xeon® central processing unit (CPU) E5-2660 v4 at 2 GHz on a virtualized linux server for a roughly one hour long calibration. By evaluating each channel independently the required CPU time is reduced to about 21 hours, such that with full parallelization the evaluation time can go down to around 40 minutes. This simplified model comes at the price of a generally smaller consistency and larger uncertainties for the weaker channels. However, the differences are negligible for reasonably sensitive channels (calibration factor changes are typically below 1 %). For less sensitive channels the difference can
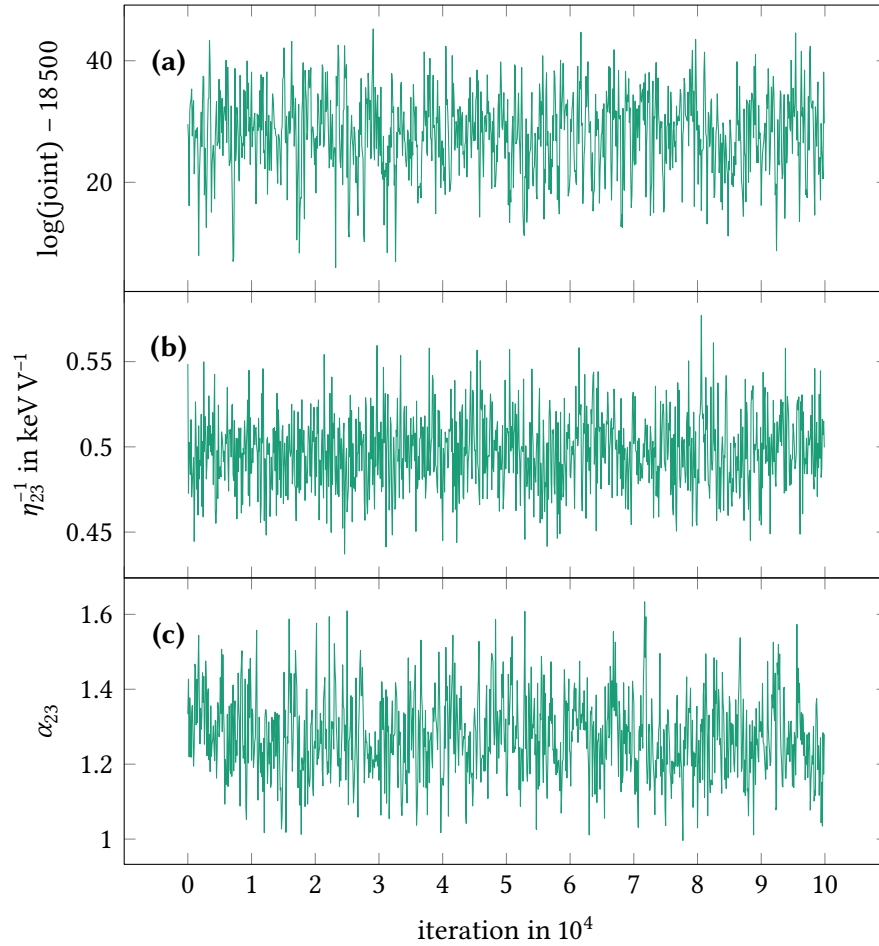
Figure 4.2.: (a) shows the logarithm of the joint (i. e. prior times likelihood) of an MCMC of the full model with variance scaling factors. **(b)** and **(c)** show the MCMC traces of the calibration factor η_{23}^{-1} and the variance scaling factor α_{23}. The values do not show jumps to a notably different phase space part, indicating that the chain converged. Only every 100[th] datapoint is shown to keep the plot size reasonable.

reach about 10 %. This is due to the more sensitive channels keeping the geometrical factors (φ_1^{geo} and φ_2^{geo}) more or less fixed, such that the impact on channels where the geometrical information is more concealed in noise profit the most. The single channel evaluation routine provides a pragmatic approach to obtain calibration factors if time requirements hinder the full model use.

4.2. Model formalism

This section introduces the required formulæ, first for W7-X and then for a generalized model with arbitrary geometries and effective temperatures.

4.2.1. Model for W7-X

To calculate the temperature difference between the room temperature and the effective cold temperature, one needs to take the influence of the water vapour emissivity and the mirror emissivity into account. The effective temperature after the radiation passes through the water vapour forming above the liquid nitrogen is calculated via

$$T_{\text{eff}}^{\text{vapour}} = T_{\text{LN}_2} + \varepsilon_{\text{H}_2\text{O}}(T_{\text{H}_2\text{O}} - T_{\text{LN}_2}), \tag{4.2}$$

where $\varepsilon_{\text{H}_2\text{O}}$ is the emissivity of the water vapour. The effective cold temperature after the mirror is given by

$$T_{\text{eff}}^{\text{cold}} = T_{\text{eff}}^{\text{vapour}} + \varepsilon_{\text{mirror}}(T_{\text{eff}}^{\text{hot}} - T_{\text{eff}}^{\text{vapour}}). \tag{4.3}$$

The emissivity of the mirror only adds to the effective cold temperature and not the effective hot temperature, as the mirror is already at the same temperature as the effective hot temperature. In case the hot reference temperature would be at a temperature different from the mirror temperature it would have to be taken into account there as well. Note that this works because the signals are recorded in the frequency range where the Jeans law approximation holds. This allows the calculation of the channel specific calibration sensitivity (respectively, the sensitivity η_i scaled with the attenuation and postdetection amplification as chosen for this channel during the calibration)

$$\tilde{\eta}_i = \frac{\Delta s_i}{\Delta T_{\text{eff}}} = \frac{\Delta s_i}{T_{\text{eff}}^{\text{hot}} - T_{\text{eff}}^{\text{cold}}}, \tag{4.4}$$

wherein Δs_i represents the signal change in bits caused by the temperature difference ΔT_{eff}. The index i denotes the channel specific parameters. Moreover, the channel specific expected effective temperature, depending on the mirror position, is given by

$$T_{\text{eff},i}(\varphi) = \frac{\int \hat{T}_{\text{eff}}(\hat{\varphi}) g_\varphi(w_i, \hat{\varphi}, \varphi) \, \mathrm{d}\hat{\varphi}}{\int g_\varphi(w_i, \hat{\varphi}, \varphi) \, \mathrm{d}\hat{\varphi}}, \tag{4.5}$$

where $\hat{T}_{\text{eff}}$ is the effective temperature at a given mirror position without taking the finite width of the Gaussian beam, defined by the horn characteristic, into account,

$$\hat{T}_{\text{eff}}(\hat{\varphi}) = \begin{cases} T_{\text{eff}}^{\text{cold}} & \text{for } \hat{\varphi} < \varphi_1^{\text{geo}} \\ T_{\text{eff}}^{\text{hot}} & \text{for } \varphi_1^{\text{geo}} \leq \hat{\varphi} \leq \varphi_2^{\text{geo}} \\ T_{\text{eff}}^{\text{cold}} & \text{for } \hat{\varphi} > \varphi_2^{\text{geo}} \end{cases}, \tag{4.6}$$

and g_φ the weight of each $\hat{T}_{\text{eff}}$ assuming a perfect Gaussian beam horn characteristic.

This in turn allows to predict the measured bit signal in dependence of the mirror angle,

$$f_i^{\text{pred}}(\varphi) = T_{\text{eff},i}(\varphi) \tilde{\eta}_i. \tag{4.7}$$

In practice, it can be useful to subtract the mean of the predicted signal to avoid dependencies on the effective temperatures where not necessary, see appendix A.

Scaling the variance σ_i^2 with a channel specific parameter α_i on the observation node allows for a realistic estimation of the uncertainties by taking the predictive capability of the model into account, see also chapter 3.

The calibration factor is given by

$$\frac{1}{\eta_i} = \frac{\lambda_i b_i G_i}{\tilde{\eta}_i} 10^{-(R_{\text{RF}} + R_{\text{IF}})/10}. \tag{4.8}$$

This takes the following quantities into account: The measured differences between the invessel and calibration optics, λ_i, the measured bit to volt conversion factor, b_i, the post detection amplification chosen during calibration, G_i, the setting of the radiofrequency waveguide attenuator right in front of

the radiometer in decibel, R_{RF}, and the setting of the attenuator at the intermediate frequency device in decibel, R_{IF}. For calibrating plasma measurements one needs to rescale this factor with the appropriate gains and attenuator settings used during the measurement. The reason for including this branch in the model is twofold: i) it allows direct extraction of the sought after quantity, without having to implement separate uncertainty propagation for λ_i and b_i and ii) preparing for future evaluations of multi diagnostic calibration factors, in which case one can simply extend this branch by supplying ECE raw plasma data to get an electron temperature profile that can also be supplied for example via Thomson scattering [47]. This would not be a simple cross calibration, but would rather combine the diagnostic specific calibration models and their corresponding plasma forward models, thus taking all information optimally into account. The calibration factors obtained that way are inherently consistent within the frame of the model.

In the Bayesian formalism this leads to (for the model shown in figure 4.1)

$$P(\boldsymbol{\varphi}^{\text{geo}}, \Delta\mathbf{s}, \mathbf{w}, \boldsymbol{\alpha}|\boldsymbol{D}) = \frac{P(\boldsymbol{D}|\boldsymbol{\varphi}^{\text{geo}}, \Delta\mathbf{s}, \mathbf{w}, \boldsymbol{\alpha})P(\boldsymbol{\varphi}^{\text{geo}}, \Delta\mathbf{s}, \mathbf{w}, \boldsymbol{\alpha})}{P(\boldsymbol{D})}, \quad (4.9)$$

with

$$\boldsymbol{\varphi}^{\text{geo}} = (\varphi_1^{\text{geo}}, \varphi_2^{\text{geo}}), \quad (4.10)$$

$$\Delta\mathbf{s} = (\Delta s_1, \dots, \Delta s_n), \quad (4.11)$$

$$\mathbf{w} = (w_1, \dots, w_n), \quad (4.12)$$

$$\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n), \quad (4.13)$$

$$\boldsymbol{D} = (\mathbf{D}_{\text{ECE},1}, \dots, \mathbf{D}_{\text{ECE},n}), \quad (4.14)$$

where n corresponds to the number of ECE channels, $\boldsymbol{\varphi}^{\text{geo}}$ to the angles at which the central line of sight switches from the cold source to the hot source and vice versa, $\Delta\mathbf{s}$ to the channel specific change in the bit signal observed when switching the temperature sources, $\mathbf{w}$ to the channel specific Gaussian beam width, $\boldsymbol{\alpha}$ to the channel specific variance scaling factors and $\boldsymbol{D}$ to the channel specific conditionally averaged measured data. The conditional averaging is done in the following way: First, the average number of data in one mirror rotation is calculated. The time series is then split at each falling edge of the chopper signal. Subsequently, resulting individual rotation measurements have their mean removed, are rescaled and sorted into the number of bins determined in the first step. Finally, dividing by the number of rotations yields

the conditionally averaged data required for the observation (cf. the grey node in figure 4.1). Note that the other models differ, for example by not using the variance scaling factors.

4.2.2. Generalised model

The model described previously uses some simplifications that can easily be dropped to generalise the model. For instance, one can drop the assumption that the problem is one dimensional and that there are only two reference temperatures. This allows easy extension to three (or more) reference temperatures for example by adding a hot ceramics hot source. Switching to cylindrical coordinates is a sensible approach for a geometry similar to the one presented here, thus introducing z along the horizontal axis of the cylinder shown in figure 2.6, and r as the radius. Assuming $i \times j$ reference temperatures leads to a definition of the effective temperature $\hat{T}_{\text{eff}}(\hat{\varphi}, \hat{z}) = T_{\text{eff},ij}$, where $T_{\text{eff},ij}$ is the effective temperature valid for $\varphi_i^{\text{geo}} \leq \hat{\varphi} \leq \varphi_{i+1}^{\text{geo}}$ and $z_j^{\text{geo}} \leq \hat{z} \leq z_{j+1}^{\text{geo}}$. The second dimension is represented by $\hat{z}$. As in the model described above, $T_{\text{eff},ij}$ can be the result of multiple layers contributing to the effective temperature at the selected coordinates, such that a dependency on r might occur as well. The calibration sensitivity $\tilde{\eta}(T)$ does not necessarily have to be linear, however, one will have to use free parameters for the temperatures and emissivities in this case, as the prediction will no longer be independent of these parameters. In general, any instrument function g can be used to calculate the appropriate weighted effective temperature that the radiometer would see by looking at (ϕ, z), therefore

$$T_{\text{eff}}(\varphi, z) = \frac{\iint \hat{T}_{\text{eff}}(\hat{\varphi}, \hat{z}) g(\hat{\varphi}, \varphi, \hat{z}, z) \, \mathrm{d}\hat{\varphi} \, \mathrm{d}\hat{z}}{\iint g(\hat{\varphi}, \varphi, \hat{z}, z) \, \mathrm{d}\hat{\varphi} \, \mathrm{d}\hat{z}}, \tag{4.15}$$

which in combination with the calibration sensitivity allows the calculation of the prediction.

4.3. Results

A typical excerpt of the calibration timetrace for a sensitive channel can be seen in figure 4.3b, while the corresponding chopper signal is shown in figure 4.3a. There are three points to consider: i) the drop in the chopper signal does not
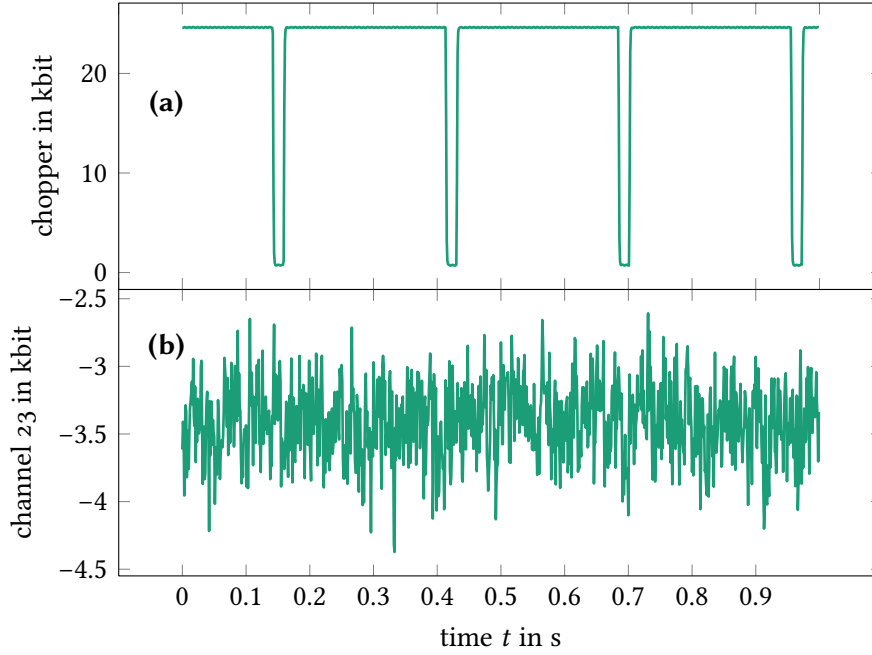
Figure 4.3.: (a) shows the chopper signal measured by a photo diode, **(b)** the raw signal of channel 23, which is a sensitive channel. The signal to noise ratio is typically of the order of 1/50. The background signal originates in broadband noise of the intermediate frequency (IF) amplifiers right after the mixer, subsequently measured for each frequency bin by the detector diode.

correspond to the full width of the cold phase, the real hot/cold duty cycle is approximately 0.2, given by the calibration unit geometry, increasing that value further would require a significant modification of the calibration unit, ii) even for the most sensitive channel the signal difference associated with the chopper channel cannot be seen directly, confirming that more elaborated analysis techniques are necessary and iii) no relevant drift within one rotation period can be observed. As said before, notable drifts were seen in some cases on timescales on the order of 10 s (cf. figure 2.5). The conditionally averaged signal in figure 4.4 supports the conclusion that there is no relevant drift within one period. The orange curve in figure 4.4a) corresponds to the measured and subsequently conditional averaged signal of the sensitive channel 23, $f_{23}^{\mathrm{meas}}(\varphi)$. Furthermore, the graph has been set to the mean values obtained from the previously run MCMC. Then, Monte Carlo samples (orange curves) have been drawn from
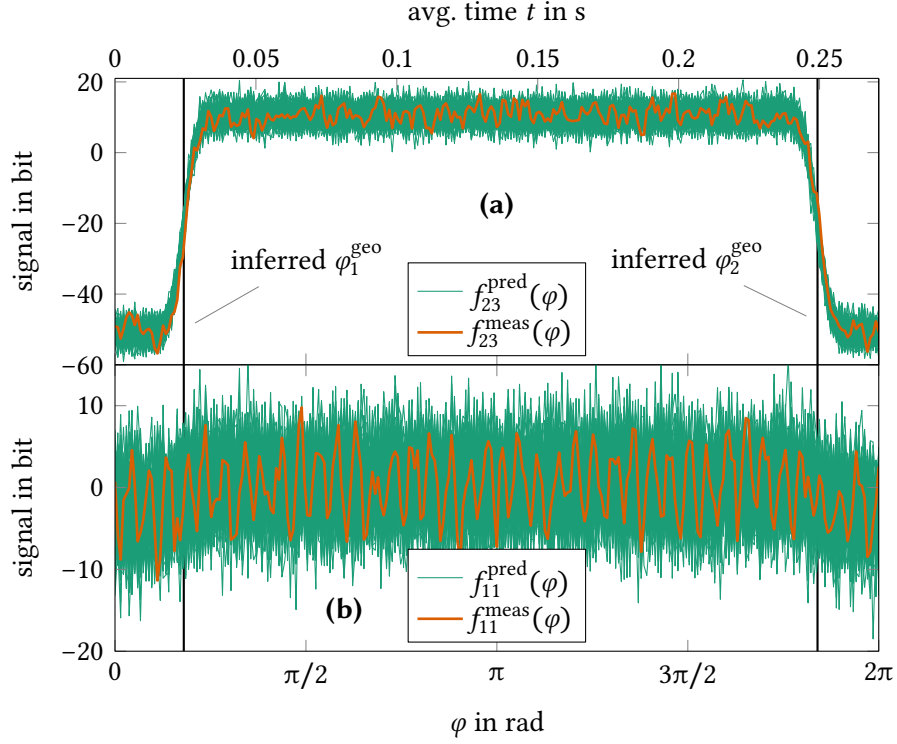
Figure 4.4.: (a) shows the sensitive channel 23, while **(b)** depicts the weakly sensitive channel 11. $\varphi = 0$ corresponds to a mirror position »looking« at the center of the LN$_2$ cold source. One can see the measured (orange) conditionally averaged and binned signal for both channels. The predictions (cyan) are 100 Monte Carlo samples, taken after the free parameters of the graph have been set to their mean values (i. e. a point estimate of the posterior predictive).

the prediction with its multivariate normal uncertainty. Consequently, these samples are predictions $f_{23}^{\mathrm{pred}}(\varphi)$ from a point estimate of the posterior predictive. Correspondingly, figure 4.4b) shows the weakly sensitive channel 11 with a barely noteable signal step between the not and cold phase. It is important to note that these samples are calculated from the model that allowed the scaling of the prediction variance, so that the uncertainties match the predictive capability of the model. Each cyan point in figure 4.4 corresponds to a predicted effective temperature scaled by the calibration sensitivity, with the offset of a whole period being removed.

channel 23
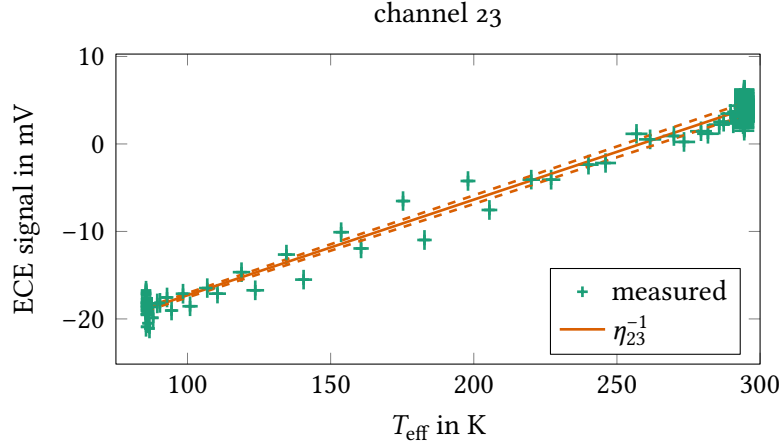
Figure 4.5.: Shown are the voltage values in dependence of the effective temperature. Do note that the plotted voltage uncertainties correspond to $b_{23}\sqrt{\alpha_{23}\sigma_{23}^2}$ with b_{23} the bit to Volt conversion factor. The temperature uncertainties correspond to the uncertainties given by the priors.

The measured bit values at each mirror angle φ are illustrated in figure 4.5, where the bit values have been scaled to represent voltages and the mirror angles were converted to their corresponding effective temperature. It should be noted that the uncertainties on the voltage axis are scaled with the channel specific variance scaling factor α_i. The orange curve shows the sensitivity with its uncertainties as calculated from the graphical model (cf. figure 4.1). Remaining deviations might be caused by 50 Hz noise (or its higher harmonics) that are not completely notched out by the bandpass filter properties of the conditional average. The plot highlights the advantage of this analysis method: While no other radiometer calibration approach known to the author uses the data that are taken when the horn pattern collects radiation from more than one reference temperature, this method allows to estimate the effective temperature (and corresponding uncertainties) reducing the overall uncertainty. However, one has to keep in mind that this is valid only as long as the Jeans law approximation is valid. For the radiometer and the reference temperatures used here this is a very good approximation. Furthermore, the frequency dependent beam width is predicted. That implies that if a measurement of the beam width is available, it can be easily supplied to the model as another observation, thus decreasing uncertainties of the sensitivities.
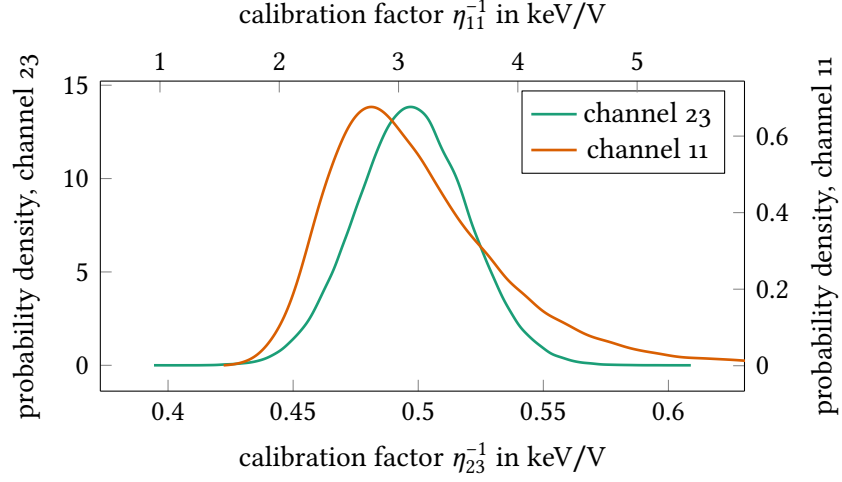
Figure 4.6.: The kernel density estimates of the MCMC samples of the calibration factors of the sensitive channel 23, η_{23}^{-1}, and the insensitive channel 11, η_{11}^{-1}. It is visible that the distribution of the insensitive channel deviates from a normal distribution by having a pronounced tail towards larger calibration factors. In contrast, the sensitive channel closely follows a normal distribution.

The kernel density estimates (i. e. the marginal posteriors) of the calibration factors as obtained from the MCMC are shown in figure 4.6. One of the strengths of an MCMC based evaluation is that one can get posterior distributions that are non-Gaussian as well. However, in the case shown here a Gaussian fit is a reasonable approximation to the posterior for the sensitive channel 23, while for channel 11 the posterior deviates notably from a Gaussian distribution. Nevertheless, for the sake of simplicity, the results of Gaussian fits to the marginal calibration posteriors are currently used for all higher level analyses. A gain in consistency could be achieved, especially for the weakly sensitive channels, if the asymmetric uncertainties originating from the MCMC would be used. To quantify the deviation for channel 11: The mean of a Gaussian fit to the kernel density estimate is at $(2.95 \pm 0.61)\,\text{keV}\,\text{V}^{-1}$, while the absolute maximum of the kernel density estimate is at $2.77^{+0.58}_{-0.78}\,\text{keV}\,\text{V}^{-1}$, corresponding to roughly 93.9 % of the Gaussian mean. However, the standard deviation shows a considerable difference to the asymmetric uncertainty. The predictive capability of the model is taken into account by scaling the variance via α_i, thus an appropriate estimation of the calibration factor uncertainties is achieved.
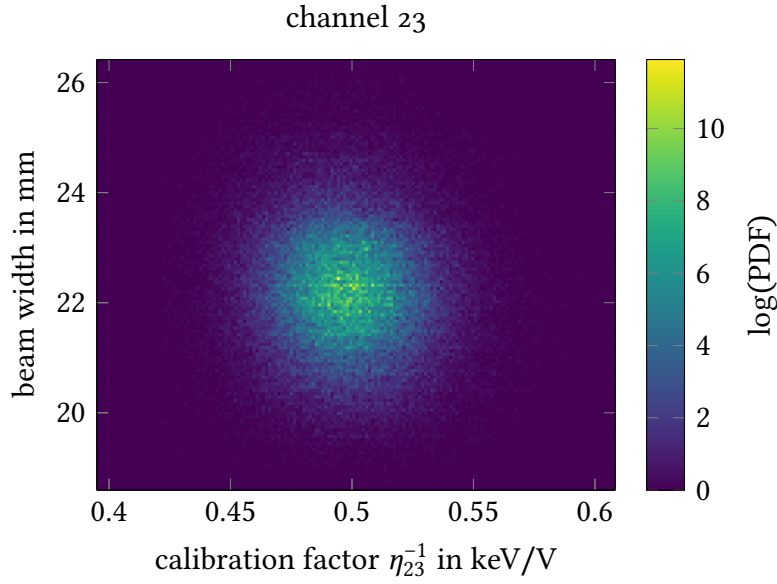
channel 23

Figure 4.7.: The posterior distribution for the sensitive channel 23 with the single channel evaluation model. One can see that the distribution is reasonably close to a 2D Gaussian distribution.

As the samples from the MCMC also allow to reveal correlations between different free parameters, figure 4.7 shows exemplarily the posterior distribution for the sensitive channel 23 for the the beam width and the calibration factor. The distribution is reasonably close to a 2D Gaussian distribution.

An example of a posterior revealing more complex relationships between two parameters can be seen in figure 4.8. There, the PDF of the angle that determines the beginning of the hot reference temperature measurement phase and the calibration factor is plotted. One can clearly see the deviation from a Gaussian distribution. It should be noted that the channel chosen here shows a relatively low signal to noise ratio. For channels with a better signal to noise ratio, the distribution resembles again a Gaussian distribution.

4.3.1. Comparison of single and multichannel evaluation

As it is not ab initio clear how large the differences between the models of varying complexity are, and thus which model is appropriate for practice, a
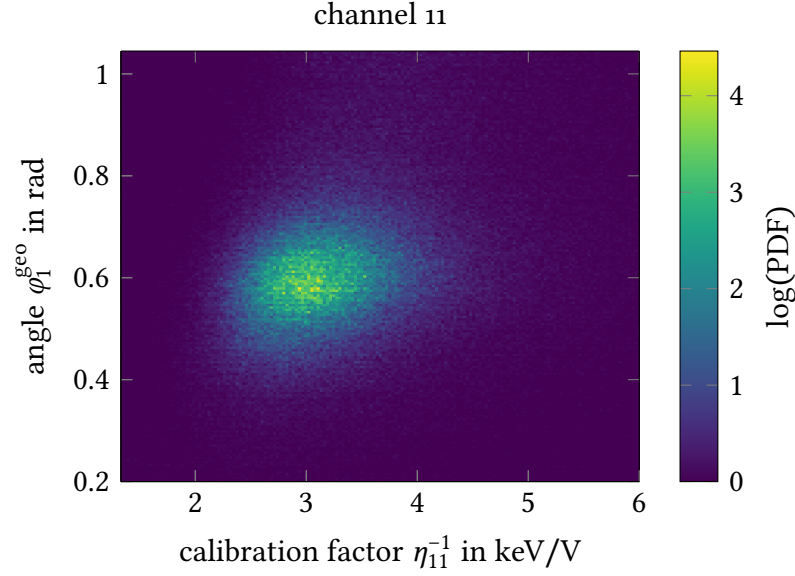
Figure 4.8.: Example for a posterior distribution between the mirror angle φ_1^{geo} and the calibration factor η_{11}^{-1} using the single channel evaluation.

careful comparison is shown in this subsection. Figure 4.9a shows the calibration factor, which is the inverse sensitivity, for each channel. One can see that the calibration factors vary over more than two orders of magnitude. A single mixer is used for the whole spectrum to allow for a better correlation analysis [52]. This is unusual as many ECE systems use multiple mixers to avoid frequencies above 18 GHz after mixing [84]. The single mixer approach leads to intermediate frequencies up to 40 GHz which need to be detected. The low sensitivities for higher frequencies might at least partially originate from different cables used for frequencies below 18 GHz (corresponding to frequencies below 140 GHz in the shown spectrum) and above 18 GHz. The cable frequency response damps higher frequencies more [36]. The conversion efficiency of the extreme broadband mixer also drops for frequencies above 140 GHz. Individual diode sensitivities are expected to play an important role as well [37].

Figure 4.9b shows the uncertainties of the different models, normalised to the multi shared model. Going from the single channel evaluation to a combined model yields substantial decreases in the calibration factor uncertainties of insensitive channels, although the uncertainties for these channels remain
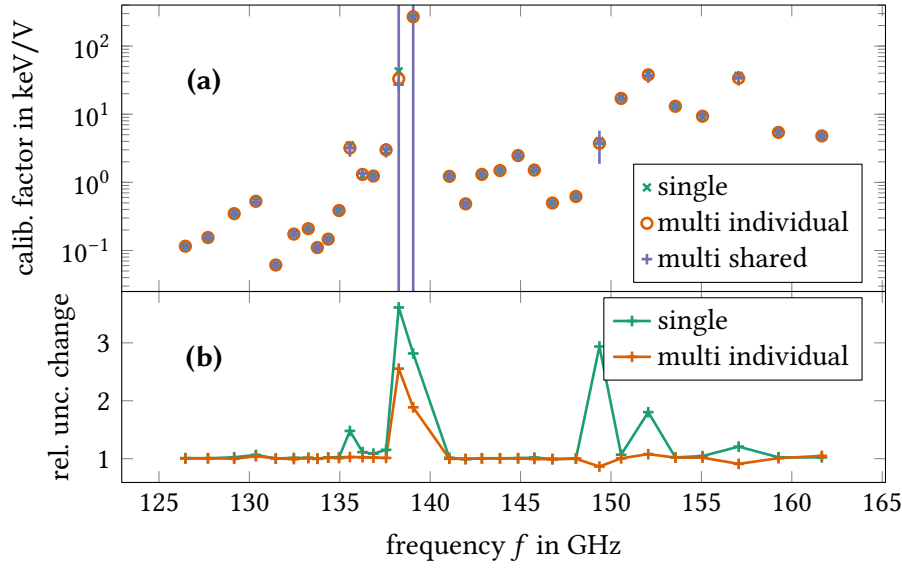
Figure 4.9.: (a) shows the frequency dependent calibration factors as obtained for single channel evaluations and the multi-channel analysis with and without a shared free parameter for the beam width. The uncertainties are only shown for the multi-channel analysis with a shared (and appropriately scaled) beam width, as this is the reference value for the relative uncertainty changes shown in **(b)**. The standard deviation of the specified models is normalised by the standard deviation of the multi shared model. The given uncertainties correspond to one standard deviation as calculated from the MCMC samples.

very large. This phenomenon is most likely caused by the additional information about the geometrical properties φ^{geo} that is mainly provided by stronger channels, helping the less sensitive channels to determine the begin and end of the hot/cold phases. Using a single beam parameter leads for a few channels to a small shift of the calibration factor, also reducing the uncertainties slightly, but less drastic than the switch from the single channel evaluation to a multi-channel evaluation model.

Figure 4.10 shows the inferred intensity Gaussian beam width for each ECE channel. The beam width has been measured in the lab with a 140 GHz source attached at the receiver end of the antenna and an 2D array of 440 intensity measurements in various fixed distances to the minimum beam waist. A simple Bayesian Gaussian squared forward model that adapts the prediction
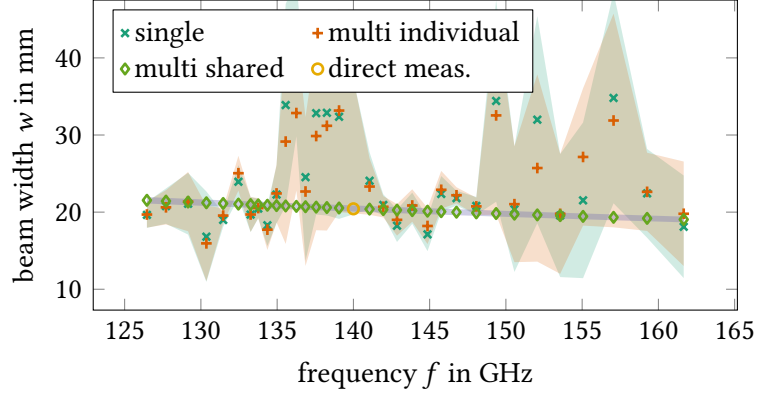
Figure 4.10.: The inferred Gaussian beam width of the microwave antenna characteristic for each channel. Reliable values can be obtained if the model using a commonly scaled beam width according to equation (4.1) is used (multi shared). The shaded areas correspond to one standard deviation (from MCMC).

uncertainties was used to infer the beam width from the intensity measurements with 10 million MCMC iterations. The inferred beam width was roughly (20.4 ± 0.2) mm, with 37.5 cm distance between the minimum beam waist and the array measurements. This does not exactly match the distance at which the microwave foam is located relative to the minimum beam waist (which would be roughly 26 cm to 32 cm), but due to the small divergence of the beam width the introduced error is small. One can see that switching from the single channel analysis to a combined channel analysis slightly decreases the uncertainties for some less sensitive channels. In most cases the beam width shifts slightly towards values closer to the directly measured width. If only a single beam width, scaled according to equation (4.1), is used, the uncertainties get drastically reduced. The measured width is roughly 40 μm away from the predicted value by the model with a single beam parameter, with prediction uncertainties on the order of 0.4 mm.

The variance scaling factors for the different channels are shown in figure 4.11. One can see that the values are not too far away from 1, which indicates that the most relevant physic effects are considered. These values were reduced from values typically around 2.6 at the begin of the first W7-X experimental campaign by two changes of the setup: i) the horn was changed to include all polarizations and ii) the container for the liquid nitrogen had a round aperture
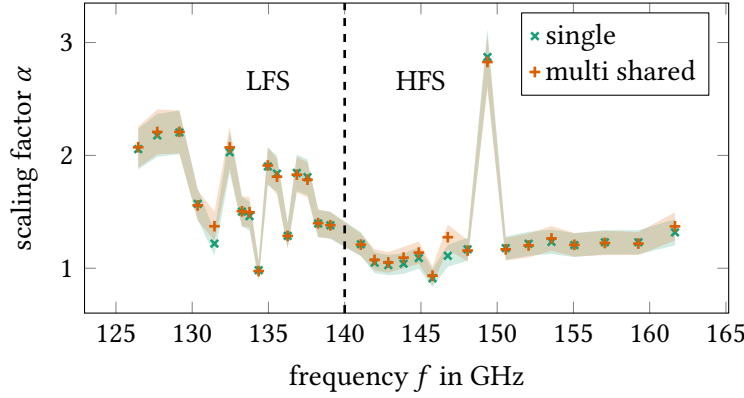
Figure 4.11.: The inferred variance scaling factors α for the covariance is a measure of the model uncertainty for each channel. For the »multi individual« case there is no visible difference to the »multi shared« case. Note that the channel at 149.36 GHz shows 150 Hz noise that dominates the Fourier spectrum even after applying the conditional averaging. Subsequent calibration measurements without this noise yielded $\alpha_{25} \approx 0.99$.

that was changed to a rectangular aperture. The round aperture system was more sensitive to misalignments of the mirror-antenna system. Indeed an offset of around 3.5 cm of the beam on the mirror was measured for the calibration used for the first experimental campaign. This offset complicates the geometry and was not reflected in the model. As the physics for each channel should be similar, α should have similar values for the different channels. A notable difference between the α values is an indication that different physics effects play a role, or at least that these different effects are of different importance. One could expect that for channels with a small sensitivity electronics effects are of larger importance, therefore changing the variance scaling. Notable differences between the models that couple the beam width of different channels directly and those that do not, and between the individual and multi-channel evaluations are not observed. The values on the low field side (below 140 GHz) scatter more and tend to be larger. This indicates that the uncertainties are underestimated for these channels, respectively, that the explanatory power of the model is *smaller* than for channels on the high field side. A potential source for this behaviour can be ascribed to implicit hardware assumptions. For example, these can be violated more strongly for low field side channels, although currently
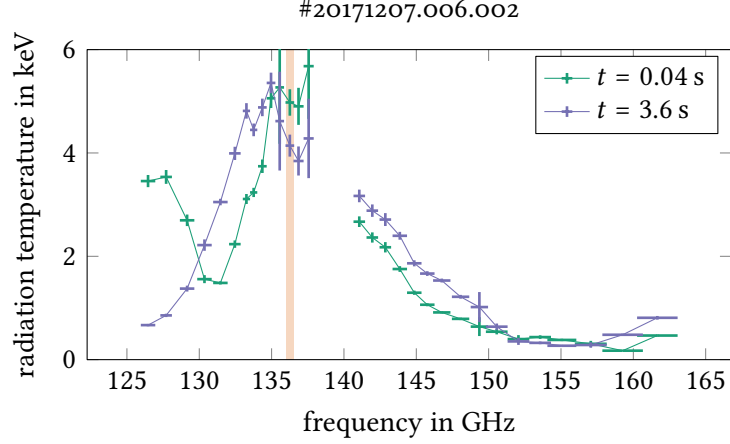
Figure 4.12.: Exemplary radiation temperature spectrum from a centrally ECRH heated W7-X plasma discharge, calculated with the calibration factors as obtained from the Minerva model that incorporates variance scaling. The channel marked in orange corresponds to the ECE timetrace shown in figure 4.13.

no such problem is observed. In summary, one can see that the single channel evaluation is satisfactory in most cases. However, if time is not a critical factor it is still beneficial to use a model combining all ECE channels in one Bayesian model.

4.3.2. ECE spectra

From the calibration procedure radiation temperature spectra can be derived from measurements done during a plasma discharge. An example is shown in figure 4.12. The data originate from a 3.8 s long plasma discharge that was centrally heated with ECRH [85, 86]. Time traces of the main plasma paremeters are shown in figure 4.13. The ECRH power was 2.5 MW in the first phase and was increased to roughly 5 MW shortly after pellet fuelling started. The line averaged electron density as measured by a single channel dispersion interferometer [87] rose during pellet fuelling up to about 7×10^{19} m^{-3}.

Channel 15 (138.26 GHz) and 16 (139.06 GHz) show a very low sensitivity, leading to radiation temperatures above 20 keV and uncertainties of several hundred percent. Consequently, they were omitted in this plot. For channel 16 this is expected, as the channels frequency band locates it in the slope of the notch
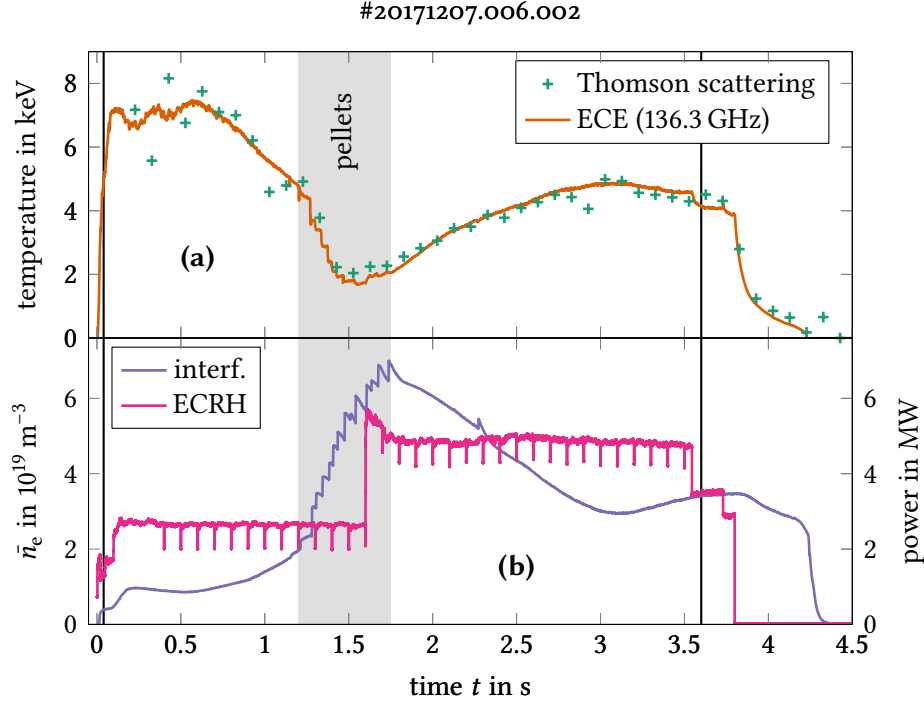
#20171207.006.002

Figure 4.13.: **(a)** shows the calibrated signal of an ECE channel close to the plasma core, compared with a central channel from the Thomson scattering system [47] for the plasma discharge also shown in figure 4.12. The lines at 0.04 s and 3.6 s indicate the spectra shown in figure 4.12. **(b)** depicts the ECRH and line averaged density from the single channel dispersion interferometry. The ECRH blips are necessary for the collective Thomson scattering (CTS) diagnostic [88].

filter. Above roughly 155 GHz contributions from the third harmonic X mode emission start to play a role. The spectrum from 0.04 s shows »shinethrough« from hot core electrons below 131 GHz. The reconstruction of the electron temperature profile from the spectrum will be discussed in chapter 5. However, figure 4.13 shows a comparison of the ECE timetrace of a channel close to the core to a Thomson scattering channel close to the core. The deviations in the first second are probably caused by the selection of filters used in the polychromators of the Thomson scattering system that lead to big uncertainties for electron temperatures above around 7 keV. The remaining seconds of the

discharge show a good agreement for a wide range of electron densities and multiple power levels. One of the optimization criteria of W7-X is the minimization of Pfirsch-Schlüter currents, which increase with plasma pressure [89, 90]. The Pfirsch-Schlüter currents give rise to the so-called Shafranov shift moving the magnetic axis outwards. Hence, having the Shafranov shift minimized makes the magnetic configuration less sensitive to plasma pressure changes. For the ECE a small Shafranov shift implies that the positions of the channels are less sensitive to pressure changes than they would be in an unoptimized stellarator. Nevertheless, the Shafranov shift combined with the diamagnetic shift (i. e. the finite plasma β effect) of the cold resonance positions of the ECE channels for the experiment shown in figure 4.13 amounts to a maximum of approximately 5.5 cm, or 0.35 in effective radii. Thus, depending on the pressure profile the ECE channel appropriate for a plasma center measurement might change. This implies that the ECE at W7-X is a diagnostic rather sensitive to the total plasma pressure.

4.4. Conclusions

The Bayesian Minerva model developed for the multi-channel consistent calibration of a generic microwave radiometer provides the sensitivities in an automated fashion, insensitive to signal drifts on timescales larger than 1 s. By using an explicit model for the calibration, the analysis can be done in a more formalized way. This also allows to obtain feedback on how well the modelled physical system is understood. In addition, the beam width can be inferred accurately. Also, classical error propagation is linear while a Bayesian analysis can handle nonlinearities.

Moreover, it allows to obtain non-Gaussian posterior distributions for the calibration factors, although for the sensitive channels in the radiometer studied a Gaussian distribution provides an excellent fit to the posterior distribution.

The use of a variance scaling factor allows to obtain uncertainties matching the predictive capability of the model.

The ECE spectra obtained from typical W7-X plasmas do not show unexpected or unphysical features. For reasonably sensitive channels the uncertainties are typically on the order of 6 %. Note that the contribution of statistical noise is small compared to the uncertainties arising from the calibration. While the high dimensional (98D) model provides in principle the highest consistency,

a comparison to the much quicker and parallelized single channel evaluation (4D for each channel) yields only little differences, justifying the use of the simpler approach in most cases. Nevertheless, the comparison with the high dimensional model was a good consistency check.

Further improvements could be achieved by applying neural networks, as described in Pavone et. al. [78] to speed up the evaluation. It also would be interesting to use the calibration model in combination with a plasma model containing predictions for Thomson scattering and ECE, guaranteeing consistent calibration factors across different diagnostic systems with the models presented in this chapter providing the tested and prepared ECE calibration branch.

Appendix B shows the code, respectively keywords, that allow customization of the developed standalone Java executable.

<div align="right">5</div>

Plasma profile inversion

Typically, many ECE analyses rely on the underlying assumption of a black body emitting plasma with clearly localized ECE positions, which is not necessarily fulfilled over the full plasma radius. In particular, for W7-X, the lowest frequencies probe nominally the plasma edge at the low field side. There, the optical depth can be low (i. e. $\tau < 3$), with the relativistically downshifted emission from hot core electrons superposed to the thermal emission from the edge. This makes the direct use of T_{rad} as a surrogate for T_{e} impossible in these regions. In W7-X this »shinethrough« effect can be on the order of several keV for some channels and thus become dominant, for example during plasma buildup in figure 4.12 at 0.04 s below 131 GHz. However, the »shinethrough« effect occurs in plasmas with a low electron density, or narrow plasmas with sufficiently hot electrons, and is not notably visible in the examples shown in this chapter. An example is depicted in section 6.2. Quantitative analysis of the underlying electron cyclotron emission and absorption and the corresponding radiation transport is handled by the TRAVIS code [43] (see also section 2.2). Note that the model itself is by no means bound to TRAVIS; any ray tracing and radiation transport code could be implemented straightforwardly.

This chapter is structured as follows: In section 5.1, a detailed description of the Minerva model and a subsequent description of the methods used to obtain the initial guess, for both the electron temperature and density profile, is given. Therein, the approach presented here is also compared to other ways used to obtain T_{e} profiles. Thereafter, section 5.2 shows and discusses results for an example timepoint in a typical W7-X discharge, while section 5.2.3 debates the influence of NBI heating on the electron temperature and density profiles as inferred from ECE. Do note that none of the chosen timepoints do exhibit

notable »shinethrough« behaviour, as this is behaviour is no longer the norm in discharges from OP 1.2a onwards. Nevertheless, the code can fully handle »shinethrough« effects, as also discussed in section 5.1. Finally, section 5.3 concludes the results and gives an outlook.

5.1. Minerva implementation

Sophisticated ECE analyses are also used at other plasma experiments, for example at JET [91–93] and AUG [94–96], both tokamaks. At JET, Minerva is also used as the embedding framework for the ECE model, while instead of the raytracing code TRAVIS the SPECE code is used [97]. A comparison to these approaches can be found at the end of this section.

The available Minerva models are shown in a simplified manner in figure 5.1. Subsequently, the structure of the graphical model will be briefly explained. The descriptions of the continuous electron temperature and density profiles are done in one of two ways in the model: The first approach consists of the parametrization that is often used in conjunction with TRAVIS,

$$f(\rho) = a_0\Big(g - h + (1 + h - g)(1 - \rho^p)^q + h\big[1 - \exp\big(-\rho^2/w^2\big)\big]\Big), \qquad (5.1)$$

with ρ the variable normalized effective radius, $\rho \in [0, 1]$. The parameters are: a_0 the value at $\rho = 0$, g the ratio of edge and core value ($f(1)/f(0)$), p and q that dominate the profile gradients, h the profile hollowness and w the profile hollowness width, the latter two being useful for hollow profile shapes as well as peaked profiles. Figure 5.2 illustrates some example profile shapes obtained with the parameterization of equation (5.1). The second approach uses a cubic interpolation between the evaluation points of a Gaussian process (cf. section 3.1) to obtain a continuous profile. For an idea of how possible profile shapes look like if a Gaussian process is used, see figure 3.1 as an example of T_e profiles.

To be able to flexibly penalize deviations from zero gradients in the core virtual observations (cf. chapter 3) have been added. Note that decreasing the uncertainty of the predicted gradients allows to increasingly force the gradients to be close to zero – at the cost of a narrowed posterior. Subsequently, the T_e and n_e profiles are fed to the TRAVIS nodes for the X mode and O mode predictions. These TRAVIS nodes get the required ECE metadata, like the line of sight or the
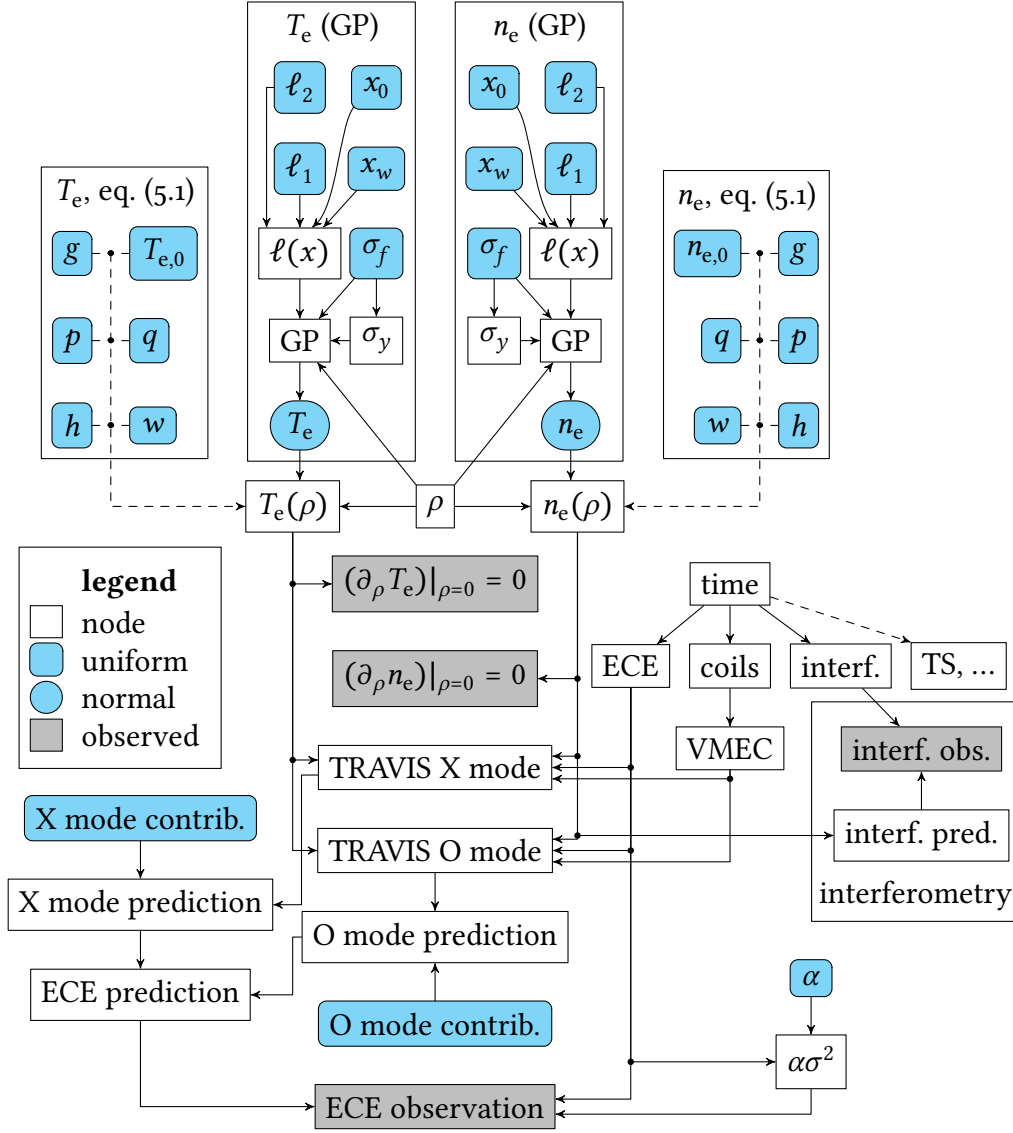
Figure 5.1.: Simplified Minerva ECE profile inversion model, cf. section 3.4. Either the Gaussian process implementation is used to describe the T_e and n_e profiles, or equation (5.1), which can also be used directly for TRAVIS. The edges do not represent data flow.
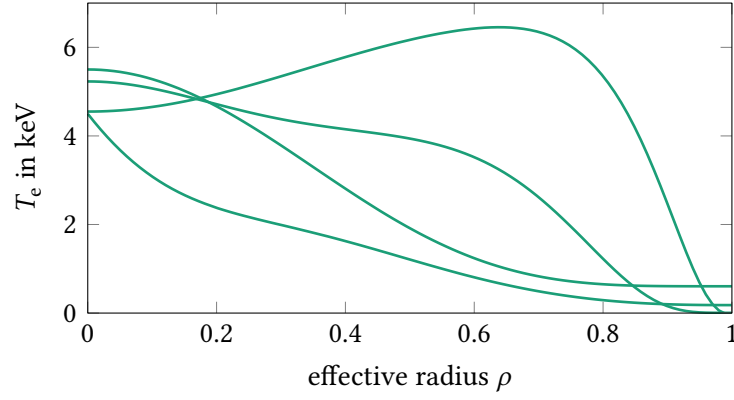
Figure 5.2.: Shown are some examples for profile shapes that can be obtained using the TRAVIS parameterisation given in equation (5.1). It is not possible to get a profile with multiple bumps.

beam diameter from the ECE datasource. The magnetic configuration, respectively the VMEC equilibrium, can be obtained in three ways: i) from a VMEC run directly provided by the user through a VMEC run identifier of the webservice, ii) from a dedicated VMEC run based on the settings of a user specified VMEC identifier, but taking the coil currents of the selected experiment directly from the corresponding coil current datasource and iii) from the VMEC identifier supplied by a datasource, also written during this thesis, that provides the best guess of the equilibrium from a precomputed set of reference equilibria, taking into account the diamagnetic energy [98]. Note that there is no automated equilibrium reconstruction running at W7-X at the moment. The resulting X mode and O mode predictions are then rescaled each by a uniform free parameter between 0 and 1. This is done to take into account that we do always measure a mixture of X mode and O mode, and, especially for the O mode, to at least partially account for contributions from multiple plasma passes. Thus, these scaling factors should not be interpreted to reflect exactly the X mode or O mode content. Experimentally, it is expected to measure more than about 96 % X mode and correspondingly less than 4 % O mode while using the experimental setup for the X mode. Adding the rescaled X mode and O mode predictions yields then the overall ECE prediction which is compared to the measured ECE spectra obtained via the datasource. A scaling parameter α on the variance σ^2 can be used to have the uncertainties match the predictive capability of the model, see
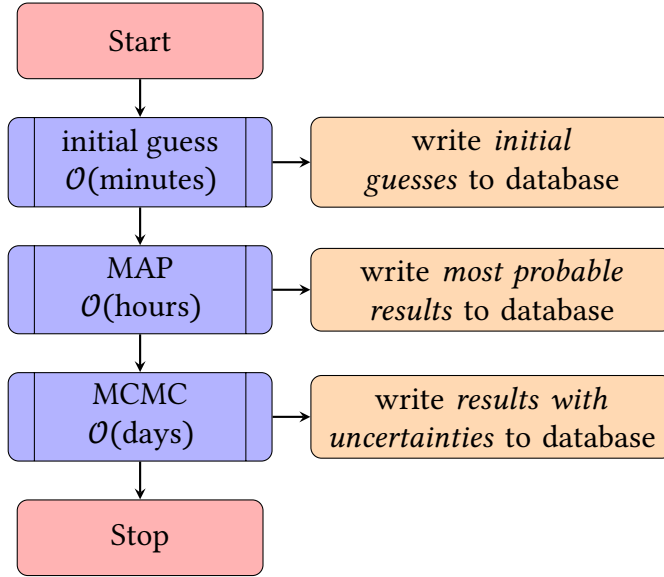
Figure 5.3.: First, an initial guess is calculated by fitting the cold resonance T_{rad} profile (which is assumed to be T_{e}). This takes on the order of few minutes. Subsequently an MAP inversion is done, finding the most probable values for the free parameters. This usually takes up to 24 h. Finally, an MCMC inversion is done, such that samples drawn from the posterior allow an estimation of non-Gaussian correlated uncertainties.

chapter 3. Furthermore, the electron density profile is integrated over the line of sight of the interferometer and the result is compared to the line integrated density as measured by interferometry, thus keeping ECE and interferometry consistent by predicting both diagnostic measurements simultaneously.

In figure 5.3 the general ECE profile analysis scheme is shown. The analysis chain is as follows: i) the initial guess is calculated after discarding channels that are most likely affected by »shinethrough«[1], followed by ii) an MAP inversion with a maximum of 150 iterations and finally iii) an MCMC inversion with a burn-in of 10^5 iterations and 3×10^4 samples drawn afterwards. After each stage in the analysis chain, the results are written to the central ArchiveDB [99] to make them available to the W7-X team.

---

[1]As a reminder: This typically affects edge channels with an optical depth $\tau < 3$, for which relativistically downshifted emission from the center is not fully reabsorbed. For an example see figure 4.12.

To determine the initial guesses in the case the T_e and n_e profiles are defined by Gaussian processes, the following steps are performed. The prior *distributions* for the T_e and n_e profile distributions are multivariate normals with a mean of zero and a covariance given by the hyperparameters ℓ_1^i, ℓ_2^i, x_0^i, x_w^i and σ_f^i [62] with i the values for either the electron temperature or density profiles, cf. section 3.1. The prior *values* are typically obtained in the following way: First, a simple Gaussian fit to the radiation temperatures at the cold resonance positions of the ECE frequencies for the given VMEC equilibrium is done. This Gaussian fit is used to determine channels which are likely to be affected by »shinethrough« by excluding channels whose radiation temperature exceeds the fit value by more than one standard deviation (dominated by the calibration uncertainty for the channel as calculated by the analysis in chapter 4)[2]. The radiation temperatures of these channels are then put as observed values in a model with the same Gaussian process setup as in the main ECE inference model (cf. figure 5.1), so an identical length scale function, depending on the same hyperparameters is used. The only difference is the always activated variance scaling of the prediction (cf. chapter 3), which has the advantage to reduce the susceptibilty to deviations of the used VMEC equilibrium to the true plasma equilibrium. As this simpler model is linear, one can use linear Gaussian inversions (cf. section 3.4.2) to solve the problem analytically for fixed hyperparameters. This allows for a Hooke and Jeeves optimization (cf. section 3.2) using a cost function that combines the probability densities of the *evidence* with the hyperparameter prior. Effectively, this leads to an optimization of the Gaussian process and its hyperparameters by applying Occam's razor, therefore taking model complexity into account by penalizing more complex models, see chapter 3. The resulting hyperparameters are then put in the main ECE model as values for the prior T_e hyperparameter distributions. Correspondingly, the resulting T_e values of the truncated multivariate normal are put in the T_e distribution of the main model as well. To make the initial guess even more robust against equilibria that for example do not take plasma pressure into account, one can specify a shift of the initial cold resonance positions, thereby changing the fitted electron temperature profile. Practically that means that the predictions of the spectrum are calculated for shifts ranging from $-\Delta\rho$ to $\Delta\rho$ in steps of

---

[2]Two further conditions have to be met: i) the channels effective radius exceeds $\rho > 0.5$ and ii) all channels further out (i. e. with a larger effective radius) were identified as »shinethrough« affected, although ii) is applied separately for the low field side and high field side.

$\delta\rho$, comparing the logarithm of the joint (i. e. prior times likelihood) probability density of the ECE observation for the shifts and selecting the shift that gives the best prediction. While this does cost some time, the additional computing time (on the order of seconds to few minutes) is small compared to the overall analysis time (on the order of days), making the gain in robustness worthwhile if the available VMEC equilibrium is expected to deviate notably from the real equilibrium. Usually, the reference VMEC equilibrium found automatically[3] is good enough that shifting the profile is not necessary. Exemplarily, the residuals of the initial guess of the time point analysed in section 5.2 are shown in figure 5.6. Note that there are other options available to determine the initial T_e profile, although not regularly used. They consist of: i) the same procedure described above, but using only channels on the low field side, ii) the same procedure described above, but using only channels on the HFS, iii) using the T_e values as determined by the Thomson scattering system to fit a Gaussian process to, iv) fitting the Gaussian process to a manually provided T_e profile and v) fitting the Gaussian process to a T_e profile parameterized by equation (5.1).

The Thomson scattering n_e profile [47] provides the observed values for the same Gaussian process procedure described above, and the result is subsequently rescaled such that the integral over the density profile along the interferometer line of sight[4] matches the measured line integrated electron density value by interferometry. Do note that the deviation between the Thomson scattering and interferometry system can be caused for example by calibration issues of the Thomson scattering system. It is known that misalignments of the Thomson scattering lasers occuring occasionally also contribute significantly to a distortion of the Thomson scattering n_e profile. This is further complicated by the Thomson scattering system using three different lasers from OP 1.2b onwards, such that potential misalignments of multiple lasers have to be taken into account if averaging over a time window long enough to contain multiple Thomson scattering laser pulses is desired. Sticking to one laser avoids prob-

---

[3]The reference equilibrium is found by first comparing the experimental coil currents with the coil currents from a precalculated set of equilibria and selecting all equilibria that have a reasonably close magnetic configuration. Amongst these equilibria, the one with the smallest deviation of their integrated plasma pressure profile to the maximum experimental diamagnetic energy during the discharge is selected.

[4]Technically, the interferometer has two slightly different lines of sight, as the laser passes through the plasma, gets reflected and passes the plasma again. The exact lines of sight are used for the calculations.

lems with different laser alignments, but reduces the frequency with which usable data is available. The rescaling is an easy method to make the n_e profile consistent with the line integrated electron density as observed by the interferometer, such that any deviation from this initial density profile in the inference is driven by information from the ECE. Other options available as reference for the initial density guess are: i) unscaled Thomson scattering profiles, ii) default profiles scaled by interferometry, iii) manually chosen profiles and iv) manually chosen profiles scaled to the line integrated density value measured by interferometry.

It is important to put the zero gradient constraints via virtual observations at the core already on the initial guesses. If one does not put them there, the low probability density for the virtual observation nodes is going to dominate the joint probability density. Hence, obtaining quickly a satisfying prediction is much more difficult as the optimization strives to reduce the virtual observation penalty, more or less regardless of the remaining profile shape.

While the option to let the Gaussian process hyperparameters (i. e. the parameters that determine the prior distribution of other parameters) free is fully implemented in the main model, typically one will keep them at their initial guesses, as the computational cost of using them as free parameters is considerable. If TRAVIS and/or the analysis routines, like the MCMC, are sped up sufficiently, leaving the hyperparameters free could help less constraining the posterior. The evaluation strategies for the MAP and MCMC inversion are detailed in section 3.4.2.

5.1.1. Limitations

Five systematic uncertainties have been identified which hamper the ECE spectrum interpretation and thus the profile inference: i) the effect of electron density outside the last closed flux surface, ii) multipass microwave radiation transport, iii) using only a central ray instead of multiple rays, iv) noise levels respectively parasitic gyrotron mode effects for individual channels and v) errors in the VMEC equilibrium.

As the density does not drop to zero at $\rho = 1$, the best approach would be to treat the density outside the last closed flux surface fully three dimensional to avoid unphysical refraction patterns caused by abrupt changes in the refractive index. However, as there is no density measurement along the line of sight of the ECE, and the computational cost of a 3D treatment would be high, a

workaround will be used in the future: A simple extrapolation of the density profile as measured by the Thomson scattering profile outside of the last closed flux surface will allow a smooth density drop off, most likely reducing the error that is made with respect to the ECE beam refraction.

Another problem is the handling of multiple passes of the beam through the plasma. After e. g. the X mode passed through the plasma once, it does get reflected. Upon hitting the plasma edge again, the polarisation does not entirely match the one that would be required for the ray to continue as pure X mode. Hence, one would have to continue with part of the power in X mode and part in O mode. In practice, TRAVIS neglects the power that would go to the O mode and continues just with the X mode. This applies to each further plasma entrance as well. In addition, the requirement for the angle with which the beam hits the plasma vessel after the first plasma pass being accurate and precise grows with the number of reflections taken into account, as slightly different angles may lead to completely different beam paths after several reflections due to the complex 3D shape of the W7-X vessel. This makes it impossible to obtain an accurate O mode prediction if more than two reflections are required. How large this effect in W7-X is remains an issue for further studies.

Furthermore, due to the 3D geometry of W7-X, many rays instead of just the central ray would be necessary to take into account reflections at the wall realistically. However, the calculations were done with a single ray passing once through the plasma due to the accompanying computational cost of multiple rays. Note especially that using i rays with j reflections each would increase the computational cost already by a factor of ij, even though the contributions from the ray polygon intersection tests that have to be performed for each reflection (cf. section 2.2) are still neglected.

It should be noted that in the last operational campaign, OP 1.2b, some channels in some discharges showed highly anomalous behaviour, most likely caused by radiation from parasitic modes from the gyrotrons [100] several GHz away from 140 GHz and thus not filtered out by the ECE notch filter. However, this is a problem that in the future will be solved on the hardware side by slightly shifting the frequency bands of the ECE channels, hopefully avoiding these parasitic mode effects.

Another huge factor are errors in the VMEC equilibrium. Due to the shallow gradient of the magnetic field strength the ECE at W7-X is highly sensitive to small changes in the absolute magnetic field. If the VMEC equilibrium e. g. fails to describe the plasma pressure profile accurately the positions from

where the radiation is expected can shift notably, cf. the discussion at the end of section 4.3.2. Currently no automatic (Bayesian) equilibrium reconstruction is running at W7-X, although preparations are being made in the Minerva framework.

5.1.2. Comparison to other ECE inference procedures

In order to be able to place the analysis presented here into context, differences and similarities to the analysis approaches used at JET as well as AUG are discussed.

The inference procedure used at JET uses Minerva as well, but uses SPECE instead of TRAVIS [91–93]. Solver for plasma electron cyclotron emission (SPECE) is specific for tokamaks: Wall reflections are approximated by assuming parallel walls on the inner- and outboard side, whereas TRAVIS can handle 3D geometries and is thus suitable for stellarators as well as for tokamaks (cf. section 2.3.2). Further differences between the Minerva approach chosen at JET and at W7-X can be found in the different kernels of the Gaussian processes. The model at JET uses a generalized squared exponential kernel [58, 93], while at W7-X the kernel based on [61] (cf. section 3.1) is used. Also, the model at JET includes a reflectometer, two Martin-Puplett interferometers and a radiometer, while the W7-X model includes the heterodyne radiometer together with the single channel dispersion interferometer. However, due to the implementation of TRAVIS within Minerva, adding reflectometers to the W7-X model is straightforward, as TRAVIS can be used to predict this diagnostic as well. In addition, virtual observations to penalize non-zero T_e and n_e gradients in the core are used in the W7-X model and the models to obtain the initial guess differ substantially. Extensive details of the JET model can be found in [93].

At AUG, data from the ECE system, the lithium beam emission spectroscopy (LiBES) system and deuterium cyanide laser interferometry (DCI) system are combined [94–96]. Uncertainties are determined via χ^2 binning, while the approach presented in this thesis uses an MCMC to obtain samples from the posterior. The ECE uncertainties for the radiation temperature at the two machines are given by

$$\Delta T_{\text{rad}}^{\text{AUG}} = 7\,\% \, T_{\text{rad,dat}} + \left(T_{\text{rad,dat}} / f_{\text{SNR}} \right)^2 + \sigma_{\text{stat}} + 15\,\text{eV}, \tag{5.2}$$

$$\Delta T_{\text{rad}}^{\text{W7X}} = \sqrt{\sigma_{\text{calib}}^2 + \sigma_{\text{stat}}^2}, \tag{5.3}$$

where $7\,\%T_{\mathrm{rad,dat}}$ includes the systematic calibration uncertainties, f_{SNR} is a factor considering the channel specific signal to noise ratio steming from the calibration, σ_{stat} are statistical uncertainties from the chosen time window. The digitization error is specified to be 15 eV. The calibration uncertainties at W7-X are denoted by σ_{calib} and are determined as described in chapter 4. For the initial guess of the electron density profile at AUG the LiBES and DCI diagnostics are used, and by default a shift determined by the electron density profile at the maximum curvature shifted to the point where the electron temperature equals 100 eV. The energy flux density used to calculate the absorption coefficient in the AUG model uses the electromagnetic energy only, while TRAVIS also accounts for the »sloshing« energy (i. e. a flux of kinetic energy caused by coherent particle motion in the ECE wave [40]). Currently, wall reflections at AUG are handled in 1D with a reflection coefficient that can be chosen as a model parameter, while TRAVIS works in 3D and reflection coefficients also defined in 3D, but not used as model parameters. Further assumptions at AUG include $T_{\mathrm{e}}(\rho > 1.02) < 50$ eV and $T_{\mathrm{e}}(\rho > \rho_{\mathrm{Wall}}) < 2$ eV. However, the evaluation at AUG takes considerably less time compared to the full analysis at W7-X ($\mathcal{O}(\mathrm{min})$ for the initial guess, $\mathcal{O}(\mathrm{hour})$ for the MAP and $\mathcal{O}(\mathrm{day})$ for the MCMC).

5.2. Results

The previously described model was used to analyse exemplarily the plasma at two times in a high mirror configuration in #20180823.016.002. Figure 5.4 depicts an overview of that discharge. In the first second, the ECRH delivers approximately 2 MW, the line averaged density rises to about $3 \times 10^{19}\,\mathrm{m}^{-3}$, a Thomson scattering system channel and an ECE channel similarly close to the plasma core show electron temperatures in reasonable agreement. From 1 s to 2 s the plasma is relatively stationary. After the ECRH power has been increased to 4 MW at 1 s, the electron temperature and density rise to an electron temperature of circa 3 keV and a line averaged density of approximately $4.6 \times 10^{19}\,\mathrm{m}^{-3}$ at 2 s. From 2 s to 4 s the ECRH is modulated with a frequency of 17 Hz to allow for heat pulse propagation measurements. The first profile analysed within this chapter was taken at 4.45 s, with an averaging over 50 ms to have the statistical contribution to the ECE uncertainty vanish, while the contribution of the calibration to the ECE uncertainty remains unchanged. Until that time, the line averaged density decreased slightly to about $4.1 \times 10^{19}\,\mathrm{m}^{-3}$. At 4.5 s NBI heating
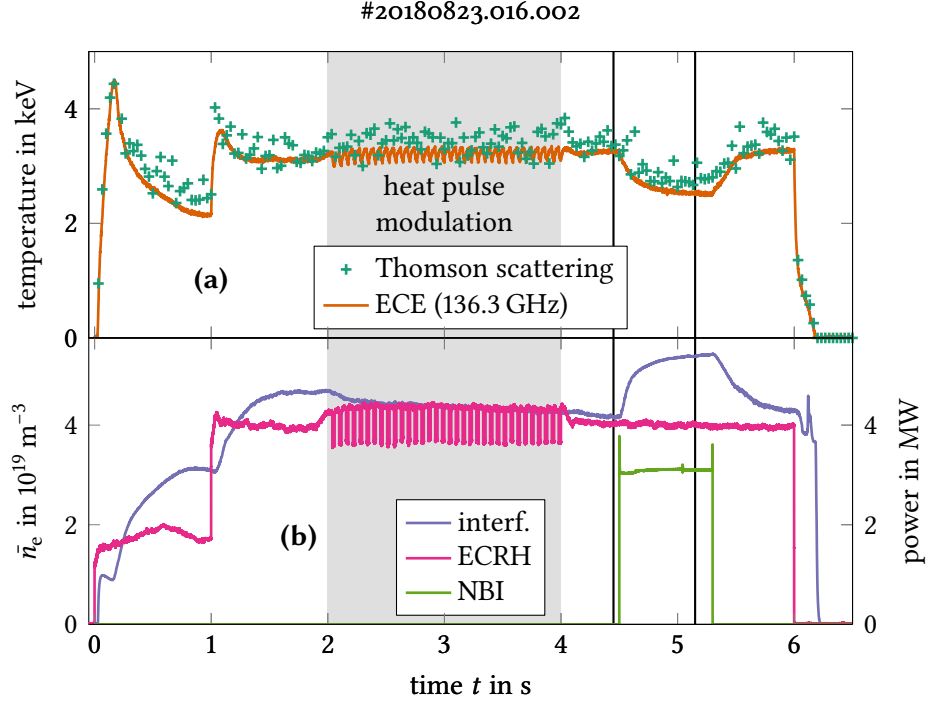
#20180823.016.002

Figure 5.4.: Shown is the plasma discharge also shown in the other plots in chapter 5. **(a)** shows the calibrated signal of an ECE channel close to the core, compared with a central channel from the Thomson scattering system [47]. A comparison to the data at 5.15 s, indicated by another vertical black line, is done in section 5.2.3. **(b)** depicts the ECRH and NBI heating power as well as the line averaged density from the single channel dispersion interferometry. The vertical black line at 4.45 s indicates the time at which the analyses shown in figures 5.5 to 5.8 and 5.11 to 5.14 was performed.

was switched on, providing roughly 3 MW more input power. In addition, the NBI acts as a considerable particle source. This leads to an increase in the line averaged density to circa 5.4×10^{19} m^{-3}. Note that not all additional particles stem from the NBI beam itself, but also from the beam duct connecting the NBI source to the plasma vessel. Consequently, these ions carry less energy. Also, the electrons contributed from the NBI beam are rather cold. In sum with potential heat transport changes due to the increased power that are currently under investigation, the electron temperature dropped to around 2.5 keV. The

#20180823.016.002 at 4.45 s
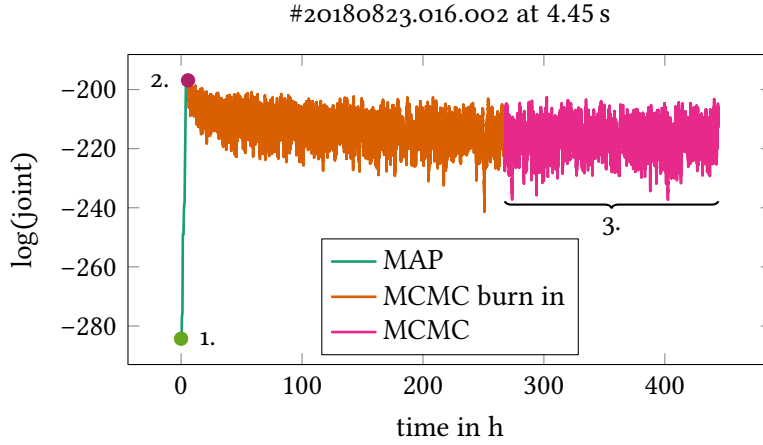
Figure 5.5.: The logarithm of the joint (i. e. prior times likelihood) of the profile inversion model with fixed hyperparameters for the whole analysis. Finding the MAP took around 4.4 h. One can see that more than 100 h are required to get the MCMC reasonably stationary. 1. corresponds to the initial value in figure 5.6b), 2. to the MAP result and 3. to the MCMC samples.

electron temperatures and densities were roughly stationary around 5.15 s, at which the second profile for this chapter has been analysed. The effects of NBI heating on the plasma parameters are briefly discussed in section 5.2.3. Finally, the discharge ends after the NBI has been switched off at 5.3 s and the electron density and temperature were relaxing back to their pre NBI heating values.

The quality of the analysis routine is described before discussing the analysis results of the spectrum taken at 4.45 s. Figure 5.5 shows the trace of the logarithm of the joint (i. e. prior times likelihood) probability density of the model during the evaluation, which serves as a quantification of the quality of the overall fit. The logarithm of the joint probability density is calculated by taking all probability nodes in the model, free and observed, and adding up the logarithms of their individual probability densities. It is clear that the MAP inversion is useful to move the initial guess of the MCMC to reasonably likely values. One can see that the number of samples used for the burn-in of the MCMC, which is used to optimize the proposal distribution adaptively as well as the location in the probability landscape, is chosen reasonably large with 100 000 samples. By eye, the chain looks converged. After the burn-in, the adaptive proposal adapter was deactivated and 30 000 iterations of the MCMC have been used to

estimate the uncertainties shown in figures 5.6 to 5.8 and 5.11 to 5.14. However, a total analysis time of more than 400 h is a considerable drawback, making this analysis unsuitable for a quick analysis between discharges. This highlights the need for optimization on the TRAVIS side, which will be parallelized on a frequency as well as on a ray basis, potentially providing a 32 fold speedup. Another twofold speedup could be achieved by parallelizing the calculations for the X mode and O mode, requiring a webservice with a properly implemented enterprise service bus (ESB). The inversion methods, namely the MCMC, can also be optimized in various ways, for example by multilayer MCMC sampler adaption showing substantial speedups in benchmarks [68].

5.2.1. Plasma profiles and observations

The observed spectrum is compared to the predicted X mode and O mode measurements in figure 5.6. This is the most crucial plot of the analysis, as it allows the clearest possible assessment of whether the result is satisfactory. In figure 5.6a) a direct comparison of the predicted radiation temperatures for the X mode and O mode to the measured values is shown. For the X mode one can see that for most of the frequencies the predictions are very close to the observed values, thus the O mode contribution is small, which was expected as the radiometer was setup to predominantly measure X mode. A notable deviation can be seen above 157 GHz, where TRAVIS does not find a resonance within the used line of sight through the plasma. The fact that the observed values are larger than zero in that frequency region are probably due to ECE of other toroidal locations and cannot be easily modelled. For the channel located at 145.76 GHz, a seemingly systematically to high value was observed throughout this day. No reason for this behaviour is currently known. The prediction uncertainties for the channels around the plasma core near 135.6 GHz have uncertainties below 10 % of their radiation temperature value. Nevertheless, the initial guess determined a length scale for the Gaussian process that leads to the relatively smooth spectrum in that frequency range, even though this means some of the observed values differ by more than a standard deviation from the predicted values. If the hyperparameters were free in the analysis, the MCMC samples would most probably be more spread out, so the width of the sample spread shown here is a lower bound. The O mode reaches its maximum contribution near 132 GHz. While the MAP inversion result yielded no contribution from O mode, the MCMC predicts a typical O mode contribution of few tens of eV.
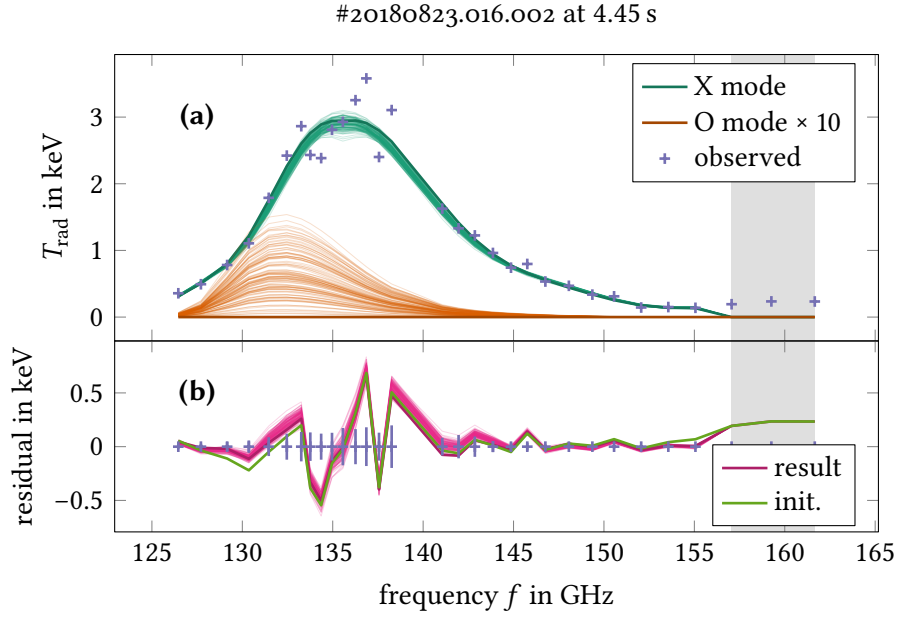
#20180823.016.002 at 4.45 s

Figure 5.6.: (a) shows the comparison of the predicted and observed T_{rad} spectra for the first vertical black line in figure 5.4. The thick, darker lines represent the MAP inversion result, while the thin lines are MCMC samples to give an idea about the uncertainties. The remaining residuals are shown in **(b)** with the predicted uncertainties in marine. TRAVIS could not find any resonance for the channels with grey background, so they should be neglected. That the ECE signal for these channels does not completely go back to zero is probably due to multiple reflections at the wall.

In order to be able to better assess systematic effects, figure 5.6b) shows the residuals of the combined X mode and O mode prediction together with the prediction uncertainties mainly dominated by the calibration uncertainties. The green line reflects the residuals of the initial guess, while the purple lines stem from the MAP (thick) and MCMC inversion (thin). One can see that there are no huge differences between the initial guess and the MAP or MCMC inversion result. For analyses that take place between discharges in the control room, the quickly available initial guess thus offers a good approximation for the complete Minerva analysis. Below 131 GHz and above 147 GHz the residuals of the inversions show a slight improvement over the initial guess. The only notable deviation in these frequency ranges is the three channels at the HFS for which
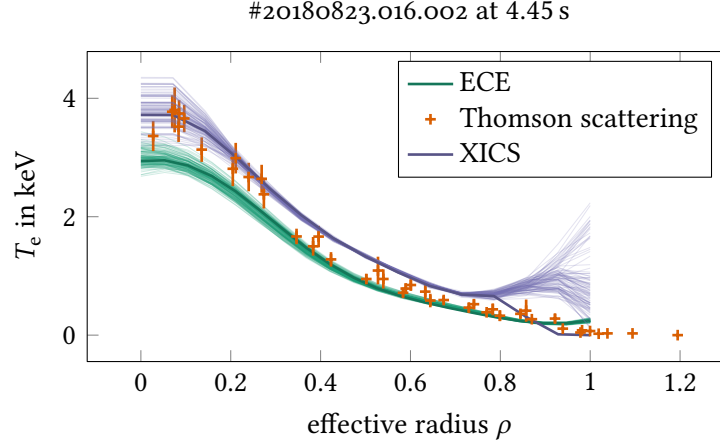
#20180823.016.002 at 4.45 s

Figure 5.7.: The plot shows the electron temperature profile as resulting from the inference done on the ECE model, as well as the values obtained from the Thomson scattering and the X-ray imaging crystal spectroscopy (XICS) [48] diagnostics. Thin lines represent MCMC samples and the darker, thick lines the MAP results. The Thomson scattering positions outside the last closed flux surface are calculated via a simple extrapolation scheme.

TRAVIS, as mentioned above, found no resonances. Besides the systematic deviation at 145.76 GHz, some channels show residuals of up to about 500 eV. A possible explanation might be systematic differences between the invessel part and the calibration unit part of the ECE, although much care has been taken to make the parts as similar as possible, cf. section 2.3.1.

The derived T_e profile is plotted in figure 5.7 along with the results from the Thomson scattering and X-ray imaging crystal spectroscopy (XICS) systems. The ECE samples are taken from the corresponding MCMC run (cf. figure 5.5). An MCMC run on the XICS Minerva model [48] produced the samples for the XICS system. The Thomson scattering data are also evaluated with an MCMC within the Minerva framework, but as only summary statistics are stored on the archive, no samples from the posterior can be shown. The uncertainties of the Thomson scattering system do not take into account calibration and systematic uncertainties, but only the statistical errors. ECE and Thomson scattering show a good agreement between an effective radius of 0.3 and 1. In the core the difference is more notable, but the inferred uncertainties are bigger for both diagnostics as well. Note that the ECE channels in the core show in general a
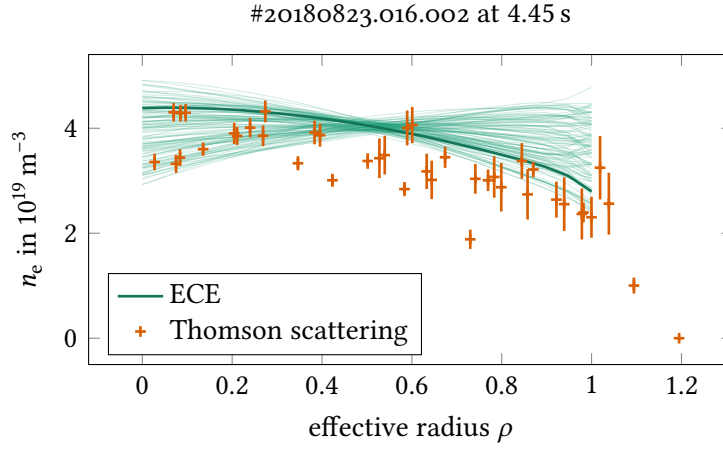
#20180823.016.002 at 4.45 s

Figure 5.8.: The plot shows the electron density profile as resulting from the inference done on the combined ECE and interferometry model, as well as the values obtained from the Thomson scattering system. For the ECE the thin lines represent MCMC samples and the thick line the MAP result. Note that the Thomson scattering system is not part of the model and consequently not part of the inversions.

larger prediction uncertainty, thus, a more flat profile shape corresponding to a lower complexity is found. XICS shows in general systematically higher values than the ECE, but the shape of the profile is similar. However, XICS does match Thomson scattering better in the core. The decrease of the XICS MAP electron temperature outside an effective radius of 0.8 is due to a virtual observation that encourages small values, while the large spread of XICS posterior samples in this region indicates that the XICS raw data contains only a small information content. This is due to the decreasing density of ionization states of the ions being detected by XICS. It is currently not known why there is a systematic deviation between XICS and the other two diagnostics.

The electron density profile as inferred is shown in figure 5.8. Keep in mind that »ECE« represents the combined model of the ECE *and* the single channel dispersion interferometry. As most of the ECE channels are reasonably optically thick and the electron density is far enough away from the cutoff of the ECE channels, the inference of the electron density profile yields a broad posterior. Interferometry ensures that the line integrated electron density remains approximately the same (cf. figure 5.9). The inferred n_e profiles exhibit a rel-
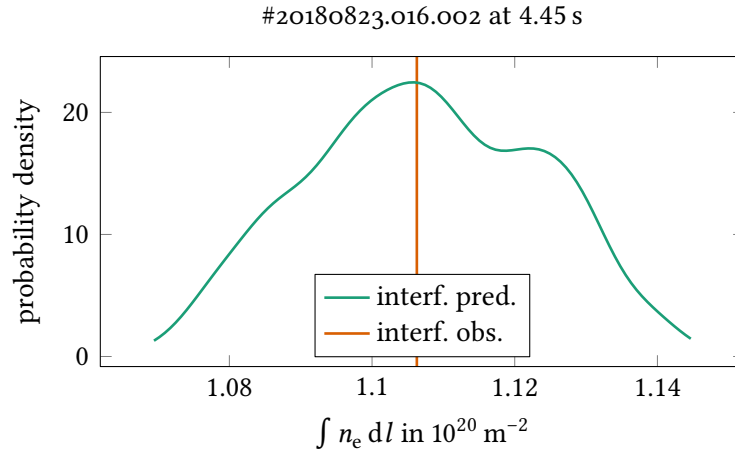
Figure 5.9.: The kernel density estimate (i. e. the marginal posterior) of the interferometry prediction as calculated from a thousand MCMC samples.

atively linear shape as this corresponds to a comparatively low complexity of the profile shape, and due to the hyperparameters being fixed for this analysis. This is the reason that the uncertainties appear to be lower in the middle of the profile. In comparison to the electron density as inferred by Thomson scattering measurements the profile by ECE and interferometry is systematically higher. This discrepancy might be caused by the Thomson scattering laser being misaligned, thus leading to a loss in collected scattered photons creating correspondingly inaccurate absolute electron density measurements including a potentially wrong profile shape. Furthermore, spectral calibration uncertainties of the Thomson scattering system are not taken into account. Chapter 6 discusses in detail the influence that densities above the cutoff of some ECE channels have on the inference of the density profile.

Figure 5.9 shows the marginal posterior estimated by the kernel density estimate of the prediction for the single channel dispersion interferometer, i. e. the electron density integrated along the line of sight of the interferometer. Due to the retroreflector used, the line of sight crosses the plasma twice, thus creating the seemingly high values of the observed value. The uncertainties of the interferometry prediction were fixed at $2 \times 10^{18} \, \mathrm{m}^{-3}$, which is estimated from calibration uncertainties and drifts. Newer estimates suggest a prediction uncertainty of $6 \times 10^{18} \, \mathrm{m}^{-3}$ to be more appropriate. One can see that there is a good agreement between the predicted and observed values.
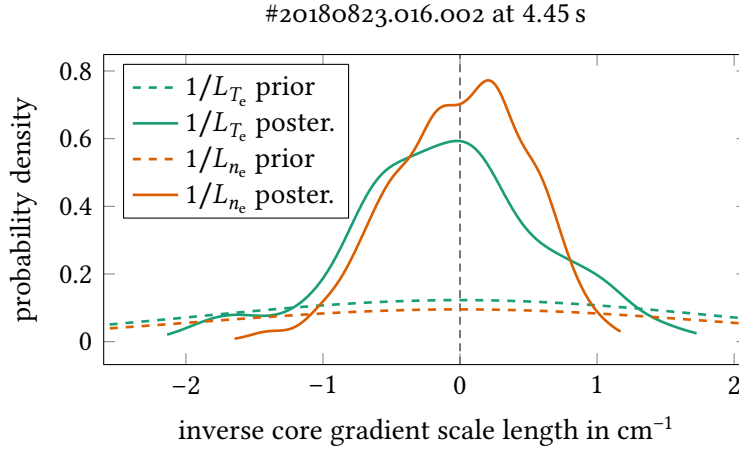
#20180823.016.002 at 4.45 s

Figure 5.10.: The kernel density estimates (i. e. the marginal posteriors) of the
virtual observation predictions as calculated from a thousand MCMC samples.
Note that for the representation the more intuitive inverse core gradient scale
length was chosen over the more technical quantities used in the model.

Figure 5.10 shows the priors and marginal posteriors of the virtual obser-
vation predictions that were used to constrain the gradients in the core to be
close to zero. The gradient was calculated by having one Gaussian process point
directly in the core, and another one very closeby such that the gradient scale
lengths at the core can be approximated by

$$\frac{1}{L_q} = \frac{\nabla q}{q} = \frac{\Delta q}{qa\Delta\rho} \quad \text{with } q \in \{T_\mathrm{e}, n_\mathrm{e}\}. \tag{5.4}$$

The minor plasma radius is denoted with *a*. Since the evaluations in this chapter
have been done the gradient as calculated directly from the cubic interpolation
is used. It should be noted that the gradient prediction has been provided with
a large uncertainty (one standard deviation) of 100 keV or $100 \times 10^{19} \, \mathrm{m}^{-3}$, re-
spectively. It can be seen that the posterior is for both the electron temperature
and density inverse core gradient scale lengths much more narrow than their
corresponding priors. This corresponds to smaller gradients in the plasma core.
Potentially, the priors could have been chosen much more narrow, reducing the
prior-posterior discrepancy. If desired, the uncertainty of the predicted gradi-
ents can be reduced, reducing the inferred gradients at the price of a more
sharply defined model posterior.

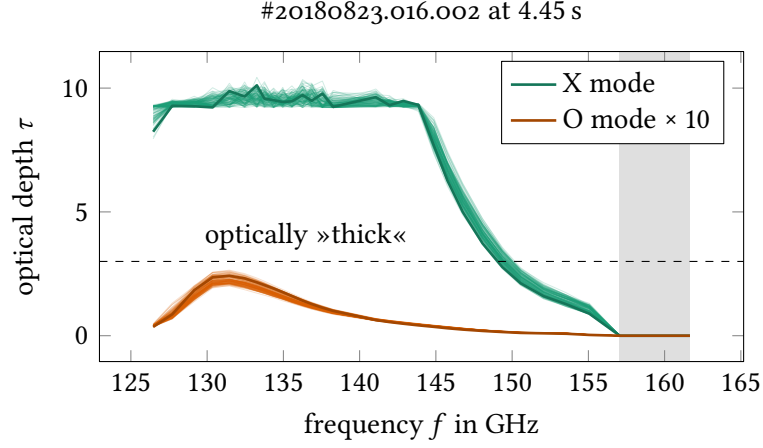#20180823.016.002 at 4.45 s

Figure 5.11.: Shown is the optical depth spectra of the X mode and O mode, the latter one upscaled by a factor of 10 to make it more visible. The thin lines represent MCMC samples and the darker, thick lines the MAP results. TRAVIS could not find any resonance for the channels with grey background, so they should be neglected.

5.2.2. Details of ECE specific parameters

Figure 5.11 shows the optical depth (cf. equation (2.23)) spectra of the X mode and O mode, the latter one upscaled by a factor of 10 to make it more visible. Optical depth is a measure of how much radiation is absorbed. One can see that the optical depth for the X mode reaches values above 3. This is used as a threshold for sufficient optical thickness to consider the plasma opaque at a given frequency. The threshold stems from estimations of the error of $\tilde{T}_e / \langle T_e \rangle$ to be lower than 15 % compared to the corresponding radiation temperature quantities. Here, $\tilde{T}_e$ represents the variation of T_e and $\langle T_e \rangle$ the time average of T_e. For estimations of $\tilde{T}_e$ directly an optical thickness of 0.7 is sufficient [101]. This can give a useful estimation on which channels can be used for the heatwave analysis preceding the timepoints analysed within this chapter. If a proper profile reconstruction for each point during the heatwave modulation would be done, one could avoid using the optical depth dependent radiation temperature spectrum as a surrogate for the electron temperature spectrum altogether. Thus no requirements on the optical depth would remain (but a low optical depth would produce a larger uncertainty in the electron temperature profiles). The

plasma for the O mode is not opaque, indicating that multiple passes of the beam through the plasma should be taken into account. The X mode plateau seen between roughly 127 GHz and 144 GHz at circa 9 is an analysis artifact caused by TRAVIS stopping the ray tracing once no relevant contributions will come from further following the ray. While this can be easily deactivated, it does give some small perfomance gains and is thus usually used. The optical depth of the second harmonic X mode is [36]

$$\tau_{X2} \propto F_X(q) \, T_e \, \omega_{c,0} \, L_B \, (\omega_{p,e}/\omega_{c,0})^2, \tag{5.5}$$

$$F_X(q) = \left(\frac{12 - 8q + q^2}{12 - 4q} \right)^{1/2} \left(\frac{6 - q}{6 - 2q} \right)^2, \tag{5.6}$$

$$q = (\omega_{p,e}/\omega_{c,0})^2. \tag{5.7}$$

For electron densities below approximately $10^{20} \, \mathrm{m}^{-3}$ increases in the electron density will enhance the optical depth τ_{X2}. However, larger densities lead to $F_X(q)$ getting smaller. Due to the direct proportionality a decrease in the electron temperature and the gradient length of the magnetic field L_B will diminish the optical depth. This is the cause of the dwindling optical depth on the ECE HFS at W7-X manifested in figure 5.11. The posterior samples indicate that the X mode optical depth for a channel is uncertain by up to 1. The optical depths of the ECE channels below 149 GHz exceed 3 even taking the uncertainties into account. Thus, they can be considered optically thick (for measurements done predominantly in X mode), which is important to know e. g. for heatwave analyses.

Figure 5.12 shows the electron momentum spectra normalized to the thermal electron momentum at the points where the radiation emission has its center of mass. The electron momentum per unit mass (i. e. a velocity) is defined as

$$u = v\gamma \tag{5.8}$$

with v the electron velocity and γ the relativistic correction. The thermal electron momentum per unit mass is defined as

$$u_{th} = \sqrt{2 T_e / m_e} \tag{5.9}$$

with T_e the electron temperature at the center of mass of the emission and m_e the electron mass. A value of approximately $u/u_{th} \approx 1$ indicates that the
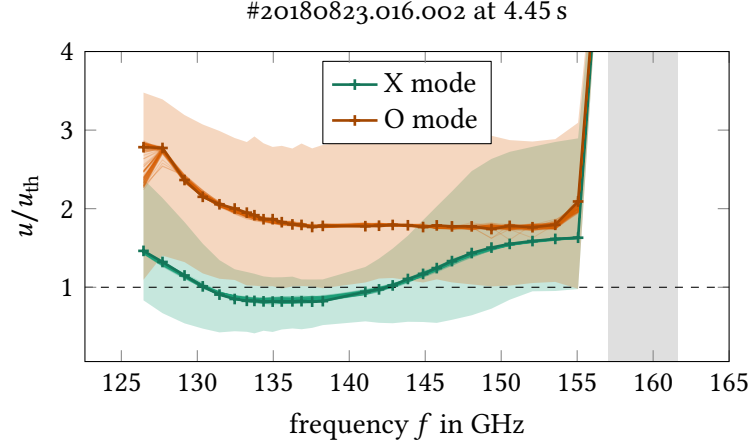
Figure 5.12.: The plot shows the electron momentum spectra normalized to the thermal electron momentum at the point where the radiation emission has its center of mass. The thin lines represent MCMC samples and the darker, thick lines the MAP results. The shaded range (only shown of the MAP result) show the normalized momenta at the lower and upper integration boundary of the radiation emission. TRAVIS could not find any resonance for the channels with grey background, so they should be neglected.

main radiation contribution stems from the thermalized bulk plasma, while higher values hint at nonthermal contributions. One can see that the MCMC samples are hardly visible except for some O mode channels, thus showing that the results are very stable[5]. For the X mode most of the radiation is from thermalized plasma, notably between roughly 130 GHz to 143 GHz. This frequency range is slightly smaller than the one derived from the optical thickness considerations discussed above. Outside of that frequency range the values of the normalized momenta increase up to about 1.6, with an increasing range of normalized momenta contributing to the predicted radiation. O mode channels show a normalized momentum component ranging from 1.75 to 2.8 at their radiation emission center of mass, indicating that there are considerable contributions from hotter core electrons. Also, the lower and upper integration boundaries show a broad range of normalized momenta contributing, from ap-

---

[5]The shaded range correspond to the normalized momenta at the lower and upper integration boundary (cut where 5 % and 95 % of the emission are contained, respectively) in normalized momentum space, here shown for the MAP result.
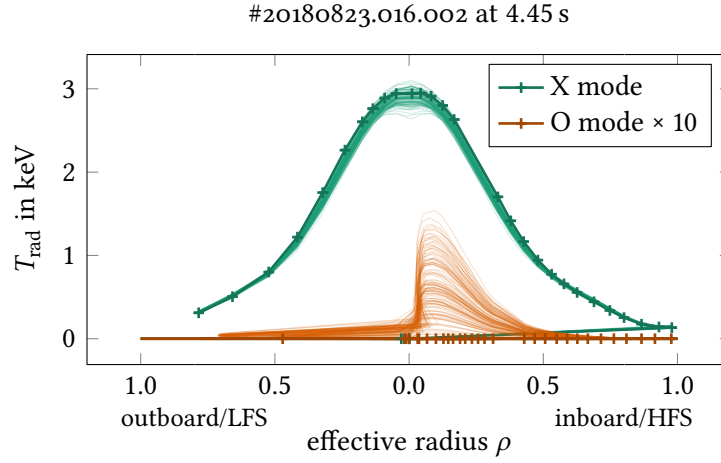
#20180823.016.002 at 4.45 s

Figure 5.13.: The plot shows the MAP predicted X mode and O mode radiation temperatures mapped to the effective radius position at which the center of mass of the radiation emission occurs. The thin lines represent MCMC samples and the darker, thick lines the MAP results. TRAVIS could not find any resonance in the plasma for the X mode channel that »jumped« back to $\rho \approx 0$. LFS represents the outerboard low field side, HFS the inboard high field side.

proximately 1 to 3, making clear that the radiation temperature predicted for a channel should not be directly understood as the electron temperature. Note that these integration boundaries are, as all results shown, calculated for the central frequency only. Taking into account the finite bandwidth of the ECE channels would change accordingly the range of momenta contributing to the signal.

Figure 5.13 shows the radiation temperature profiles mapped via TRAVIS and the best known VMEC equilibrium for that discharge, hence with radiation transport taken into account. The effective radius uncertainties correspond to the effective radius of the lower and upper integration boundary of the radiation emission[6]. One can see that the X mode channels exhibiting »shinethrough«

---

[6]However, they are hardly visible, as these uncertainties are shown for the MAP result only, corresponding to the thick lines. The integrations thresholds are chosen such that the integral over the emission includes 90 % of the intensity, such that the lower integration boundary cuts off at 5 %, and 95 % for the upper integration boundary respectively, cf. footnote 5 on page 84.

on the edge of the low field side (outboard side) have the main intensity contribution coming from the core, although the broad region from which they collect radiation prevents a reasonable localisation. The O mode predictions show for most low field side channels the major contribution coming from the core with bad localisation. Taking multiple plasma passes into account would further impair the localisation due to the wall being nonperpendicular to the line of sight, as the reflected ray potentially will collect radiation originating from the resonance layer at a different effective radius. This mapping was calculated only for the central frequencies, larger bandwidths correspond to worse localization as radiation from a broader part of the plasma is collected.

The difference between the emission positions as determined from the cold resonance (cf. equation (2.3)) and by TRAVIS is shown in figure 5.14a). One can see that for the X mode both ways to determine the position yield similar results. Deviations are still visible, which correspond to the 1 cm to 2 cm that the emission is coming from behind the cold resonance, cf. section 2.1 and figure 5.14b). For many applications it can be expected that the cold resonance mapping is sufficient for a rough guess of the ECE channels emission position. This is also important for the analysis in this chapter, as too large deviations would prevent the ECE from being used to determine the initial electron temperature profile guess. In principle, one could use the electron temperature profile from the Thomson scattering system as a replacement, but as it has a sampling rate of at most 30 Hz and collects the intensity only over few nanoseconds, changes of the equilibrium profile on the order of few milliseconds might not be visible in the Thomson scattering profile. The O mode emission of the low field side comes dominantly from the core, virtually approaching the cold resonance position on the inboard side with vanishing electron temperatures. However, both modes were calculated using only one 3D plasma pass. As the O mode has a low optical depth (considering one plasma pass), these effective radius values only possess indicative character. Large deviations can occur e. g. for plasmas with a very low or very high electron density. The former requires, due to low optical depth, reflections at the wall to be taken into account and multiple plasma passes, effectively collecting radiation from potentially many different plasma locations. The latter might lead to refraction such that radiation from a completely different part of the plasma might be collected. Note that the channel marked with a grey background in figure 5.14b) shows a deviation of approximately 1.25 cm probably due to numerical inaccuracies of the magnetic field strength as calculated by VMEC near the core, as only few flux surfaces are calculated in the
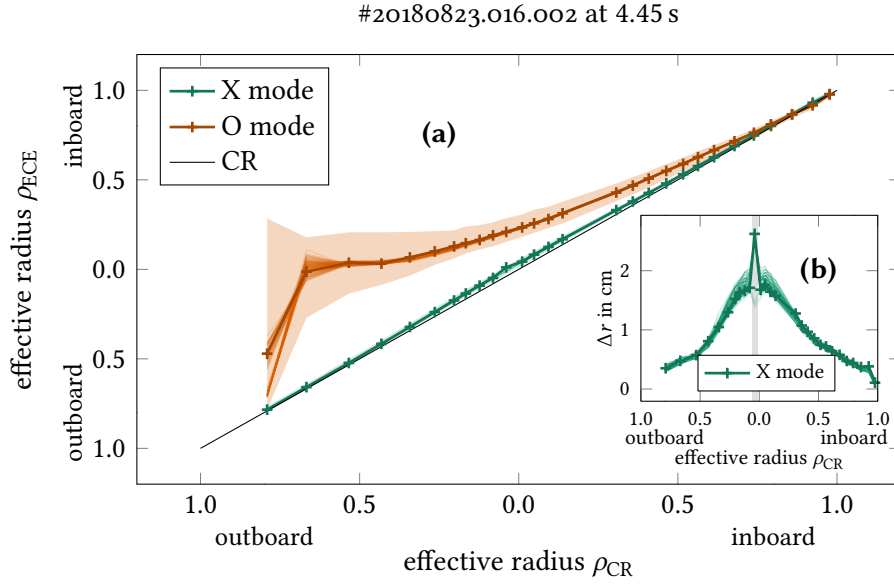
#20180823.016.002 at 4.45 s

Figure 5.14.: (a) shows the deviations of the ECE channel positions as calculated from the model compared to the cold resonance (CR) positions. One can see that the X mode cold resonance is a good estimate for the emission position calculated by TRAVIS. The broad O mode low field side (outboard) emission stems largely from the plasma core, the high field side (inboard) approaches the cold resonance but measures very low intensities. **(b)** shows the difference of the X mode between the model and cold resonance channel positions Δr, cf. section 2.1. The grey area in the core is numerically unstable. The thin lines represent MCMC samples and the darker, thick lines the MAP results.

vicinity (in effective radius) of the magnetic axis. The factors, with which the calculated X mode and O mode intensities are scaled to add up to the final T_{rad} prediction, are free parameters due to the aforementioned reasons and small inaccuracies of the polarisation setting of the ECE system with respect to the inclination angle of the magnetic field angle in the used VMEC equilibrium.

The marginal posterior probability densities of the scaling factors for the calculated X mode and O mode intensities (called »X mode contrib.« and »O mode contrib.« in figure 5.1) are depicted in figure 5.15, where they are called »X mode« and »O mode«. As only one plasma pass with one central ray was used in the calculations, the contribution for the O mode is an upper boundary. If multiple plasma passes would be used, the predicted radiation intensity would increase,
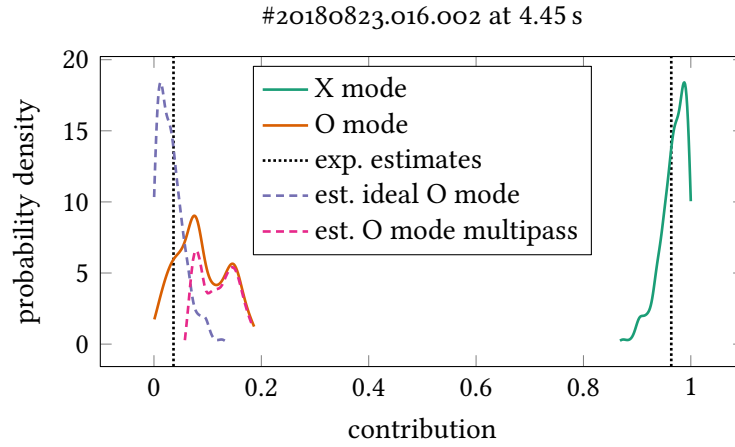
Figure 5.15.: The kernel density estimates (i. e. the marginal posterior) of the X mode and O mode contributions as calculated from 10^3 MCMC samples.

therefore allowing for a smaller O mode scaling factor. Thus, these free parameters allow to estimate how large of an effect the single plasma pass restriction has. This is why the X mode and O mode contributions have been chosen to not necessarily add up to 1. However, the X mode is for the most part optically thick. Hence, the X mode scaling factor is approximately the same as one would expect if the TRAVIS prediction is correct (in the sense that no intensity is »lost« due to missing reflections, etc.) and scaled down according to the polarisation matching of the ECE diagnostic. This also allows to estimate the O mode factor one would expect if multiple plasma passes were taken into account, see the marine dashed curve (»est. ideal O mode«). One infers the most probable contribution of the X mode to the overall intensity to be circa 98.8 % with the center of mass at 97.5 %. From physics considerations (i. e. the position of the wire grid selecting the mode) one would expect the O mode contribution to be lower than 4 % and the X mode contribution higher than 96 %, cf. the vertical, dotted black lines and section 5.1. As there is a considerable probability density mass of the O mode marginal posterior above 4 %, it is clear that the reconstruction would profit from calculations including multiple plasma passes. While this poses no technical problem and is fully implemented, it is computationally expensive to take multiple plasma passes into account, thus slowing down the inference, cf. section 5.1.1. Another prohibitive factor is to get the line of sight accurately after multiple reflections, which is practically not possible.
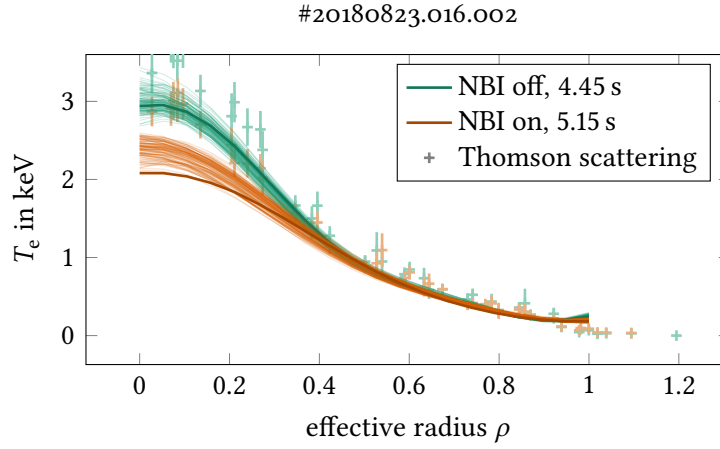
Figure 5.16.: Comparison of inferred T_e profiles with NBI heating on and off. The thin lines represent MCMC samples and the darker, thick lines the MAP results. The chosen timepoints are indicated by vertical black lines in figure 5.1.

5.2.3. Effects of NBI

In this subsection, the effect of NBI heating on the electron temperature and density profile is discussed. An overview of the discharge with the two time points being compared is shown in figure 5.4. The analyses were done while the plasma was reasonably stationary, i. e. no big changes in plasma temperature or density were occuring. Figure 5.16 shows the electron temperature profiles as inferred before and after NBI heating was active. Outside of an effective radius of 0.4 there is virtually no difference between the two profiles. In the very core, the electron temperature decreases by about 1 keV when the NBI heating is active, although the NBI increases the total heating power by about 3 MW. This is caused by the NBI heating also providing notable amounts of particles, thus increasing the density. Note that the difference between the MAP result and the MCMC result at 5.15 s hints at a non-Gaussian posterior, that is most of the probability mass is found at electron temperatures above the MAP solution.

The density profiles as inferred are shown in figure 5.17. As expected the electron density profiles show large uncertainties, which is to be expected as neither »shinethrough« nor cutoff effects play a role that would increase the information that the ECE measurements contain with respect to the density profile. The total electron density increases when NBI heating is switched on,
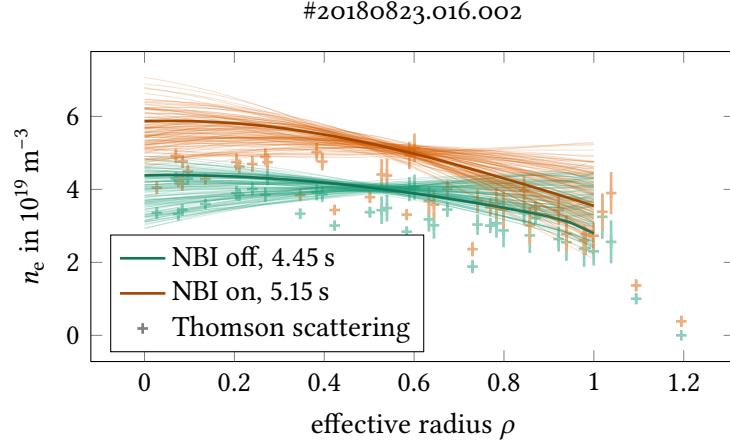
#20180823.016.002

Figure 5.17.: Comparison of inferred n_e profiles with NBI heating on and off. The thin lines represent MCMC samples and the darker, thick lines the MAP results. The chosen timepoints are indicated by vertical black lines in figure 5.1.

as it acts as a particle source. A slight steepening of the profile is observed with the NBI being active. Note that the apparent smaller uncertainties around an effective radius of 0.5 stems from the combination of a (more or less) fixed total electron density by the interferometer and fixed hyperparameters, such that near $\rho = 0.5$ a kind of »pivot point« forms. Chapter 6 discusses the effects that ECE channels in cutoff have on the inference of the electron density profile.

5.3. Conclusions

The holistic evaluation scheme presented in this chapter showed an improvement over the traditional evaluation schemes. Compared to the spline based integrated data analysis (IDA) analysis done at AUG (cf. [102]), the approach in this chapter is more generic as Gaussian processes put less constraints on the profile shape. Furthermore, TRAVIS works both for stellarators (using a VMEC equilibrium) and tokamaks (using an EFIT equilibrium or providing the necessary quantities directly) and can handle wall reflections in full 3D, adding to the generality. Compared to the Minerva approach undertaken at JET, the model presented in this chapter allows the full 3D treatment required for a stellarator, adding to the generality of the approach.

As the time an evaluation takes is too long for many practical purposes, ways to speed up the evaluation are highly anticipated. The plan is to take advantage of the generic Minerva model structure to train a neural net in a similar fashion as was done for X-ray imaging crystal spectroscopy at W7-X [78]. The author regards optimizations specific for this ECE model as problematic, logically optimizations should be done either within TRAVIS to speed up forward calculations or within the inversion methods of Minerva.

A possible improvement of the interpretation/modelling of the ECE could be achieved by using a viewing dump, as noted by [101], as reflections and mode scrambling effects at the wall could be removed which in turn eases the forward calculation. Given the steady-state requirements of W7-X, the feasibility of a microwave beam dump is problematic with the current line of sight, as the target tiles of the ECE suffer already potentially from high heat loads.

Appendix C shows part of the code, or rather keywords, that allow customization of the developed standalone Java executable. The model detailed in this chapter is used to analyse special cases with very low and high electron densities in the next chapter, which encompass »shinethrough« effects in the former and partial cutoff effects in the latter case.

6

Higher level applications of ECE

The inference of the electron temperature and density profiles discussed in chapter 5 extended the domain in which an ECE analysis is possible to regions with more complex radiation transport behaviour. Examples are i) the reconstructions at densities approaching the ECE cutoff (at W7-X: refraction starts to play a role above approximately 8×10^{19} m^{-3} and ultimatively reflection at densities above approximately 12×10^{19} m^{-3}) and ii) the reconstructions with low electron density (at W7-X: below approximately 10^{19} m^{-3}) as redshifted emission of strongly relativistic core electrons starts to get relevant for some channels. Extremely high electron temperatures and distortions of the Maxwellian electron energy distribution function are expected to also hamper analysis considerably. In the following, section 6.1 will explore ECE spectroscopy at high densities as a direct application of this analysis to W7-X high density discharges, a major goal of this project. Section 6.2 will shortly present and discuss findings at low plasma densities, while section 6.3 will give an outlook on possible Bayesian high level physics approaches, here specifically for a heatwave analysis to obtain heatpulse diffusivities.

6.1. Application on a High Density Plasma

This section contrives an example analysis of a measurement, with some ECE frequency channels exhibiting cutoff effects. The improvement of the uncertainty of the inferred electron density profile due to the channels being close to their cutoff is assessed. The limitations and complications coming with the analysis of this edge case are, amongst others: i) the line of sight may no longer be a straight line with the magnetic field strength increasing monotonically

along it, ii) the lines of sight might differ substantially between the X mode and the O mode, iii) the lines of sight for different frequencies can differ significantly and iv) the strong influence of the electron density profile on the optical thickness.

The profile inversion model explained in chapter 5 was used to analyse exemplarily a plasma discharge with electron densities high enough to achieve (partial, i. e. limited to a few channels) second harmonic X mode cutoff for the ECE channels. This particular experiment, #20171115.039.002, was the first time that a high performance plasma was sustained by 2^{nd} harmonic O mode ECRH alone [35]. An overview of the discharge showing the traces of the electron temperature, density and heating is shown in figure 6.1. After 2 s all gyrotrons were heating in O mode. This plasma scenario reached one of the highest diamagnetic energies in OP 1.2a, around 800 kJ at around 4.5 s. One can see that the agreement between a Thomson scattering channel and an ECE channel, both measuring around an effective radius of 0.1, is reasonably good throughout the lower density parts of the discharge. A possible reason for remaining deviations could be found in the VMEC equilibrium, as one single equilibrium was used throughout the whole discharge. It is optimized for the plasma at the high density phase, around 3.75 s. Thus, the ECE channel best suited to match a corresponding Thomson scattering channel might differ in the course of the discharge. Via a pellet injection system [49, 103] a density ramp is driven, such that the selected ECE channel shown in the graph is pushed into cutoff around 2.85 s, with its signal returning at around 4.15 s. The remaining ECE signal can stem from multiple origins, amongst them: i) a small chunk will originate from O mode contributions due to the inherently imperfect selection of the X mode due to the finite polarization mismatch at the last closed flux surface, ii) the refracted ECE line of sight might have to consider Doppler broadening as the ray is no longer perpendicular to the flux surfaces, iii) potentially a different region in normalized momentum space is sampled and iv) the X mode might get reflected from the cutoff layer in the plasma, again reflected at the vessel wall, and, due to the line of sight now »looking« at another part of the plasma where the resonance surface is at another electron density, radiation intensity can be collected again. Note that in a stellarator, due to its helical axis and magnetic mirror fields, a toroidal shift of the ECE line of sight changes the relation between observation frequency and observed flux surface. The Thomson scattering system delivers electron temperatures reliably even if the ECE is in cutoff as expected. The huge increase in particles through pellet injection lowers the
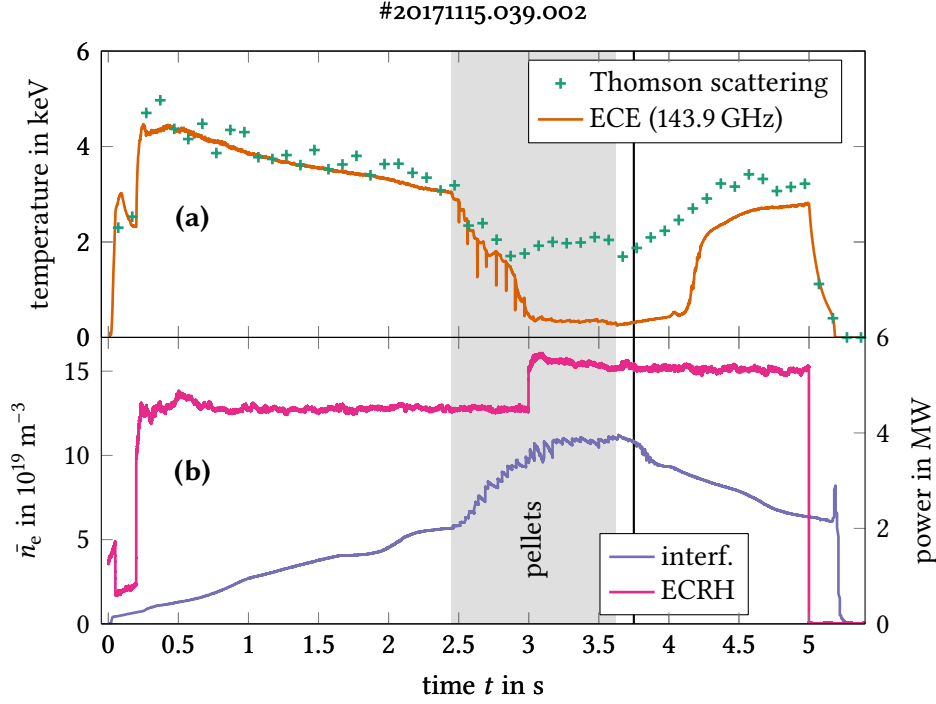
Figure 6.1.: High performance plasma discharge with 2^{nd} harmonic O mode ECRH. **(a)** shows the radiation temperature of an ECE channel close to the core, compared with a central channel from the Thomson scattering system [47]. The vertical black line indicates the time at which the analysis shown in section 6.1 was performed. **(b)** depicts the ECRH heating power as well as the line averaged density from the single channel dispersion interferometry. No other heating method was used.

electron temperature from about 4 keV to about 2 keV, as the ECRH power is increased only from circa 4.5 MW to 5.4 MW[1]. A profile inversion of the combined ECE interferometry model was performed at the time point indicated by the vertical black line. The electron density decrease during the post pellet injection phase offers a plasma with less abrupt changes when compared to the pellet injection phase in which the dynamics of the pellet fuelling complicate

---

[1]One can see the power drop at around 100 ms. This were two gyrotrons in X mode for plasma startup, which were shut down to change their polarisation to O mode. They came back online at 3 s.
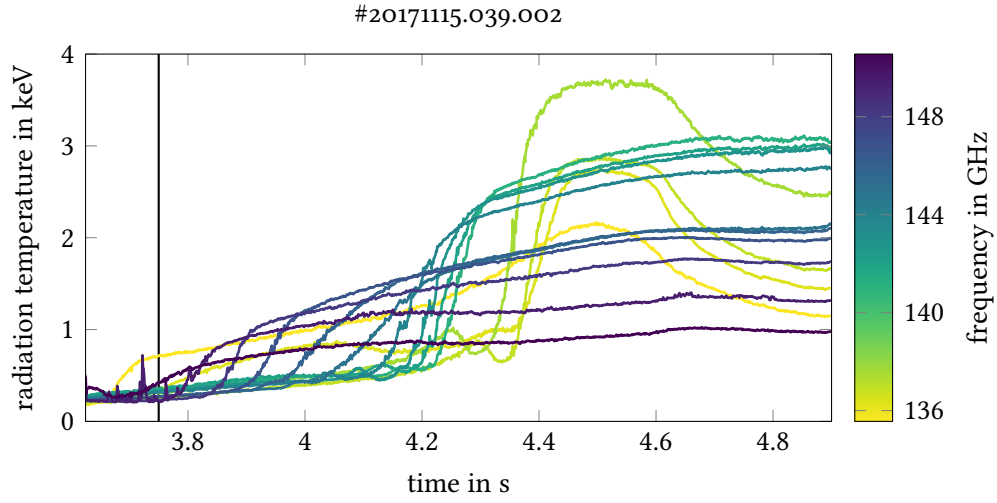
#20171115.039.002

Figure 6.2.: ECE channels coming out of cutoff during the density decrease in the high performance discharge #20171115.039.002, see figure 6.1. One can see that first a channel with a low frequency comes out of cutoff, followed sequentially by high frequency channels with decreasing frequencies. The black vertical line indicates the time point of the profile inversion.

the plasma behaviour. Hence, the former is better suited for an assessment of how much electron density profile information is concealed in an ECE spectrum in partial cutoff.

Figure 6.2 shows the timetraces of the ECE channels between 135.5 GHz and 150.6 GHz during the time window in which the electron density drops gradually from around 1.09×10^{20} m^{-3} to around 0.65×10^{20} m^{-3}, until no ECE channel is in cutoff anymore. The black vertical line indicates at which timepoint the profile inversion as described in this section was ran. One can see that the first channel to come out of the cutoff is the channel with the lowest frequency. However, subsequently channels at the *upper* end of the displayed frequency range drop out of the cutoff first. Two factors play a role for this to happen: i) at high densities the line of sight might be deflected strongly such that the resonance layer is never hit (at least on the first plasma pass), even though the density at the resonance layer is *below* the cutoff density and ii) the density profile is not just a box, therefore it does play a role at which effective radius a channel is collecting radiation as further out the density tends to be notably lower.
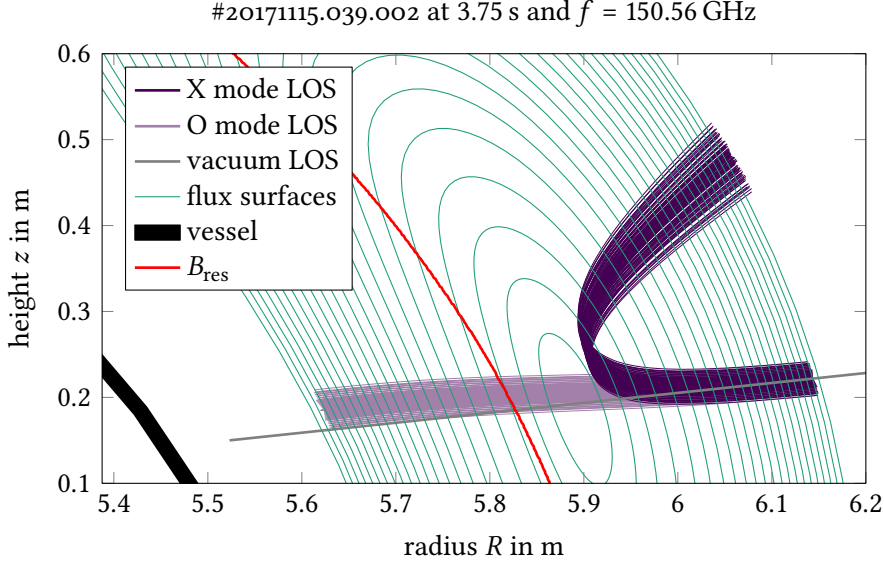
Figure 6.3.: Ray traced ECE channel with an electron density profile that reaches values close to the cutoff density of the channel. The channel corresponds to the highest frequency channel shown in figure 6.2. To account for the finite beam width the beam is split in multiple rays.

To understand these phenomena it is helpful to visualize the X mode and O mode lines of sight for a specific frequency in the experiment investigated in this section. Figure 6.3 shows exemplarily how the beam of channel 26, split up in 120 rays for each the X mode and the O mode, propagates through the plasma. It should be noted that the X mode rays are deflected not only along R and z, but also along the toroidal direction, ϕ. A widening of the beam can be observed for the X mode, while the O mode does not show a strong deviation from the vacuum line of sight. One can imagine now, that, if the frequency bandwidth is large enough and the electron density is close to the point where the ray gets reflected, small changes of the electron density can lead to drastically different lines of sight of the rays representing the beam. Even for a single frequency, the ECE signal does get highly sensitive to small electron density changes under these circumstances. Therefore, it is very difficult to obtain a satisfactory initial guess, as the ECE itself cannot be used in such high density scenarios for the initial guess of the electron temperature; and the Thomson scattering system at W7-X provides an electron density and temperature profile only every 33 ms to
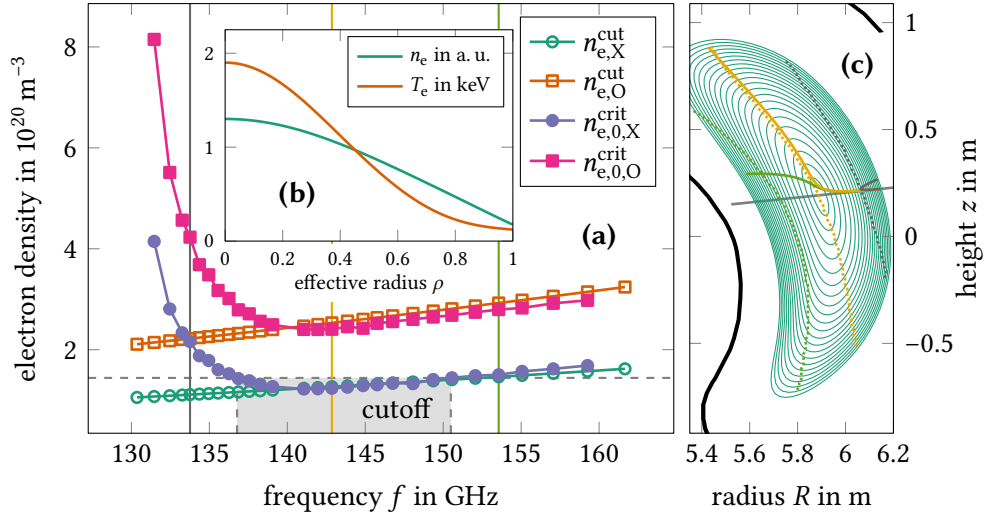
Figure 6.4.: (a) shows the critical central electron densities and cutoff densities (cf. equation (2.5)) for the ECE at W7-X for profile shapes shown in **(b)** corresponding to #20171115.039.002 at 3.75 s where many channels are deflected due to cutoff effects. To indicate the central density reached in this discharge a grey dashed line is used, which also signifies the ECE X mode channels to be expected in cutoff. The line of sight of the three colored vertical lines in **(a)** is shown in **(c)**.

100 ms. This in turn leads to problems in the analysis; the MAP often gets stuck in a local maximum and the MCMC will take an inconveniently long time to get its chain stationary.

Figure 6.4a) gives an estimation on which ECE channels can be expected to be in cutoff. The calculations used the profiles shown in figure 6.4b), where the electron temperature profile corresponds approximately to the one of the discharge used within this section, #20171115.039.002, at 3.75 s. For the electron density profile the approximate shape at that time point was taken, with a to-be-fitted factor scaling the density profile up respectively down. Then, the factor (that is, the central electron density, henceforth called »critical central density«) for which the scaled profiles deflect the ray so much that it does not reach the cold resonance surface[2] is determined by a simple optimization scheme. Cor-

---

[2]For the intersection tests a modified Möller–Trumbore ray-triangle intersection algorithm [104] is used.

respondingly, figure 6.4c) shows the flux surfaces for the VMEC equilibrium used in this section, as well as the rays (solid lines) and resonance layers at the magnetic axis (dotted) for three frequencies (low field side: dark gray, center: yellow, high field side: green) of the X mode. The straight, light grey line indicates the vacuum line of sight. While the dark gray ray, corresponding to the vertical line at 133.76 GHz, is reflected back directly, the yellow ray, corresponding to the vertical line at 142.86 GHz is deflected upwards. The yellow ray at the top of figure 6.4c) looks like it intersects the resonance layer, but that is due to the projection onto the (R, z) plane as the ray changes the toroidal angle sufficiently to see a slightly differently positioned resonance layer at this toroidal angle. The same holds true for the green ray (cf. the vertical line at 153.56 GHz in figure 6.4a), which is deflected strongly in the toroidal direction, an effect that is not visible in this (R, z) projection. By using a central electron density circa corresponding to #20171115.039.002 at 3.75 s we can estimate which frequencies of the ECE are in cutoff. This density is indicated in figure 6.4a) by the horizontal grey dashed line. It intersects the calculated critical central densities of the X mode shown in marine at 136.78 GHz and 150.49 GHz, so that the frequencies in between can be expected to be in cutoff, indicated by the grey background. The O mode critical central densities, indicated by the purple curve and filled squares, do not show any intersections, hence no channel's O mode contribution is expected to vanish. Thus, the ECE spectrum for that time is expected to exhibit a crater like structure, being in cutoff only in central channels. It also explains why the channels subsequently coming out of cutoff in figure 6.2 are not monotonically increasing in frequency. Furthermore, figure 6.4a) (or correspondingly newly calculated graphs if different temperature or density profile shapes are required) allows to tailor discharges where certain ECE channels are close to the cutoff, but still provide useful X mode information. Similarly, the graph can be used for estimations which channels do not reach the O mode cold resonance surface, although this is only relevant at extremely high electron densities above $20 \times 10^{19}\,\mathrm{m}^{-3}$.

As expected after these simulations, the observed radiation temperature spectrum shows a crater-like structure, see figure 6.5. The ECE spectrum was averaged over 5 ms. Note that a VMEC equilibrium manually tailored for this particular timepoint was used[3], that uses the experimental coil currents instead of the coil currents of the precalculated equilibrium. However, the pressure pro-

---

[3]EJM, see minerva-vmec-a87872131bcf5422ac2b858a41ceb7e5.
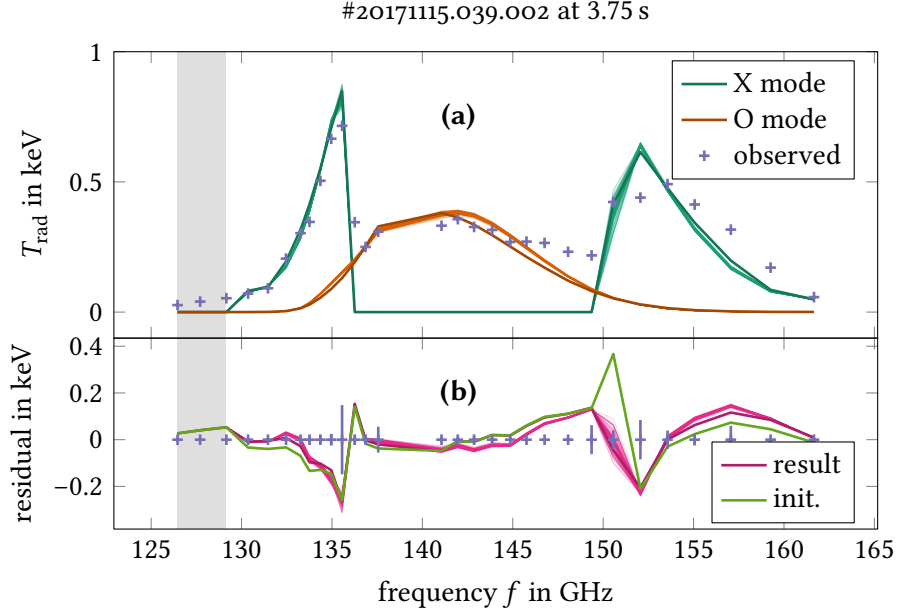
#20171115.039.002 at 3.75 s

Figure 6.5.: (a) shows the comparison of the predicted and observed T_{rad} spectra for the time defined by the vertical black line in figures 6.1 and 6.2. The thick, darker lines represent the MAP inversion result, while the thin lines are MCMC samples to give an idea about the uncertainties. The remaining residuals are shown in **(b)** with the predicted uncertainties in marine. For the channels with a grey background the resonance surface was not hit or their optical depth is below 10^{-5}, so they should be neglected. That the ECE signal for these channels does not completely go back to zero is probably due to multiple reflections at the wall.

files from the reference equilibrium (cf. footnote 5 on page 84) are used, which worked well enough for this case. Furthermore, in contrast to the analysis in chapter 5, one reflection of the O mode at the inner torus wall was taken into account. This was done as the intensity from the X mode was expected to be much lower in the cutoff region, increasing the relative importance of the O mode contribution, which has a single pass absorption at these densities above 30 %. As the calibration factors of channels 15 and 16, corresponding to 138.26 GHz and 139.06 GHz, possess very high uncertainties they were neglected during the analysis. The O mode prediction captures the central channels, in which the X mode is in cutoff and thus not contributing directly, reasonably well, although

with degrading quality above 145 GHz. However, it has to be noted, that the O mode prediction serves as a *proxy* here: If the X mode gets reflected at the plasma, hits the vessel wall and is reflected back to the plasma, it can potentially reach the cold resonance surface, as the reflection might have redirected the line of sight to some other part of the plasma. This implies that even if a channel is in cutoff for the first plasma pass, it might still collect radiation after taking its reflections into account. As no reflections for the X mode were taken into account in the TRAVIS calculations, potential X mode contributions are subsummed in the O mode proxy. Note that taking reflections for the X mode into account is difficult as, depending on the density, large parts of the plasma vessel can potentially act as a mirror due to the high sensitivity of the X mode on the electron density. Hence, more polygons to describe the plasma vessel are needed if an X mode reflection needs to be taken into account when compared to an O mode reflection. In turn, more ray-polygon intersection checks are necessary, which would make an analysis with X mode reflections taken into account slower. The contribution of the O mode proxy was fixed at 100 %, and the values match at least roughly quantitatively. The low field side channels below the ECE channels that went into cutoff are decently predicted by the X mode, although there are deviations especially at the transition into cutoff. On the high field side, where the ECE channels are not yet in cutoff the deviations between predictions and observations are stronger. This might very well hint at the MCMC not being completely converged.

Figure 6.6 shows the trace of the logarithm of the joint (i. e. prior times likelihood) of the inversion at 3.75 s. There are several things to note: i) The MAP failed to find the global maximum, as the logarithm of the joint increased further during the MCMC, ii) even though a burn-in of 125 000 iterations was chosen, it is doubtful that the MCMC completely converged[4]. As a consequence, the results obtained from such a MCMC have to be taken cum grano salis, even though the electron density profile discussed later has proved to be remarkably stable throughout the burn-in and thereafter. Unfortunately, due to multiple server restarts and other events that required to pick up the analysis at the last saved point, it is not possible to use accurate times on the abscissa. The total analysis time was over three months, mainly due to the inclusion of one

---

[4]While the MCMC is very robust and will eventually converge, it is usually slow to find the MAP, especially if the MAP analysis providing the starting position got stuck in a local maximum.
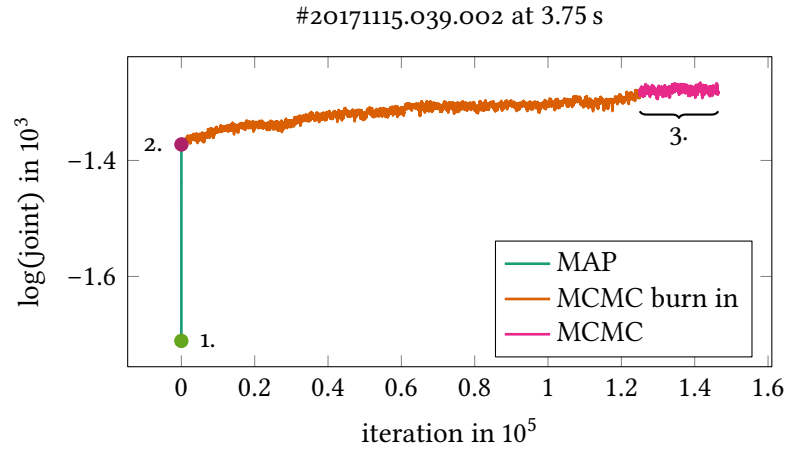
#20171115.039.002 at 3.75 s

Figure 6.6.: The log(joint), where the joint is the product of prior and likelihood, of the profile inversion done for the spectrum in figure 6.5a). The MAP inversion got stuck in a local maximum. Hence, the MCMC continued to climb up the joint probability and did not get stationary. 1. corresponds to the initial value in figure 6.5b), 2. to the MAP result and 3. to the MCMC samples.

(potential) reflection for the O mode. This requires costly collision checks to find the spot at which the ray hits the vessel wall, even though the representation of the carbon wall panel was optimized beforehand to contain only 54 polygons. The semi-automatic optimization was performed by selecting only polygons that with some probability are relevant to reflections, effectively reducing the number of polygons by more than 50 %. Note further that the small hole in the wall panel on the high field side[5] was discarded, as calculations with multiple rays were computationally prohibitively expensive and using the hole might have caused jumps in the signal depending on the distortion of the line of sight[6] (either the ray »disappears« in the high field side horn, or it does get reflected and contributes to the predicted signal if the ray is slightly more distorted and hits the wall around the horn). To address the long evaluation time a Minerva model based on figure 5.1 was developed to be able to generate training data for a convolutional neural network by sampling from the priors [106]. Once trained, the convolutional neural network will help to reduce the analysis time by several orders of magnitude, making the analysis much more feasible.

---

[5]This hole corresponds to the high field side ECE horn [105].

[6]The resulting file can be found here: 🖫.
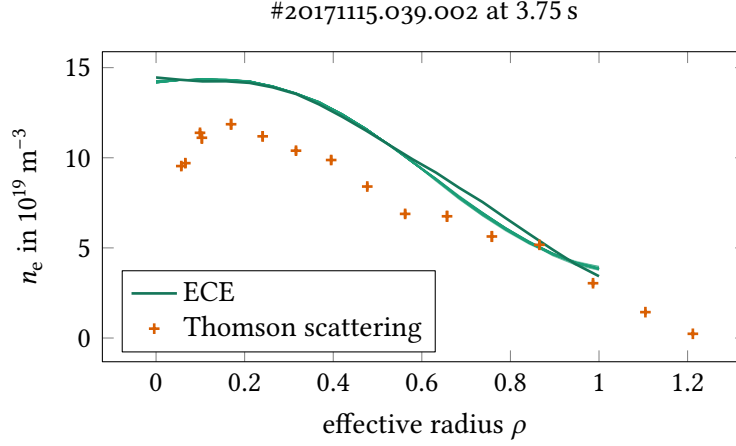
#20171115.039.002 at 3.75 s

Figure 6.7.: The plot shows the electron density profile as resulting from the inference done on the combined ECE and interferometry model, as well as the values obtained from the Thomson scattering system. For the ECE the thin lines represent MCMC samples and the thick line the MAP result.

Nevertheless, the inferred electron density profile as shown in figure 6.7 in particular yielded an interesting result. The MAP and MCMC results are similar, although there is a notable deviation between an effective radius ρ of 0.6 to 0.9. The main difference to the Thomson scattering measurement points is probably again caused by the different line integrated electron densities, which may originate in systematic problems in the alignment of the Thomson scattering system. For the future, it is planned to calibrate the Thomson scattering system with respect to slightly different laser lines of sight, which would allow to interpolate the required calibration more appropriately. Another source of error, however, is the density outside the last closed flux surface, which is currently neglected in the combined ECE and interferometry model. This can lead to an overestimation of the density. There is a distinct contrast between the electron density profile as inferred during plasma scenarios in which neither a partial cutoff nor noteworthy »shinethrough« occurs and the electron density as inferred during partial ECE cutoff; In the former case, a large uncertainty of the electron density profile is found (cf. figure 5.17), while in the latter the uncertainty is reduced manifold. Therefore, the ECE, in combination with the single channel interferometer, can provide useful information about the electron *density* profile, given the right circumstances – which is also the justification for

the extensive and comparably slow analysis of the complex model in figure 5.1. However, it is clearly necessary to increase the speed of a forward calculation to be able to reach a safely converged MCMC within reasonable timespans. For this, the convolutional neural network for which generation of the training data has been started, is essential, although a parallelization of TRAVIS could already provide major speedups. There are two options what one can do with the training set, once available: i) one can use the convolutional neural network trained on the training data to approximate the *forward* model (of the part involving TRAVIS), which is the more flexible option, as the rest of the model and its analysis methods can be changed, ii) one can train the convolutional neural network on the inversion, which would be faster overall, but more locked into the exact model structure. For a first approach the former option is advisable, as one can compare the results directly with corresponding TRAVIS runs and no comparisons with full inversions, which might require an MCMC run and are consequently slow, are necessary.

6.2. Application on a Low Density Plasma

In contrast to the application on a high density plasma as shown in the previous section, this short section focusses on a proof of principle example of a plasma with low electron density. In low density plasmas, »shinethrough« effects (cf. chapter 5) start to play a role and can dominate the intensity for several channels. As a reminder, the effect called »shinethrough« corresponds to the relativistically downshifted core plasma emission that is not (completely) reabsorbed on its way to the ECE horn due to the low optical depth at the emitted frequency. An example spectrum exhibiting this behaviour was shown in figure 4.12. The analysis of this spectrum is shown in figure 6.8. Note that the absolute values of the spectrum shown in figure 4.12 and the spectrum shown in figure 6.8 differ by a small amount, as both the averaging time and the calibration were slightly different. The huge discrepancy at 129.16 GHz is an indication that the MAP did not converge and thus only represents a local maximum of the posterior. Note that the channel at 129.16 GHz is located exactly in the steep edge region of the gradient of the optical depth (grey shaded), such that small changes in the electron temperature and density profile can inflict large changes in the amount of »shinethrough« expected to contribute to the signal. Note further, that the other channel showing a large discrepancy at 136.26 GHz is similarly located in
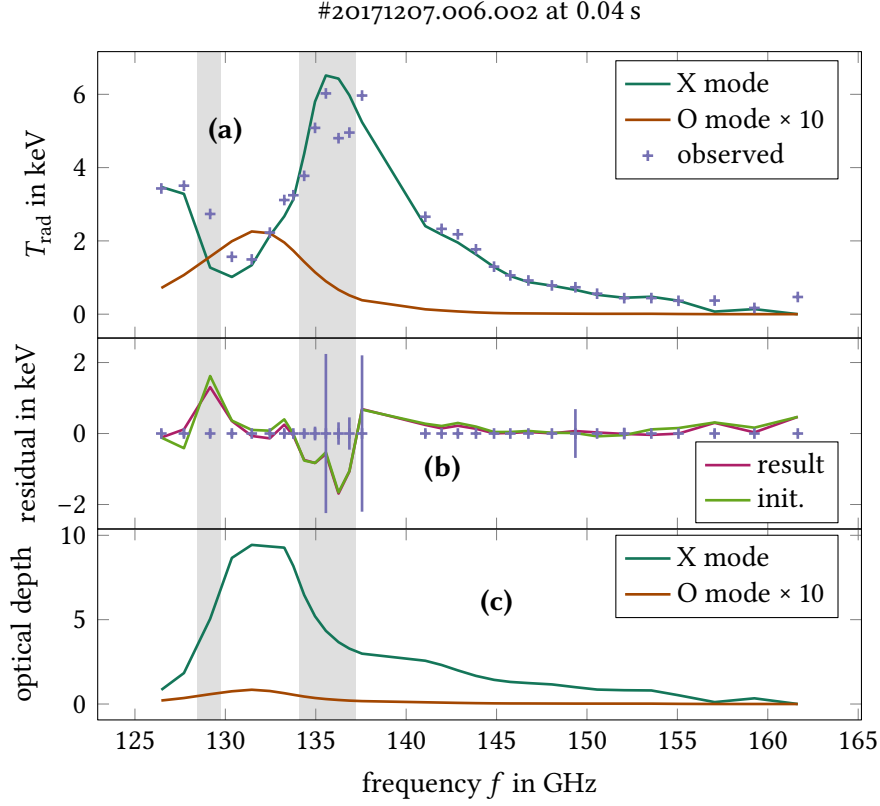
Figure 6.8.: (a) shows the comparison of the predicted and observed T_{rad} spectra for the time defined by the left vertical black line in figure 4.13. The thick, dark lines represent the MAP inversion result. The remaining residuals are shown in **(b)** with the predicted uncertainties in marine. **(c)** shows the optical depth.

the gradient region of the optical depth, again shaded in grey. Inbetween these two channels the plasma is optically thick as $\tau > 3$, cf. section 5.2.2. The optical depth gradient at the side of the channel at 136.26 GHz is less steep than on the other side, hence explaining the smaller, but notable deviations starting to build up from 134.36 GHz onwards. The large sensitivity on the optical depth, which depends notably on the electron density in the shown example, implies that the electron density information content in the ECE signal is comparably large. An MCMC was started, but did not finish in time to make it into the thesis. Therefore, no uncertainties can be provided.

6.3. A Bayesian approach to heatwave analysis

An important quantity to determine the quality of the confinement is the (electron) heat diffusivity χ in $\mathrm{m}^2\,\mathrm{s}^{-1}$. One way to determine this value experimentally is by so-called heat pulse propagation (HPP) experiments [38, 107]. In these experiments, the power of at least one of the heating sources is modulated. The temporally varying heating power at the ECRH deposition region acts as a small local perturbation. This local perturbation causes »ripples« on top of the equilibrated electron temperature profile that propagate through the plasma. These »ripples« are called heatpulses/heatwaves. The speed and amplitude of these pulses yields information about the heat diffusivity, χ^{HP}. A heating method especially well suited for HPP experiments is ECRH, as it allows for a well defined power deposition region within an effective radius of $r_{\mathrm{eff}} = 1\,\mathrm{cm}$. At W7-X, the NBI system deposits the power so broadly and acts so strongly as a particle source that one cannot use it for HPP, at least not for the model presented in this section. Note that the frequency of the modulation determines whether the observations are better suited to infer the heat diffusivity or the power deposition profile of the ECRH. The former corresponds at W7-X to modulation frequencies below 100 Hz (typically about 17 Hz are used), while the latter necessitates frequencies around 500 Hz to be able to explicitly neglect transport effects. Heat pulses originating from modulation frequencies above 100 Hz are damped more and hence the analysis of heatwaves across the effective radius becomes more and more problematic, as less and less ECE channels are capable of detecting the resulting heat pulse. ECE is apt for measuring HPP as the sampling rate can reach up to 2 MHz at W7-X. Note that the low temporal resolution of the other main electron temperature diagnostics at W7-X, the Thomson scattering and X-ray imaging crystal spectroscopy systems, prevents them from observing heat waves in a way that can be easily analysed.

The theory of HPP is a comprehensively developed field[b]. Subsequently, a schematic derivation of a simple equation governing heat transport is shown that closely follows [108, 109]. First, energy conservation in cylindrical geometry mandates

$$\frac{3}{2}\frac{\partial}{\partial t}(n_{\mathrm{e}} T_{\mathrm{e}}) + \vec{\nabla} \cdot \vec{q}_{\mathrm{e}} = \hat{S}, \tag{6.1}$$

in which t is the time, n_{e} the electron density, T_{e} the electron temperature, $\vec{q}_{\mathrm{e}}$ the electron heat flux and $\hat{S}$ the effective source term. Now, the quantities are

split in »background« quantities that do not change with time and a superposed time dependent quantity indicated by a tilde over the variable,

$$n_e = n_{e,0}(r) + \tilde{n}_e(r, t), \tag{6.2}$$

$$T_e = T_{e,0}(r) + \tilde{T}_e(r, t), \tag{6.3}$$

$$q_e = q_{e,0}(r) + \tilde{q}_e(r, t), \tag{6.4}$$

$$\hat{S} = S_{e,0}(r) + \tilde{S}_e(r, t). \tag{6.5}$$

In there, r refers to the radius in cylindrical geometry. Expanding the ECRH modulation depth $\epsilon = \tilde{S}_e / S_e$ in first order allows to find

$$\frac{3}{2} n_{e,0} \frac{\partial \tilde{T}_e}{\partial t} + \vec{\nabla} \cdot \tilde{q}_e = \tilde{S}_e. \tag{6.6}$$

Finally, away from the sources (so that $\hat{S}$ can be neglected) and switching to coordinates of the effective radius ρ yields

$$n_{e,0} \frac{3}{2} \frac{\partial \tilde{T}_e}{\partial t} + \frac{1}{V'} \frac{\partial}{\partial \rho} \left(n_{e,0} V' \langle |\nabla \rho|^2 \rangle \chi^{\mathrm{HP}} \frac{\partial \tilde{T}_e}{\partial \rho} \right) = 0, \tag{6.7}$$

in which $V' = \partial_\rho V$ and $\langle \cdot \rangle$ denotes a flux surface average. The electron heat diffusivity in $\mathrm{m}^2 \, \mathrm{s}^{-1}$ as obtained from HPP corresponds to χ^{HP}.

Figure 6.9 shows the simplified version of the model used to gain radially resolved profiles of the electron heat diffusivity. The model works as follows. First the user provides the time, i. e. the start and end within a discharge. These time points are provided to the ECRH and ECE datasources. Then, the power trace as fetched from the ECRH is upsampled, as it has a lower sampling rate than the ECE radiation temperature traces. Subsequently, a »jump finder« selects each falling flank in the power trace, allowing for a conditional averaging of the temporal dependence of the radiation temperatures of each ECE channel during the ECRH modulation cycles. As the offset is removed in this procedure, the result is taken to represent the time dependent perturbation of the electron temperature. Note that, in principle, a profile inversion as described in chapter 5 would be ideal to be performed for each point in time. However, the computational cost prohibits this approach. Thus, the radiation temperature perturbation is used as a surrogate for the electron temperature perturbation. As the ECE sampling rate is high, even the conditionally averaged signal can
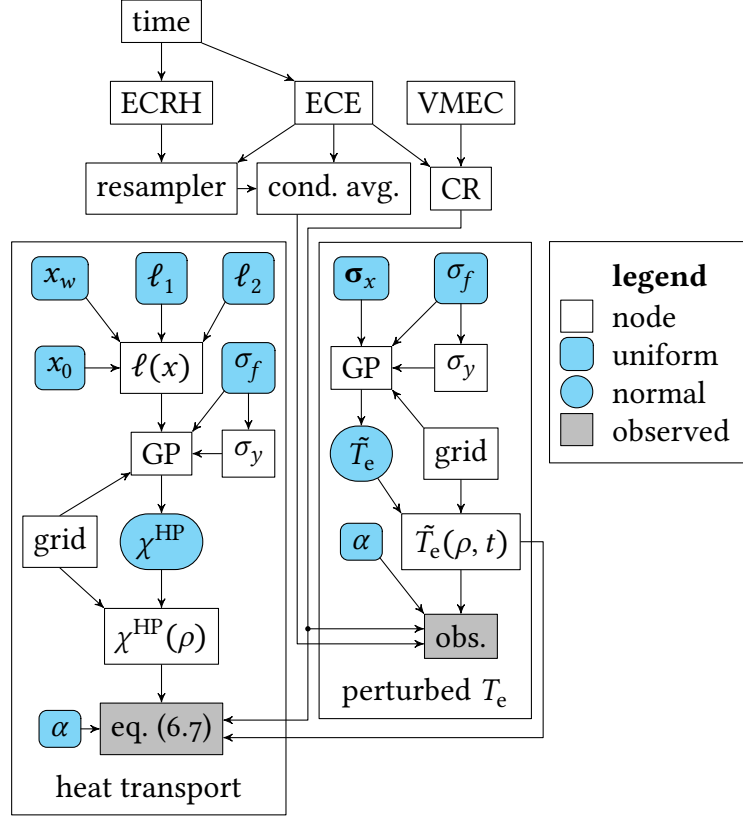
Figure 6.9.: The simplified Minerva heat transport model showing the dependencies used in the model. Free parameters have a cyan background, their shape as specified in the legend corresponds to the prior distribution used. Edges indicate the dependencies between nodes. Note that α and σ_f are different for each subgraph.

contain several thousand data points. The sheer amount of data can slow the analysis down considerably, which is why a downsampled signal is used in the heat transport model.

A 2D squared exponential kernel is used for the Gaussian process to describe the electron temperature perturbation, $\tilde{T}_e$. It requires $\boldsymbol{\sigma}_x$, i. e. the respective kernel scale length in both the effective radius and the time direction, the electron temperature perturbation variance σ_f and σ_y, which is set to 1 ‰ of σ_f and is used to avoid kernel degeneration. Furthermore, a grid provides the points in effective radius and time at which the kernel is evaluated. By default, a 20 × 20

grid is used. The free multivariate normal of the electron temperature perturbation at the grid points uses the 2D squared exponential kernel and a mean of zero. The values (not the mean!) of the multivariate normal are combined with the grid to do a bicubic interpolation, which provides the continuous $\tilde{T}_e(\rho, t)$ function. In the observation node the conditionally averaged signals at the cold resonance positions, determined from the ECE channels' center frequencies and a VMEC equilibrium, are compared to the predicted bicubically interpolated Gaussian process. The prediction uncertainties are based on the estimated uncertainties of the conditional averaging, scaled with a variance scaling factor α. Note that the posterior is sensitive to the chosen interpolation method, as the virtual observation depends on the first and second derivatives.

The Gaussian process describing the heatpulse diffusivity χ^{HP} is structurally the same as used for the electron temperature and density profiles in chapter 5, see section 3.1 for an explanation of the hyperparameters. By default, a grid of 20 points is used. A free truncated multivariate normal using the kernel defined above and a mean of zero allows to infer the heat diffusivity at the grid positions. The continuous $\chi^{\mathrm{HP}}(\rho)$ profile is obtained via a cubic interpolation. This allows to calculate the prediction, which corresponds to the left hand side of equation (6.7). Consequently, the virtual observation corresponds to the right hand side and is consequently set to zero. The prediction uncertainties are based on a fixed value scaled with a variance scaling factor α. To take into account that equation (6.7) is only valid where no heating occurs, the following improvement is suggested: The mirror positions of the ECRH allow, combined with an electron temperature and density profile, to predict the power deposition profile via TRAVIS. As, by default, we use the virtual observations with the same fixed (albeit scaled) value for all radial positions alike, equation (6.7) is enforced with the same strength within the power deposition region as outside of it. The violation of it will be most pronounced within the power deposition region. Thus, one can use the power deposition profile as an additional scaling for the prediction uncertainties of the virtual observation, such that the variance is no longer $\alpha\sigma^2$, but rather $\alpha P_{\mathrm{dep}}\sigma^2$, automatically giving less importance to the virtual observations where the expected violations are strongest. Note that this would, in principle, also allow to infer at least the rough power deposition region by having the variance scaling depend on the effective radius.

As the analysis of the model with all 430 free parameters is computationally expensive, the following evaluation strategy that differs from the approach described in section 3.4.2 is used. First, the free parameters in the heat trans-
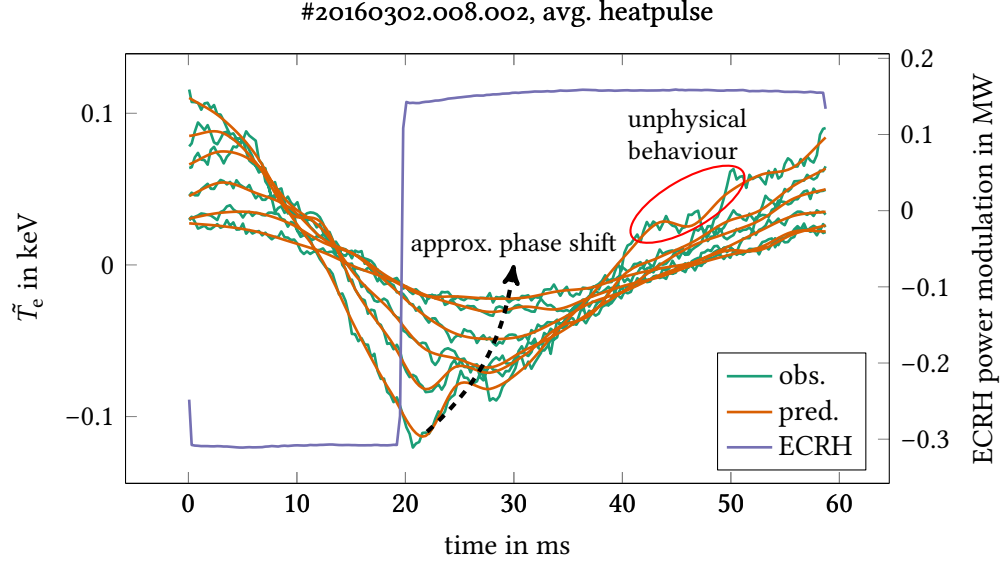
Figure 6.10.: Shown are the observed (from the conditional average) and predicted (from the Gaussian process) electron temperature perturbations for the analysed ECE channels, as well as the conditionally averaged ECRH trace. One can see that the plasma responses of the different channels show phase (cf. the dashed black curve) and amplitude differences. The red ellipse shows an example for unphysical behaviour in the fit.

port submodel are deactivated, such that they keep their momentarily value. The remaining free parameters are all located in the electron temperature submodel. They are linear with respect to the corresponding observation, with σ_f, σ_x and α being treated as hyperparameters. This allows to use the power of the linear Gaussian inversion combined with the Hooke and Jeeves algorithm (cf. section 3.2) operating on the *posterior* probability, cf. section 3.4.2. The default number of iterations amounts to 10^3 and is automatically cut short if 15 sequential iterations yield an improvement of the logarithm of the posterior probability of less than 1.

The average heatpulse responses obtained during an exemplary HPP experiment are shown in figure 6.10. Here, a total of 57 modulation cycles was used for the conditional averaging. Thus, the *x*-axis represents the duration of the average modulation cycle. One can see that the modulated fraction of the ECRH was switched off at 0 ms. It was switched back on at 20 ms, leading to an increase

of the heating power by about 450 kW. Correspondingly, the measurements of the selected ECE channels[7] depicted in cyan show a declining radiation temperature. Channels with larger amplitudes correspond to channels closer to the power deposition region. Noteworthily, a slight shift by about 10 ms in the minimum of the channels can be seen, indicating the delayed response of channels further away from the perturbation. This delay contains information about the heat transport in the plasma. Note that the traces of the selected ECE channels had to be effectively downsampled. In principle, one could use also the raw data directly. As the conditional average uses the mean number of data points per period for the number of bins, the number of predictions grows to such large values that Java cannot handle the arrays without specific implementations of the linear Gaussian inversion, which has not been attempted during this thesis, effectively necessitating the use of downsampled data. The fit on the downsampled data via the method described above provides a decent continuous electron temperature perturbation fit and is shown in orange.

One could assume that this fit is good enough to continue directly with the electron heat diffusivity fit and an deactivated electron temperature perturbation subgraph to keep the dimensionality low. However, some of the features visible in figure 6.10 are unphysical, e. g. the sudden drop in the electron temperature shown in the red ellipse. This leads to deviations from the virtual observation (cf. equation (6.7)) that cannot be avoided by only changing χ^{HP}. Hence, to avoid unphysical behaviour, it is necessary to optimize the whole model at once. In the example shown in figure 6.10 this would probably imply a stronger smoothing of the traces, such that behaviour like the transient drop in the electron temperature rise indicated by the red ellipse during the equilibration at the higher heating power level is supressed. However, it is still a reasonable approach to optimize $\tilde{T}_{\mathrm{e}}(\rho, t)$ separately first to get a reasonable starting point. In the future the analysis will be done with an MCMC to avoid getting trapped in local maxima. This implies, due to the then high dimensional model (430 free parameters) to be analyzed, that a lot of iterations are probably necessary until convergence is reached. First tests indicate that after more than 2.2×10^5 MCMC iterations the chain is still moving up the posterior and no convergence has been reached.

---

[7]Shown are the channels 20 to 25 (corresponding to 143.9 GHz to 149.4 GHz, or cold resonance positions in effective radii between 0.35 to 0.67), hence all are on the high field side. Consequently, both O mode and »shinethrough« related effects are neglectable.

6.4. Conclusions

The evaluation of the ECE as an electron density diagnostic (even if in combination with an interferometer) was, to the best knowledge of the author, attempted for the first time. Here, the inferred density profiles agree well with the overall shape of the density profile determined by local measurements as available from the Thomson scattering diagnostic. It is clear that work on the speed of the forward calculations is still required, as analysis times on the order of months is unpractical for most usecases. An artificial neural network is planned to help to alleviate the long analysis times. The model for generating the training data is prepared, although getting the artificial neural network to behave sensibly in the vicinity of densities where »shinethrough« and cutoff effects play a role will be challenging. Rough estimations for the required time with an artificial neural network range on the order of minutes, at most hours, allowing to potentially make practical use of the ECE as a density diagnostic. If possible this would make a very valuable contribution to at least W7-X, as the Thomson scattering system as the main provider of electron density profiles operates at most at a frequency of 30 Hz, while an ECE based system would allow far finer temporal control, potentially reaching MHz. However, the ECE can only fulfill this role while the electron density is in a range close to the cutoff densities for several channels – or the electron density is so low that »shinethrough« starts to become dominant in some channels.

As an outlook on a Bayesian analysis of the electron heat diffusivity a model has been implemented, but not yet comprehensively analysed.

7

Résumé

In the context of furthering the understanding of magnetically confined plasmas it is important to model the available diagnostics comprehensively, and combine them via Bayesian forward modelling. The knowledge gained by the analysis of these models will help to design future experiments and improve theoretical understanding. Thus, in this work, the electron cyclotron emission (ECE) diagnostic, which is one of the main diagnostics (especially if high temporal resolutions are desired) for the electron temperature profile is modelled.

For the first time, the modelling of the ECE absolute calibration is done in a Bayesian way (cf. chapter 4) and includes the geometry of the calibration unit (cf. section 2.3.2). The classical approach only takes into account the measured signal difference during observation of two black body emitters at different reference temperatures. The periodic change of the emission source allows averaging over the performed cycles to improve the signal to noise ratio. In addition to that allows the developed model to use the whole time series of the average cycle, including the phases that collect radiation from multiple reference temperature radiation sources due to the finite width of the ECE line of sight. Furthermore, it allows to infer the sensitivities, i. e. the calibration factors, for all ECE channels simultaneously, allowing to calibrate weakly sensitive channels and to reduce the uncertainties. In addition, the width of the Gaussian antenna characteristic can be inferred, which deviates by less than 1 mm from the independently inferred beam width for the case shown in section 4.3.1. Note that the analysis of the model also provides the marginalised posteriors for the sensitivities, allowing to reconstruct asymmetric uncertainties.

The application of the determined calibration factors allows to obtain only radiation temperature spectra, however, in general, the desideratum is the elec-

tron temperature profile. Hence, a Bayesian forward model of the ECE radiation transport in the plasma was constructed (cf. chapter 5 and figure 5.1). Internally, the ray tracing code TRAVIS is used via a webservice. The forward model permits the inference of a multitude of parameters, including the posteriors of the electron temperature profile, the optical depth spectra and the normalized electron momentum spectra. Furthermore, the model is kept generic, such that it would be applicable for tokamaks as well. Noteworthily, the radiation transport depends on the electron density profile as well. Consequently, the ECE measurements contain information about the density as well, although the amount of information is small for optically thick plasmas in the »classical« operational regime (i. e. at W7-X: $10^{19}\,\mathrm{m}^{-3} < n_{\mathrm{e},0} < 8 \times 10^{19}\,\mathrm{m}^{-3}$). To allow a more robust inference without the knowledge of the exact electron density profiles in these cases, the interferometer was added to the ECE model, constraining primarily the line integrated electron density. Hence, the electron density profile is one of the parameters that can be inferred.

Finally, chapter 6 presents applications of the ECE system at W7-X. In the first part, the ECE model (cf. figure 5.1) is aiming to push to the boundaries of the operational space where ECE may be interpreted by applying it to a plasma with high electron density, such that the ECE is partially in X mode cutoff. The inferred electron density profile exhibits little variance as the channels close to their cutoff provide a lot of information on the density due to the high sensitivity of their lines of sight on the exact density. The main drawback of this analysis is the required time, which is well over three months. However, the groundwork for generating an artificial neural network of the time critical part of the model, namely the TRAVIS ray tracing and radiation transport calculations, has been layed. If further pursued this should bring enough speed improvements to eventually open the possibility to use the ECE (in combination with the interferometer) in partial cutoff as a *density* diagnostic for the first time. A proof of concept for an analysis with low electron density plasmas was shown and dicussed as well. Subsequently, a simple heat transport model was implemented, although the long analysis time caused by the high dimensionality of the model (430 free parameters) prohibited an analysis within the scope of this thesis.

Analyses of Bayesian forward models have been performed exemplarily for the ECE in this thesis. Such models allow easy combination of measurements from a variety of diagnostics and the inference of the posteriors of the corresponding plasma parameters. In preparation for the Bayesian analysis of models that combine several diagnostics multiple datasources have been written dur-

ing this thesis, notably the ECRH, ECE, X-ray imaging crystal spectroscopy, snifferprobe, single channel dispersion interferometry, NBI, Langmuir probe system and VMEC reference identifier datasources for W7-X. Generic nodes for obtaining the electron temperature and density profiles were added to the Minerva framework at W7-X as well. The general model for W7-X, combining nine diagnostics, has been created, although the analysis of it remains an open issue.

A

Temperature dependence of the calibration model prediction

For a specific channel i the predicted signal in bit is described via equation (4.7). Given a linear sensitivity, this equation can be rewritten to

$$f_i^{\text{pred}}(\varphi) = T_{\text{eff},i}(\varphi)\tilde{\eta}_i = \frac{\Delta s_i T_{\text{eff},i}(\varphi)}{T_{\text{eff}}^{\text{hot}} - T_{\text{eff}}^{\text{cold}}}. \tag{A.1}$$

As we do subtract the offset of the measured signal, we end up with

$$f_i^{\text{pred}}(\varphi) - \bar{f}_i^{\text{pred}} = \Delta s_i \left(\frac{T_{\text{eff},i}(\varphi)}{T_{\text{eff}}^{\text{hot}} - T_{\text{eff}}^{\text{cold}}} - \underbrace{\left(\beta - \frac{\gamma}{n} \right)}_{\text{offset}} \right), \tag{A.2}$$

where

$$\beta = \frac{T_{\text{eff}}^{\text{hot}}}{T_{\text{eff}}^{\text{hot}} - T_{\text{eff}}^{\text{cold}}},$$

γ is the number of entries for the cold vector and n the number of entries for the cold and the hot vector. As all effective temperatures are between $T_{\text{eff}}^{\text{cold}}$ and $T_{\text{eff}}^{\text{hot}}$, they can be expressed by $\beta - \delta$ with $\delta \in [0, 1]$. Thus, the predicted signal can be written as

$$f_i^{\text{pred}}(\varphi) - \bar{f}_i^{\text{pred}} = \Delta s_i \left(\beta - \delta - \left(\beta - \frac{\gamma}{n} \right) \right) \tag{A.3}$$

$$= \Delta s_i \left(\frac{\gamma}{n} - \delta \right). \tag{A.4}$$

Thus, the predicted signal does not depend on the absolute values of the effective temperature for a linear sensitivity.

B
ECE calibration code

Here, the main class for performing an ECE calibration at W7-X from within Java is explained. The code is taken from `Calibrate.java` in the corresponding Minerva package and corresponds to [110]. Obviously, not all code of the calibration analysis can be explained here. Nevertheless, the code is fully Javadoc documented, so should be accessible to the interested reader.

The implementation of the `main` method as the entry point for the Java `Calibrate` class is presented in the subsequent code snippets. Code B.1 starts with the beginning of the `main` method.

Code B.1: Instantiation of calibration class, and first setter of the ECE Minerva calibration code.

```
27  /**
28   * Intended to run the selected radiometer calibration model
       with the specified
29   * settings.
30   *
31   * @param args empty
32   */
33  public static void main(String[] args) {

35      W7xEceQmeCalibrate calib = new W7xEceQmeCalibrate();

37      calib.setCalibrationSegments(QME_20181115_2());
```

Line 35 creates an instance of the calibration class. The setter in line 37 takes a `long[][]` array, such that multiple calibration segments (as W7-X time stamps)

can be passed to the inference routines. This allows »stitching together« multiple calibration segments which is useful for example if one had to refill liquid nitrogen during the calibration and wants to discard the associated data. *QME_20181115_2()* is just an example method that provides the W7-X time stamps.

Code B.2 shows the setters related to the hardware setup during the calibration and the logging settings.

Code B.2: Hardware and logging setters of the ECE Minerva calibration code.

```
39      calib.setChannel(8);
40      calib.setChopperChannel(16);
41      calib.setChopperZoom(false);
42      calib.setLogToFile(true);
43      calib.setMode(Modes.HotColdLoad);
```

Line 39 sets the channel to evaluate. This setter is only of importance if a model is selected that evaluates a single channel. Counting starts at 1. The channel containing the chopper signal is specified via the setter in line 40. Additionally, line 41 allows to specify whether the chopper channel is located in the zoom or the standard ECE system. Line 42 allows to specify whether logging information should be written to a file. This is particularly recommended for evaluations that take a long time to get a quick overview of the current status. The mode can be selected via the setter in line 43. The *Modes* enum has two possible values, *HotColdLoad* and *NoiseSource*. Note that the *NoiseSource* mode is experimental, and not as deeply tested as the *HotColdLoad* mode. In principle, the *NoiseSource* mode allows an absolute (sic!) calibration of the ECE, although at the moment the effect of the transmission line is not taken into account therein.

Code B.3 shows how to select the desired model. As explained in section 4.1, there is a plethora of calibration models to choose from.

Code B.3: The calibration model selection of the ECE Minerva calibration code.

```
45      calib.setModel(RadiometerHotColdLoadModel.STANDARD_HYPER);
```

In fact, due to technical reasons, there are more models available than described above. As the zoom system can in principle be calibrated the same way as the standard ECE system, options are also available to calibrate the

former, or both systems combined. Do note that the setter takes an interface object of type `EceCalibration` as an argument, which is implemented by both *RadiometerHotColdLoadModel* and *RadiometerNoiseSourceModel*. Again, the code of the noise source models is kept similar to the code of the hot cold load models (not the structure of the model, though), but it is less thoroughly tested which is why chapter 4 does not mention the noise source models. However, it could be worthwhile to finish the implementation of the noise source models in the future, as this would allow a much quicker calibration procedure (on the order of few minutes). The disadvantage of the noise source models are currently twofold: i) the noise source is located *after* the transmission line (cf. figure 2.4), thus not taking the transmission line into account, and ii) the accuracy of the excess noise ratio of the current noise source is only certain within 1 dB, thus leading inevitably to large uncertainties in the resulting calibration factors.

Code B.4 highlights the available analysis settings. The settings shown in code B.4 correspond to a calibration performed after OP 1.2b of W7-X.

Code B.4: Setters for the analysis of the ECE Minerva calibration code, part I.

```
47    calib.setNumBins(0);
48    calib.setMultithreading(false);
49    calib.setNumThreads(4);
50    calib.setNumIntegralSteps(0);
51    calib.setUseMap(true);
52    calib.setNumMapIterations(500);
53    calib.setUseLaplaceApproximation(false);
54    calib.setUseMcmc(true);
55    calib.setNumMcmcBurnIn(500_000);
56    calib.setNumMcmcKeepEvery(3);
57    calib.setNumMcmcIterations(50_000);
58    calib.setScan(true);
59    calib.setNumScanPoints(100);
60    calib.setUseSinglePeriodAnalysis(false);
61    calib.setCompareToReferenceDischarges(true);
62    calib.setValidSince("2018-11-15");
63    calib.setReason("calibration after op1.2b");
64    calib.setNumSamplesToPlot(100);
```

Line 47 allows you to select the number of bins to be used for the conditionally averaged signal. A reasonable default is chosen if 0 is the argument. Be aware

that for high sampling rates the default will determine a large number of bins, which will slow down the analysis correspondingly. If a model has been chosen that does multiple single channel evaluations, line 48 allows to activate multi-threading. The corresponding number of threads can be chosen in line 49. The integrals in equation (4.5) are performed numerically, thus the number of steps to be done can be chosen via the setter in line 50. Again, 0 chooses a default that performed well for the calibrations done for this thesis. Choosing to few steps can lead to a posterior with »steps«. This is not relevant for noise source models. If an MAP inversion should be performed, the setter in line 51 should be set to `true`. This is generally recommended. The corresponding number of iterations is specified in line 52. In general, a value around 500 seems to behave well. Do note that internally the MAP methods are stopped once either these number of iterations have been performed, or the change in the logarithm of the prior times likelihood probability after 10 iterations is less than 1. If the setter in line 53 is set to `true` a Laplace approximation (LA) is attempted to estimate the initial proposal distriubution for the MCMC. While fully implemented, this has not been tested in practice and thus should be considered an experimental option. Similar to line 51, the MCMC inversion can be activated by the setter in line 54. The number of burn-in iterations is specified in line 55. Typically, `500_000` should be more than enough. As there can be a slight correlation between subsequent points in an MCMC (cf. section 3.3), one can set a number of points to skip before the next value gets stored. This can be done as shown in line 56. The number of MCMC iterations to be performed after the burn-in can be specified via the setter in line 57. If the setter in line 58 is set to `true` a scan is done in which the free parameters are set to the mean values obtained from an MCMC performed beforehand. Therein, each free parameter is scanned independently from -3σ to $+3\sigma$ around the MCMC mean, with the number of points for the scan determined via the setter in line 59. Thus, the distributions obtained that way will be generally too narrow, as uncertainty contributions from the other free parameters are neglected. Usually, the scan is not necessary. Also only relevant for very specific usecases is the single period analysis that can be activated via the setter shown in line 60. In this analysis the inversions are executed as usual. Subsequently, the free parameters in the model are set to their MCMC means and are deactivated. The only parameters that remain free are the bit dip and the variance scaling (if active). Then, another short MAP and MCMC inversion is done *for each period in the calibration segment*. The inferred sensitivites are stored. This analysis method can be used to detect some sort of

sensitivity drifts, although it is generally better to still use multiple periods as without averaging there will be a huge spread of values, and the analysis will take much longer. To see quickly after a calibration which channels changed, an option to compare previously calculated ECE spectra with spectra calculated with the results from this analysis, see line 61. The validity of the calibration values can be set as shown in line 62. This is not automized, as the validity might predate the actual calibration, if the hardware was left unchanged beforehand. Line 63 contains the code to give a reason for the calibration (analysis) that is saved to the W7-X database if writing to it is enabled. If plots are to be saved the number of samples used can be specificated as shown in line 64.

Code B.5: Setters for the analysis of the ECE Minerva calibration code, part II.

```
65    calib.setOutputFolder(null);
66    calib.setPath(null);
67    calib.setPlot(true);
68    calib.setReducedNumSamples(50);
69    calib.setRemarks(null);
70    calib.setRemarksZoom(null);
71    calib.setWriteToDatabase(true);
72    calib.setDeactivateNonnecessaryFreeParameters(true);
```

The output folder can be specified with the setter shown in line 65. If you set the value to `null` a sensible default folder is created. Line 66 shows the setter to specify the name of the log file. The path of the output folder is prepended. Again, `null` provides a default. Saving plots automatically can be done via the setter in line 67. As the calibration values stored on the W7-X database do not require all MCMC samples for most subsequent analyses, a reduced number of samples as specified in line 68 is stored directly with the calibration values. The full MCMC samples are not discarded, but kept separately to keep access time for the mean calibration values small. For each calibration uploaded to the W7-X database remarks can be added for the standard ECE system (cf. line 69) as well as for the zoom system (cf. line 70). If, like for the example shown here, the setters set the remarks to `null`, no remarks are added. Via the setter in line 71 one can determine whether to write the results automatically (including the full model structure) to the immutable W7-X database. A huge speedup can be achieved be deactivating free parameters on which the prediction does not explicitly depend (cf. appendix A), as this decreases the effect of the »curse of

dimensionality« notably, especially for the combined models. The uncertainty of these deactivated free parameters is still taken into account for the MCMC, as their values are sampled anew for each MCMC iteration.

Finally, code B.6 shows in line 73 the method to run the calibration with the options selected above. Line 74 closes the `main` method.

Code B.6: The method to run the W7-X ECE calibration and the end of the `main` method begun in code B.1.

```
73      calib.run();
74  }
```

All the setters shown here have analogous keywords that can be set for the standalone executable version directly via the command line or can be provided via a `.properties` file (example: 🖫), allowing a high degree of flexibility.

C
ECE profile inversion code

Here, the main class for performing an ECE profile inversion at W7-X from within Java is explained. The code corresponds to [111] from `History.java` and `Inference.java` in the corresponding Minerva package. `History.java` contains the settings for previously run inferences with manually set settings to allow for easy reproducibility. Obviously, not all code of the profile inversion analysis package can be explained here. Nevertheless, the code is fully Javadoc documented to be accessible to the interested reader. Example settings for a profile inference are shown in code C.1 which starts with the enum that lists analyses done for a specific person.

Code C.1: Enum with evaluated shots.

```
920   /**
921    * Evaluated profiles for Udo Hoefel.
922    *
923    * @author  Udo Hoefel
924    */
925   enum UdoHoefel implements EvaluatedProfiles {
              ⋮
937       shot_20180823_016_time_4_45,
938       shot_20180823_016_time_5_15,
939       shot_20181009_034_time_3_00,
940       ;
```

There is an enum for each of the people that requested analyses. The beginning of the implementation of the `EvaluatedProfiles` interface is shown in code C.2.

Code C.2: Implementation of the `EvaluatedProfiles` interface, part I.

```
942     @Override
943     public W7xEceQmeProfileInversion getSettings() {
944         W7xEceQmeProfileInversion inference = new
        W7xEceQmeProfileInversion();

946         switch (this) {
                           ⋮
1521        case shot_20180823_016_time_4_45: {
1522            inference.setYear(2018);
1523            inference.setMonth(8);
1524            inference.setDay(23);
1525            inference.setExperiment(16);
1526            inference.setScenario(2);
1527            inference.setTime(4.45);

1529            inference.setTimeWindowForAveraging(50);
1530            inference.setDataVersion(0);

1532            inference.setEceChannelsToDeactivate(16);
1533            inference.
        setEceChannelsToDeactivateForInitialGuess(16);

1535            inference.setTeInitialGuess(TeInitialGuess.ECE);
1536            inference.setNeInitialGuess(NeInitialGuess.
        THOMSON_SCATTERING_SCALED);
```

The year of the discharge to be analyzed is given in line 1522, the month in line 1523 and the day in line 1524. On the specified date, the lines 1525 to 1527 allow to determine the nanosecond at which to analyze the combined ECE and interferometry data. Line 1529 contains the command to set the time window over which the ECE and interferometry signal is averaged. An example of how to use the most recent data version is shown in line 1530. As this is not guaranteed to be reproducible the version that the most recent version corresponds to is written explicitly to the output files. Sometimes some ECE channels are problematic due to various problems. Specific ECE channels can be deactivated for the inference as shown in line 1532 or only for the initial guess as shown in line 1533. The initial guess for the electron temperature pro-

file is set exemplarily in line 1535. The enum *TeInitialGuess* allows to select the **ECE** to use all ECE channels for the initial guess, **ECE_LFS** restricts the channels for the initial guess to the low field side, **ECE_HFS** correspondingly to the channels on the HFS. If **THOMSON_SCATTERING** is selected the initial guess uses the measurements from the Thomson scattering system. Furthermore, **TRAVIS_PARAMETERIZATION** uses equation (5.1) to represent the profile. The parameters can be set via `setTeTravisParamsInitialGuess`. If a manual profile should be applied use **MANUAL**. The `setManualTeInitialGuess` method allows to set the corresponding profile and requires a `ScalarFunction1D`.

Correspondingly, line 1536 shows how to set the initial guess for the electron density profile. Similar as for the electron temperature profile initial guess, an enum, *NeInitialGuess*, allows the selection of the type of initial electron density profile guess. Possible options are **THOMSON_SCATTERING** using the density profile as measured by Thomson scattering, **THOMSON_SCATTERING_SCALED**, which uses the same profile as for **THOMSON_SCATTERING** but scaled to match the line integrated density as measured by the single channel dispersion interferometry. A default profile scaled to the interferometry measurements can be used via **INTERFEROMETRY_SCALED_DEFAULT_PROFILE** if Thomson scattering data are not available. Similar to the initial guess for the electron temperature profile, **MANUAL** and **TRAVIS_PARAMETERIZATION** are available for the electron density profile initial guess and can be used with their corresponding setters. To scale a manual profile such that consistency with the line integrated electron density measured by the single channel dispersion interferometer is guaranteed one can use **INTERFEROMETRY_SCALED_MANUAL_PROFILE**. The next part of the implementation is shown in code C.3.

Code C.3: Implementation of the `EvaluatedProfiles` interface, part II.

```
1538            inference.setNumRho(20);
1539            inference.setAdditionalPointNearCore(true);
1540            inference.setDensityProfileActive(true);
1541            inference.setUseHyperparams(false);
1542            inference.setDeltaRho(0.0);
1543            inference.setStepsizeRho(0.005);
1544            inference.setInitialGuessIterations(0);

1546            inference.setModel(EceProfileInversion.
        ECE_PROFILE_INVERSION_OX);
```

The number of points at which the Gaussian processes are evaluated is set in line 1538. Increasing this number increases the dimensionality of the model. By setting the setter in line 1539 to true, one will get an additional Gaussian process evaluation close to the core (at $\rho = 10^{-3}$), such that the virtual observations that limit the gradients in the core do not have to depend as much on the interpolation mode used between the Gaussian process points. However, this increases the dimensionality of the model by 2. In line 1540 it is shown how one can deactivate the density profile and just keep it at the initial guess. Due to the computational cost that free hyperparameters entail, they are normally deactivated. In case they are still wanted as free parameters, one can set the setter in line 1541 to **true**. If a shift of the initial guess of the electron temperature profile is desired, line 1542 allows to specify by how much the profile is shifted. The stepsize for this shift is set in line 1543, such that effectively the prior times likelihood of the whole model is checked for an initial guess electron temperature profile shifted by $-\Delta\rho, -\Delta\rho + \delta\rho, \ldots, \Delta\rho$ and the one with the highest probability is used for the inference. Iterating on the initial electron temperature profile guess will be done the number of times specified in line 1544. Usually, this can be left at zero, more than one iteration typically make the initial guess worse. Selecting the type of model can be done like is shown in line 1546. The enum *EceProfileInversion* contains the available models. The standard model ***ECE_PROFILE_INVERSION_OX*** uses Gaussian processes, ***ECE_PROFILE_INVERSION_OX_HYPER*** uses in addition a variance scaling factor. For the parameterization described by equation (5.1) two corresponding models are available: The first available model is ***ECE_PROFILE_INVERSION_OX_TRAVIS_PARAM***. The one including the variance scaling is ***ECE_PROFILE_INVERSION_OX_HYPER_TRAVIS_PARAM***. The settings for the MAP and MCMC inversion are specified in code C.4.

Code C.4: Implementation of the `EvaluatedProfiles` interface, part III.

```
1548            inference.setUseMap(true);
1549            inference.setNumMapIterations(150);
1550            inference.setMapAutoConverge(true);
1551            inference.setUseMcmc(true);
1552            inference.setNumMcmcBurnIn(100_000);
1553            inference.setNumMcmcKeepEvery(3);
1554            inference.setNumMcmcIterations(10_000);
```

If an MAP inversion should be done the setter in line 1548 has to be set to **true**. The maximum number of MAP iterations done can be set as shown in line 1549. If one wants to stop the MAP automatically once the logarithm of the prior times likelihood does not change by more than 1 over 7 iterations the setter in line 1550 should be set to **true**. This might lead to fewer MAP iterations being done than specified in line 1549. To perform an MCMC inversion set the setter in line 1551 to **true**. The number of MCMC iterations to be used for the burn-in is specified as shown in line 1552. To make sure that the MCMC samples are uncorrelated, you can choose to keep only every n^{th} sample, cf. line 1553. The number of MCMC iterations to be done aside the burn-in iterations is set as shown in line 1554. Further settings are shown in code C.5.

Code C.5: Implementation of the `EvaluatedProfiles` interface, part IV.

```
1557            W7xVmecIdDataSource vmecId = new
         W7xVmecIdDataSource();
1558            ArchiveDBFetcher adb = ArchiveDBFetcher.
         defaultInstance();
1559            long t1 = adb.quiet().getT1("20180823.016.002");
1560            vmecId.setNanosecond(t1 + 4_450_000_000L);
1561            inference.setFixedVmecId(vmecId.getBestVmecId());
1562            VmecInfoType type = vmecId.getBestVmecIdType();
1563            double scaling = vmecId.getMagneticFieldScaling(
         type);
1564            inference.setMagneticFieldRescaling(scaling);

1566            inference.setMirrorId("W7X-Vessel/
         W7X_ECE_targetTileOnly_noHFShorn.xml");

1568            inference.setFixedXModeContribution(null);
1569            inference.setFixedOModeContribution(null);
```

Line 1557 instantiates a datasource that can provide automatically select best guess VMEC equilibria for a given timepoint. Then, line 1558 instantiates the class used for fetching data. Subsequently, line 1559 fetches the T1 trigger of, exemplarily, #20180823.016.002 in W7-X nanoseconds. Thereafter, the VMEC equilibrium datasource is set to T1 + 4.5 s in line 1560. This allows to set the profile reconstruction to the best automatically determinable VMEC equilibrium, see line 1561. As there may be different VMEC equilibria supplied by different

analysis types, the analysis type that is estimated by the VMEC datasource to match best to the experimental data is fetched in line 1562. But the precalculated VMEC equilibrium may not have the appropriate absolute *B* field to which the ECE at W7-X is highly sensitive due to the shallow magnetic field gradient along the line of sight of the ECE. Therefore, a factor that rescales the magnetic field to match to the field as expected from the experiment is fetched in line 1563. The shape of the flux surfaces stays the same. If reflections, for example of the O mode, are to be taken into account TRAVIS needs to have a way to determine where the ray gets reflected, what the reflection coefficient is and so on. The file that contains this information is supplied as shown in line 1566. X mode and O mode contributions are typically free parameters, in which case the setters in lines 1568 and 1569 have to be set to **null**. If they are set to any number between 0 and 1 they are no longer treated as free parameters and kept fixed. The remaining settings for the inference are shown in code C.6.

Code C.6: Implementation of the `EvaluatedProfiles` interface, part V.

```
1572            inference.setInterpolationMode(InterpolationMode.
        CUBIC);
1573            inference.setLogToFile(true);
1574            inference.setShowPlots(false);
1575            inference.setPlotLevel(Level.INFO);
1576            inference.setNumSamplesToPlot(100);
1577            inference.setNumSamplesToSave(100);
1578            inference.setOutputFolder(null);
1579            inference.setPath(null);
1580            inference.setNoOutput(false);
1581            inference.setPlot(true);
1582            inference.setBeamInfo(true);
1583            inference.setWriteToDatabase(false);
1584            inference.setDatabase(Database.W7X_ARCHIVE);
1585            inference.setWriteAliases(false);
1586            break;
1587        }
                                ⋮
1696        return inference;
1697        }
1698    }
```

The interpolation mode can be set to either *LINEAR* or *CUBIC*, see line 1572 and determines how the inferred Gaussian process points get turned to a continuous profile. If one wants to have the logging messages written to a log file, the setter in line 1573 should be set to **true** (recommended). Showing each plot interactively (which blocks the evaluation until closed) can be achieved by setting the setter in line 1574 to **true**. The plot level can be set as shown in line 1575. This allows to avoid plots with lower importance. The number of samples to be plotted can be set as shown in line 1576. Correspondingly, the number of samples to be saved is determined by the setter in line 1577. Usually, the output folder is determined automatically. However, the setter in line 1578 can be used to use a different output folder. If a file is used to store the logged messages its path can be specified as shown in line 1579, **null** uses a default. For automatic routines that upload for example the initial guess to the archive, automatically created local plots and files are usually just a waste of time. Thus, all output can be deactivated if the setter in line 1580 is set to **true**. In the case that only plots should be avoided set the setter in line 1581 to **false**. For a very detailed analysis of a time point, activate additional information about the ECE beam by setting the setter in line 1582 to **true**. The results can be automatically written to the database if the setter in line 1583 is set to **true**. Line 1584 shows the setter allowing to select to which database to write results to. Possible options are *W7X_ARCHIVE* for the productive archive and *W7X_TEST* for the test archive. Finally, one can force rewrite the aliases pointing to the data if the setter in line 1585 is set to **true**. By default, the aliases will be written automatically if required. In the `Inference` class one can then call the code shown in code C.7 in the `main` method.

Code C.7: Code to run the profile inversion from within Java.

```
37        UdoHoefel.shot_20180823_016_time_4_45.run();
```

An example input file for the compiled Java executable as used for the analysis in section 5.2 can be found here: 🖫.

Acknowledgements

First I have to thank Prof. Wolf for giving me the opportunity to do this PhD thesis at IPP.

I am deeply indebted to Matthias Hirsch as well, as I could always ask for advice, always got a lot of useful comments and was very gently directed back to not stray away too much from the topics at hand. I also have to mention Torsten Stange. Not only did he introduce me to IPP, but he also always helped me whenever I needed help and the light in his office was often the only »guiding star« when I left IPP late. Nikolai Marushchenko and Yuriy Turkin were also very helpful as they provided theoretical insights, the webservice for TRAVIS and never seemed to get tired of my questions. A large thank you is also deserved by Michael Grahl who always quickly helped if I didn't understand something with the archive or the webservices. Keeping my sanity was made a lot easier by the shared suffering of the quirks of the archive by Kai-Jakob Brunner and especially Oliver Ford, who were also fostering my physics understanding tremendously.

A thank you also goes out to Fabian Wilde (I am curious how long it will take you to get the Tesla) and Peter Traverso. The office was never again so much fun after you left the office!

I already miss the Squash sessions with Adnan Ali, Daniel Böckenhoff and occasionally Tullio Barbui and Lukas Rudischhauser. The gatherings with Valeria Perseo and Priyanjana Sinha were amongst the most enjoyable time I had in Greifswald. Merci beaucoup and mea culpa for not coming to Lindy Hop anymore. The discussions with Humberto Trimiño Mora were amazing, although sometimes slightly scary. You are too good at them!

A special thanks goes out to Jakob Svensson. Starting the PhD I did not expect major revisions of my attitude towards how an analysis should be done to be on the table, but lo and behold, your view on Bayes was able to do so. The other Minervarians, especially Sehyun Kwak (or was it Kwak Sehyun?) and Andrea Pavone are both amazingly kind and clever. I learned a lot from you. I am happy

that I had the opportunity to work with Jonathan Schilling, whose speed and efficiency is simply stunning.

Many of the people above I consider good friends, I am glad that I met you. Of course I also have to mention my family which, to a large part, made me who I am today. Also, everyone very much appreciated the »care packages«, even though they were not really necessary ☺. And last but not least: Thank you for everything, Sandy.

Statutory declaration

I hereby declare in accordance with the examination regulations that I myself have written this document, that no other sources as those indicated were used and all direct and indirect citations are properly designated, that the document handed in was neither fully nor partly subject to another examination procedure or published and that the content of the electronic exemplar is identical to the printing copy.

Greifswald, 24th February 2020

UDO HÖFEL

List of Figures

List of Codes

Publications as first author

Peer-reviewed articles

[1] **U. Höfel** et al. »Bayesian modelling of the electron cyclotron emission diagnostic at Wendelstein 7-X«. In: *Nuclear Fusion*, (2020). In preparation.

[2] **U. Höfel** et al. »Bayesian Modelling of Microwave Radiometer Calibration on the example of the Wendelstein 7-X Electron Cyclotron Emission diagnostic«. In: *Review of Scientific Instruments*, Vol. 90.4, Art. 04 3502 (11th Apr. 2019). doi: 10.1063/1.5082542.

Publications as coauthor

Peer-reviewed articles

[1] T. STANGE et al., amongst them **U. HÖFEL**. »First demonstration of
     fully equilibrated ions and electrons in a Magnetically Confined High
     Temperature Plasma sustained by 2nd harmonic O-mode polarized ECRH
     only«. In: *Physical Review Letters*, (2020). Submitted.

[2] N. A. PABLANT et al., amongst them **U. HÖFEL**. »Investigation of the
     neoclassical ambipolar electric field in ion-root plasmas on W7-X«. In:
     Nuclear Fusion, (22nd Jan. 2020). DOI: 10.1088/1741-4326/ab6ea8.

[3] J. W. OOSTERBEEK et al., amongst them **U. HÖFEL**. »Michelson In-
     terferometer design in ECW heated plasmas and initial results«. In:
     Fusion Engineering and Design, Vol. 146 (Sept. 2019), pages 959–962. DOI:
     10.1016/j.fusengdes.2019.01.124.

[4] J. BALDZUHN et al., amongst them **U. HÖFEL**. »Pellet fueling experi-
     ments in Wendelstein 7-X«. In: *Plasma Physics and Controlled Fusion*,
     Vol. 61.9, ART. 095012 (13th Aug. 2019). DOI: 10.1088/1361-6587/
     ab3567.

[5] R. C. WOLF et al., amongst them **U. HÖFEL**. »Performance of Wendel-
     stein 7-X stellarator plasmas during the first divertor operation phase«.
     In: *Physics of Plasmas*, Vol. 26.8, ART. 082504 (13th Aug. 2019). DOI:
     10.1063/1.5098761.

[6] D. ZHANG et al., amongst them **U. HÖFEL**. »First Observation of a
     Stable Highly Dissipative Divertor Plasma Regime on the Wendelstein 7-X
     Stellarator«. In: *Physical Review Letters*, Vol. 123.2, ART. 025002 (9th July
     2019). DOI: 10.1103/PhysRevLett.123.025002.

[7]   A. Pavone et al., amongst them **U. Höfel**. »Neural network approximation of Bayesian models for the inference of ion and electron temperature profiles at W7-X«. In: *Plasma Physics and Controlled Fusion*, Vol. 61.7, Art. 075012 (30th May 2019). DOI: 10.1088/1361-6587/ab1d26.

[8]   A. Dinklage et al., amongst them **U. Höfel**. »Plasma termination by excess pellet fueling and impurity injection in TJ-II, the Large Helical Device and Wendelstein 7-X«. In: *Nuclear Fusion*, Vol. 59.7, Art. 07601 0 (24th May 2019). DOI: 10.1088/1741-4326/ab17fd.

[9]   S. C. Liu et al., amongst them **U. Höfel**. »The effects of magnetic topology on the scrape-off layer turbulence transport in the first divertor plasma operation of Wendelstein 7-X using a new combined probe«. In: *Nuclear Fusion*, Vol. 59.6, Art. 066001 (24th Apr. 2019). DOI: 10.1088/1741-4326/ab0d29.

[10]   M. Hirsch et al., amongst them **U. Höfel**. »ECE Diagnostic for the initial Operation of Wendelstein 7-X«. In: *European Physical Journal Web of Conferences*, Vol. 203, Art. 03007 (25th Mar. 2019). DOI: 10.1051/epjconf/201920303007.

[11]   M. Zanini et al., amongst them **U. Höfel**. »ECCD operations in the second experimental campaign at W7-X«. In: *European Physical Journal Web of Conferences*, Vol. 203, Art. 02013 (25th Mar. 2019). DOI: 10.1051/epjconf/201920302013.

[12]   N. Chaudhary et al., amongst them **U. Höfel**. »Investigation of Optically Grey Electron Cyclotron Harmonics in Wendelstein 7-X«. In: *European Physical Journal Web of Conferences*, Vol. 203, Art. 03005 (25th Mar. 2019). DOI: 10.1051/epjconf/201920303005.

[13]   J. W. Oosterbeek et al., amongst them **U. Höfel**. »Assessment of ECH stray radiation levels at the W7-X Michelson Interferometer and Profile Reflectometer«. In: *European Physical Journal Web of Conferences*, Vol. 203, Art. 03010 (25th Mar. 2019). DOI: 10.1051/epjconf/201920303010.

[14]   G. G. Plunk et al., amongst them **U. Höfel**. »Stellarators Resist Turbulent Transport on the Electron Larmor Scale«. In: *Physical Review Letters*, Vol. 122.3, Art. 035002 (25th Jan. 2019). DOI: 10.1103/PhysRevLett.122.035002.

[15] D. Moseev et al., amongst them **U. Höfel**. »Collective Thomson scattering diagnostic at Wendelstein 7-X«. In: *Review of Scientific Instruments*, Vol. 90.1, Art. 013503 (8th Jan. 2019). DOI: 10.1063/1.5050193.

[16] R. C. Wolf et al., amongst them **U. Höfel**. »Electron-cyclotron-resonance heating in Wendelstein 7-X: A versatile heating and current-drive method and a tool for in-depth physics studies«. In: *Plasma Physics and Controlled Fusion*, Vol. 61.1, Art. 014037 (27th Nov. 2018). DOI: 10.1088/1361-6587/aaeab2.

[17] A. Langenberg et al., amongst them **U. Höfel**. »Impurity transport studies at Wendelstein 7-X by means of x-ray imaging spectrometer measurements«. In: *Plasma Physics and Controlled Fusion*, Vol. 61.1, Art. 014030 (23rd Nov. 2018). DOI: 10.1088/1361-6587/aaeb74.

[18] G. Fuchert et al., amongst them **U. Höfel**. »Global energy confinement in the initial limiter configuration of Wendelstein 7-X«. In: *Nuclear Fusion*, Vol. 58.10, Art. 106029 (29th Aug. 2018). DOI: 10.1088/1741-4326/aad78b.

[19] A. Pavone et al., amongst them **U. Höfel**. »Bayesian uncertainty calculation in neural network inference of ion and electron temperature profiles at W7-X«. In: *Review of Scientific Instruments*, Vol. 89.10, Art. 10K102 (6th Aug. 2018). DOI: 10.1063/1.5039286.

[20] A. Dinklage et al., amongst them **U. Höfel**. »Magnetic configuration effects on the Wendelstein 7-X stellarator«. In: *Nature Physics*, Vol. 14.6 (21st May 2018). DOI: 10.1038/s41567-018-0141-9.

[21] B. P. van Milligen et al., amongst them **U. Höfel**. »Study of radial heat transport in W7-X using the Transfer Entropy«. In: *Nuclear Fusion*, (19th Apr. 2018). DOI: 10.1088/1741-4326/aabf5d.

[22] N. A. Pablant et al., amongst them **U. Höfel**. »Core radial electric field and transport in Wendelstein 7-X plasmas«. In: *Physics of Plasmas*, Vol. 25.2, Art. 022508 (12th Feb. 2018). DOI: 10.1063/1.4999842.

[23] S. C. Liu et al., amongst them **U. Höfel**. »Observations of the effects of magnetic topology on the SOL characteristics of an electromagnetic coherent mode in the first experimental campaign of W7-X«. In: *Nuclear*

Fusion, Vol. 58.4, Art. 046002 (7[th] Feb. 2018). DOI: 10.1088/1741-4326/aaa98e.

[24] T. Stange et al., amongst them **U. Höfel**. »Advanced electron cyclotron heating and current drive experiments on the stellarator Wendelstein 7-X«. In: *European Physical Journal Web of Conferences*, Vol. 157, Art. 02008 (23[rd] Oct. 2017). DOI: 10.1051/epjconf/201715702008.

[25] S. A. Bozhenkov et al., amongst them **U. Höfel**. »The Thomson scattering diagnostic at Wendelstein 7-X and its performance in the first operation phase«. In: *Journal of Instrumentation*, Vol. 12.10, Art. P 10004 (10[th] Oct. 2017). DOI: 10.1088/1748-0221/12/10/P10004.

[26] H.-S. Bosch et al., amongst them **U. Höfel**. »Final integration, commissioning and start of the Wendelstein 7-X stellarator operation«. In: *Nuclear Fusion*, Vol. 57.11, Art. 116015 (3[rd] Aug. 2017). DOI: 10.1088/1741-4326/aa7cbb.

[27] R. C. Wolf et al., amongst them **U. Höfel**. »Major results from the first plasma campaign of the Wendelstein 7-X stellarator«. In: *Nuclear Fusion*, Vol. 57.10, Art. 102020 (27[th] July 2017). DOI: 10.1088/1741-4326/aa770d.

[28] M. Hirsch et al., amongst them **U. Höfel**. »Confinement in Wendelstein 7-X limiter plasmas«. In: *Nuclear Fusion*, Vol. 57.8, Art. 086010 (14[th] June 2017). DOI: 10.1088/1741-4326/aa7372.

[29] D. Moseev et al., amongst them **U. Höfel**. »Inference of the microwave absorption coefficient from stray radiation measurements in Wendelstein 7-X«. In: *Nuclear Fusion*, Vol. 57.3, Art. 036013 (22[nd] Dec. 2016). DOI: 10.1088/1741-4326/aa4f13.

[30] M. Krychowiak et al., amongst them **U. Höfel**. »Overview of diagnostic performance and results for the first operation phase in Wendelstein 7-X (invited)«. In: *Review of Scientific Instruments*, Vol. 87.11, Art. 11D304 (27[th] Oct. 2016). DOI: 10.1063/1.4964376.

[31] A. Köhn et al., amongst them **U. Höfel**. »Schemes of microwave heating of overdense plasmas in the stellarator TJ-K«. In: *Plasma Physics and Controlled Fusion*, Vol. 55.1, Art. 014010 (17[th] Dec. 2012). DOI: 10.1088/0741-3335/55/1/014010.

Acronyms

ADC analog digital converter
AUG ASDEX Upgrade

CPU central processing unit
CTS collective Thomson scattering

DAQ data acquisition
DCI deuterium cyanide laser interferometry

ECCD electron cyclotron current drive
ECE electron cyclotron emission
ECRH electron cyclotron resonance heating
ESB enterprise service bus

HFS high field side
HPP heat pulse propagation

IDA integrated data analysis
IF intermediate frequency
ISS04 international stellarator scaling 2004
ISS95 international stellarator scaling 1995
ITER International Thermonuclear Experimental Reactor

JET Joint European Torus

LA Laplace approximation
LiBES lithium beam emission spectroscopy
LO local oscillator

MAP maximum a posteriori
MAST Mega-Ampere Spherical Tokamak
MCMC Markov chain Monte Carlo
MHD magnetohydrodynamic

NBI neutral beam injection

OP operation phase
O mode ordinary mode

PDF probability densitiy function
PLL phase locked loop

SOA service-oriented architecture
SPECE solver for plasma electron cyclotron emission

TRAVIS tracing visualized

VMEC variational moments equilibrium code

W7-X Wendelstein 7-X

X mode extraordinary mode

Bibliography

[1] T. Hamacher et al. »Nuclear fusion and renewable energy forms: Are they compatible?« In: *Fusion Engineering and Design*, Vol. 88.6-8 (Oct. 2013), pages 657–660. DOI: 10.1016/j.fusengdes.2013.01.074.

[2] J. Ongena and G. Van Oost. »Energy for Future Centuries: Will Fusion Be an Inexhaustible, Safe, and Clean Energy Source?« In: *Fusion Science and Technology*, Vol. 45.2T (Mar. 2004), pages 3–14. DOI: 10.13182/fst04-a464.

[3] J. Scheffran. »Climate Change and Stability: The Case of the Mediterranean Region«. In: *International Federation of Automatic Control Proceedings Volumes*, Vol. 43.25 (Oct. 2010), pages 35–40. DOI: 10.3182/20101027-3-xk-4018.00009.

[4] A. A. Solomon et al. »How much energy storage is needed to incorporate very large intermittent renewables?« In: *Energy Procedia*, Vol. 135 (Oct. 2017), pages 283–293. DOI: 10.1016/j.egypro.2017.09.520.

[5] D. J. Ward. »Fusion as a future energy source«. In: *Europhysics News*, Vol. 47.5-6 (21st Nov. 2016), pages 32–34. DOI: 10.1051/epn/2016505.

[6] »Advantages of Fusion«. URL: https://www.iter.org/sci/Fusion (visited on 21st Apr. 2019).

[7] P. H. Rebut, R. J. Bickerton and B. E. Keen. »The Joint European Torus: installation, first results and prospects«. In: *Nuclear Fusion*, Vol. 25.9 (Sept. 1985), pages 1011–1022. DOI: 10.1088/0029-5515/25/9/003.

[8] H. Shirai et al. »Recent progress of the JT-60SA project«. In: *Nuclear Fusion*, Vol. 57.10, Art. 102002 (June 2017). DOI: 10.1088/1741-4326/aa5d01.

[9]   M. Hirsch et al. »Major results from the stellarator Wendelstein 7-AS«. In: *Plasma Physics and Controlled Fusion*, Vol. 50.5, Art. 053001 (10^th Mar. 2008). DOI: 10.1088/0741-3335/50/5/053001.

[10]   A. H. Boozer. »Theory of tokamak disruptions«. In: *Physics of Plasmas*, Vol. 19.5, Art. 058101 (May 2012). DOI: 10.1063/1.3703327.

[11]   R. J. Buttery et al. »Neoclassical tearing modes«. In: *Plasma Physics and Controlled Fusion*, Vol. 42.12B (Dec. 2000), B61–B73. DOI: 10.1088/0741-3335/42/12b/306.

[12]   R. J. La Haye. »Neoclassical tearing modes and their control«. In: *Physics of Plasmas*, Vol. 13.5, Art. 055501 (May 2006). DOI: 10.1063/1.2180747.

[13]   A. Bock et al. »Advanced tokamak investigations in full-tungsten ASDEX Upgrade«. In: *Physics of Plasmas*, Vol. 25.5, Art. 056115 (May 2018). DOI: 10.1063/1.5024320.

[14]   F. Warmer et al. »System Code Analysis of HELIAS-Type Fusion Reactor and Economic Comparison With Tokamaks«. In: *IEEE Transactions on Plasma Science*, Vol. 44.9 (Sept. 2016), pages 1576–1585. DOI: 10.1109/tps.2016.2545868.

[15]   U. Stroth et al. »Energy confinement scaling from the international stellarator database«. In: *Nuclear Fusion*, Vol. 36.8 (Aug. 1996), pages 1063–1077. DOI: 10.1088/0029-5515/36/8/i11.

[16]   H. Yamada et al. »Characterization of energy confinement in net-current free plasmas using the extended International Stellarator Database«. In: *Nuclear Fusion*, Vol. 45.12 (Dec. 2005), pages 1684–1693. DOI: 10.1088/0029-5515/45/12/024.

[17]   C. Beidler et al. »Physics and Engineering Design for Wendelstein VII-X«. In: *Fusion Technology*, Vol. 17.1 (1990), pages 148–168. DOI: 10.13182/FST90-A29178.

[18]   F. L. Hinton and R. D. Hazeltine. »Theory of plasma transport in toroidal confinement systems«. In: *Reviews of Modern Physics*, Vol. 48.2 (Apr. 1976), pages 239–308. DOI: 10.1103/RevModPhys.48.239.

[19] R. C. Wolf et al. »Wendelstein 7-X Program—Demonstration of a Stellarator Option for Fusion Energy«. In: *IEEE Transactions on Plasma Science*, Vol. 44.9 (Sept. 2016), pages 1466–1471. DOI: 10.1109/tps.2016.2564919.

[20] K. Lackner and R. Schneider. »The role of edge physics and confinement issues in the fusion reactor«. In: *Fusion Engineering and Design*, Vol. 22.1-2 (Mar. 1993), pages 107–116. DOI: 10.1016/s0920-3796(05)80012-2.

[21] D. Zhang et al. »First Observation of a Stable Highly Dissipative Divertor Plasma Regime on the Wendelstein 7-X Stellarator«. In: *Physical Review Letters*, Vol. 123.2, Art. 025002 (9[th] July 2019). DOI: 10.1103/PhysRevLett.123.025002.

[22] S. Sudo et al. »Scalings of energy confinement and density limit in stellarator/heliotron devices«. In: *Nuclear Fusion*, Vol. 30.1 (Jan. 1990), pages 11–21. DOI: 10.1088/0029-5515/30/1/002.

[23] P. Zanca et al. »A unified model of density limit in fusion plasmas«. In: *Nuclear Fusion*, Vol. 57.5, Art. 056010 (Mar. 2017). DOI: 10.1088/1741-4326/aa6230.

[24] T. Sunn Pedersen et al. »First results from divertor operation in Wendelstein 7-X«. In: *Plasma Physics and Controlled Fusion*, Vol. 61.1, Art. 014035 (27[th] Nov. 2018). DOI: 10.1088/1361-6587/aaec25.

[25] M. Greenwald. »Density limits in toroidal plasmas«. In: *Plasma Physics and Controlled Fusion*, Vol. 44.8 (Aug. 2002), R27–R53. DOI: 10.1088/0741-3335/44/8/201.

[26] J. Nührenberg and R. Zille. »Stable stellarators with medium β and aspect ratio«. In: *Physics Letters A*, Vol. 114.3 (10[th] Feb. 1986), pages 129–132. DOI: 10.1016/0375-9601(86)90539-6.

[27] T. Klinger et al. »Performance and properties of the first plasmas of Wendelstein 7-X«. In: *Plasma Physics and Controlled Fusion*, Vol. 59.1, Art. 014018 (18[th] Oct. 2016). DOI: 10.1088/0741-3335/59/1/014018.

[28] M. Krychowiak et al. »Overview of diagnostic performance and results for the first operation phase in Wendelstein 7-X (invited)«. In: *Review of Scientific Instruments*, Vol. 87.11, Art. 11D304 (27th Oct. 2016). DOI: 10.1063/1.4964376.

[29] R. C. Wolf et al. »Major results from the first plasma campaign of the Wendelstein 7-X stellarator«. In: *Nuclear Fusion*, Vol. 57.10, Art. 102020 (27th July 2017). DOI: 10.1088/1741-4326/aa770d.

[30] V. Erckmann et al. »Electron Cyclotron Heating for W7-X: Physics and Technology«. In: *Fusion Science and Technology*, Vol. 52.2 (1st Aug. 2007), pages 291–312. DOI: 10.13182/fst07-a1508.

[31] N. J. Fisch and A. H. Boozer. »Creating an Asymmetric Plasma Resistivity with Waves«. In: *Physical Review Letters*, Vol. 45.9 (1st Sept. 1980), pages 720–722. DOI: 10.1103/physrevlett.45.720.

[32] Y. Turkin et al. »Current Control by ECCD for W7-X«. In: *Fusion Science and Technology*, Vol. 50.3 (Oct. 2006), pages 387–394. DOI: 10.13182/fst06-5.

[33] M. Zanini et al. »ECCD operations in the second experimental campaign at W7-X«. In: *European Physical Journal Web of Conferences*, Vol. 203, Art. 02013 (25th Mar. 2019). DOI: 10.1051/epjconf/201920302013.

[34] J. I. Paley et al. »From profile to sawtooth control: developing feedback control using ECRH/ECCD systems on the TCV tokamak«. In: *Plasma Physics and Controlled Fusion*, Vol. 51.12, Art. 124041 (12th Nov. 2009). DOI: 10.1088/0741-3335/51/12/124041.

[35] T. Stange et al. »First demonstration of fully equilibrated ions and electrons in a Magnetically Confined High Temperature Plasma sustained by 2nd harmonic O-mode polarized ECRH only«. In: *Physical Review Letters*, (2020). Submitted.

[36] H. J. Hartfuss and T. Geist. *Fusion Plasma Diagnostics with mm-Waves: An Introduction.* Weinheim: Wiley-VCH, Sept. 2013. DOI: 10.1002/9783527676279.

[37] M. Hirsch et al. »ECE Diagnostic for the initial Operation of Wendelstein 7-X«. In: *European Physical Journal Web of Conferences*, Vol. 203, Art. 03007 (25th Mar. 2019). doi: 10.1051/epjconf/20192030300 7.

[38] A. Jacchia et al. »Determination of diffusive and nondiffusive transport in modulation experiments in plasmas«. In: *Physics of Fluids B: Plasma Physics*, Vol. 3.11 (1st Nov. 1991), pages 3033–3040. doi: 10.1063/ 1.859781.

[39] B. P. van Milligen et al. »Study of radial heat transport in W7-X using the Transfer Entropy«. In: *Nuclear Fusion*, (19th Apr. 2018). doi: 10.1088/1741-4326/aabf5d.

[40] M. Bornatici et al. »Electron cyclotron emission and absorption in fusion plasmas«. In: *Nuclear Fusion*, Vol. 23.9 (Sept. 1983), pages 1153–1257. doi: 10.1088/0029-5515/23/9/005.

[41] U. Höfel et al. »Bayesian Modelling of Microwave Radiometer Calibration on the example of the Wendelstein 7-X Electron Cyclotron Emission diagnostic«. In: *Review of Scientific Instruments*, Vol. 90.4, Art. 043502 (11th Apr. 2019). doi: 10.1063/1.5082542.

[42] V. D. Shafranov. »Equilibrium of a toroidal pinch in a magnetic field«. In: *Soviet Atomic Energy*, Vol. 13.6 (Nov. 1963), pages 1149–1158. doi: 10.1007/bf01312317.

[43] N. B. Marushchenko, Y. Turkin and H. Maassberg. »Ray-tracing code TRAVIS for ECR heating, EC current drive and ECE diagnostic«. In: *Computer Physics Communications*, Vol. 185.1 (2014), pages 165–176. doi: 10.1016/j.cpc.2013.09.002.

[44] D. Walsh. »On Zheleznyakov's Equation of Radiative Transfer in a Magnetoactive Plasma«. In: *The Astrophysical Journal*, Vol. 159 (Feb. 1970), page 733. doi: 10.1086/150349.

[45] S. Bozhenkov et al. »Service oriented architecture for scientific analysis at W7-X. An example of a field line tracer«. In: *Fusion Engineering and Design*, Vol. 88.11 (1st Aug. 2013), pages 2997–3006. doi: 10.1016/ j.fusengdes.2013.07.003.

[46]   M. Grahl et al. »Archive WEB API: A web service for the experiment data archive of Wendelstein 7-X«. In: *Fusion Engineering and Design*, Vol. 123 (1st Mar. 2017). Proceedings of the 29th Symposium on Fusion Technology, pages 1015–1019. DOI: 10.1016/j.fusengdes.2017.02.047.

[47]   S. A. Bozhenkov et al. »The Thomson scattering diagnostic at Wendelstein 7-X and its performance in the first operation phase«. In: *Journal of Instrumentation*, Vol. 12.10, Art. P10004 (10th Oct. 2017). DOI: 10.1088/1748-0221/12/10/P10004.

[48]   A. Langenberg et al. »Forward Modeling of X-Ray Imaging Crystal Spectrometers Within the Minerva Bayesian Analysis Framework«. In: *Fusion Science and Technology*, Vol. 69.2 (Apr. 2016), pages 560–567. DOI: 10.13182/FST15-181.

[49]   H. Damm. *Upgrade of the Wendelstein 7-X Thomson Scattering Diagnostic to Study Short Transient Plasma Effects. A Demonstration on Pellet Injection for Stellarator Fuelling.* Master's thesis. University of Greifswald, 5th June 2018, pages 46–49. 199 pages. DOI: 10.17617/2.2596741.

[50]   S. Schmuck et al. »Design of the ECE diagnostic at Wendelstein 7-X«. In: *Fusion Engineering and Design*, Vol. 84.7–11 (June 2009). Proceeding of the 25th Symposium on Fusion Technology, pages 1739–1743. DOI: 10.1016/j.fusengdes.2008.12.094.

[51]   D. Wagner et al. »Sub-THz notch filters based on photonic bandgaps in overmoded waveguides«. In: *40th International Conference on Infrared, Millimeter, and Terahertz waves (IRMMW-THz)*. IEEE, Aug. 2015. DOI: 10.1109/irmmw-thz.2015.7327739.

[52]   C. Fuchs and H.-J. Hartfuss. »Extreme broadband multichannel ECE radiometer with "zoom" device«. In: *Review of Scientific Instruments*, Vol. 72.1 (3rd Jan. 2001), pages 383–386. DOI: 10.1063/1.1309005.

[53]   X. Liu et al. »Absolute intensity calibration of the 32-channel heterodyne radiometer on experimental advanced superconducting tokamak«. In: *Review of Scientific Instruments*, Vol. 85.9, Art. 093508 (Sept. 2014). DOI: 10.1063/1.4896047.

[54]   J. L. Ségui et al. »An upgraded 32-channel heterodyne electron cyclo-
       tron emission radiometer on Tore Supra«. In: *Review of Scientific In-
       struments*, Vol. 76.12, Art. 123501 (14th Dec. 2005). doi: 10.1063/1.
       2140225.

[55]   E. T. Jaynes. *Probability theory. The Logic of Science.* Edited by G. L.
       Bretthorst. Cambridge University Press, 10th Apr. 2003. 727 pages.
       doi: 10.1017/CBO9780511790423.

[56]   D. Sivia and J. Skilling. *Data Analysis. A Bayesian Tutorial.* 2nd ed.
       Oxford University Press, 11th June 2006. 260 pages.

[57]   J. Svensson and JET EFDA contributors. *Non-parametric Tomo-
       graphy Using Gaussian Processes.* Research report. Abingdon, United
       Kingdom: Culham Centre for Fusion Energy, 2011. 15 pages. url: http:
       //www.euro-fusionscipub.org/wp-content/uploads/
       eurofusion/EFDP11024.pdf (visited on 20th Feb. 2019).

[58]   C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Ma-
       chine Learning.* MIT Press, 23rd Nov. 2005. 272 pages. url: https://
       mitpress.mit.edu/books/gaussian-processes-machine-
       learning (visited on 23rd Feb. 2019).

[59]   M. N. Gibbs. *Bayesian Gaussian Processes for Regression and Classific-
       ation.* PhD thesis. University of Cambridge, Sept. 1997. 128 pages. url:
       http://citeseerx.ist.psu.edu/viewdoc/download?doi=
       10.1.1.147.1130&rep=rep1&type=pdf (visited on 26th Mar.
       2019).

[60]   D. Higdon, J. Swall and J. Kern. »Non-Stationary Spatial Modeling«.
       In: *Bayesian Statistics. Proceedings of the Sixth Valencia International
       Meeting.* Edited by J. M. Bernardo et al. Vol. 6. Oxford University
       Press, 12th Aug. 1999. url: https://global.oup.com/academic/
       product/bayesian-statistics-6-9780198504856 (visited on
       22nd Feb. 2019).

[61]   C. J. Paciorek. *Nonstationary Gaussian processes for regression and
       spatial modelling.* PhD thesis. Carnegie Mellon University, May 2003.
       236 pages. url: https://www.stat.berkeley.edu/~paciorek/
       diss/paciorek-thesis.pdf (visited on 25th Feb. 2019).

[62] M. A. CHILENSKI et al. »Improved profile fitting and quantification of uncertainty in experimental measurements of impurity transport coefficients using Gaussian process regression«. In: *Nuclear Fusion*, Vol. 55.2, ART. 023012 (23rd Jan. 2015). DOI: 10.1088/0029-5515/55/2/023012.

[63] M. AVRIEL. *Nonlinear Programming: Analysis and Methods*. Dover Publications, 11th Sept. 2003. 512 pages.

[64] R. HOOKE and T. A. JEEVES. »"Direct Search" Solution of Numerical and Statistical Problems«. In: *Journal of the Association for Computing Machinery*, Vol. 8.2 (Apr. 1961), pages 212–229. DOI: 10.1145/321062.321069.

[65] N. METROPOLIS et al. »Equation of State Calculations by Fast Computing Machines«. In: *The Journal of Chemical Physics*, Vol. 21.6 (1st June 1953), pages 1087–1092. DOI: 10.1063/1.1699114.

[66] W. K. HASTINGS. »Monte Carlo sampling methods using Markov chains and their applications«. In: *Biometrika*, Vol. 57.1 (1st Apr. 1970), pages 97–109. DOI: 10.1093/biomet/57.1.97.

[67] H. HAARIO, E. SAKSMAN and J. TAMMINEN. »An adaptive Metropolis algorithm«. In: *Bernoulli*, Vol. 7.2 (Feb. 2001), pages 223–242. URL: https://projecteuclid.org:443/euclid.bj/1080222083 (visited on 3rd Jan. 2019).

[68] D. NGUYEN et al. »Automatic adaptation of MCMC algorithms«. In: *ArXiv e-prints*, (24th Feb. 2018). arXiv: 1802.08798 [stat.ME].

[69] C. J. GEYER. »Introduction to Markov Chain Monte Carlo«. In: *Handbook of Markov Chain Monte Carlo*. Edited by S. BROOKS et al. Handbooks of Modern Statistical Methods. Chapman and Hall/CRC, 10th May 2011. Chapter 1. DOI: 10.1201/b10905-2.

[70] G. O. ROBERTS, A. GELMAN and W. R. GILKS. »Weak convergence and optimal scaling of random walk Metropolis algorithms«. In: *The Annals of Applied Probability*, Vol. 7.1 (Feb. 1997), pages 110–120. DOI: 10.1214/aoap/1034625254.

[71]   M. Bédard. »Optimal acceptance rates for Metropolis algorithms: Moving beyond 0.234«. In: *Stochastic Processes and their Applications*, Vol. 118.12 (31st Dec. 2008), pages 2198–2222. DOI: 10.1016/j.spa.2007.12.005.

[72]   E. Parzen. »On Estimation of a Probability Density Function and Mode«. In: *The Annals of Mathematical Statistics*, Vol. 33.3 (Sept. 1962), pages 1065–1076. DOI: 10.1214/aoms/1177704472.

[73]   B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. 1st ed. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1st Apr. 1986. 176 pages. URL: https://www.crcpress.com/Density-Estimation-for-Statistics-and-Data-Analysis/Silverman/p/book/9780412246203 (visited on 17th Mar. 2019).

[74]   J. Svensson, A. Werner and JET EFDA contributors. »Large Scale Bayesian Data Analysis for Nuclear Fusion Experiments«. In: *2007 IEEE International Symposium on Intelligent Signal Processing* (3rd–5th Oct. 2007). Alcalá de Henares, Madrid, Spain, Oct. 2007. DOI: 10.1109/WISP.2007.4447579.

[75]   J. Svensson et al. »Modelling of JET Diagnostics Using Bayesian Graphical Models«. In: *Contributions to Plasma Physics*, Vol. 51.2-3 (16th Mar. 2011), pages 152–157. DOI: 10.1002/ctpp.201000058.

[76]   M. J. Hole et al. »The use of Bayesian inversion to resolve plasma equilibrium«. In: *Review of Scientific Instruments*, Vol. 81.10, Art. 10E127 (27th Oct. 2010). DOI: 10.1063/1.3491044.

[77]   C. M. Bishop. *Graphical Models*. In: *Pattern Recognition and Machine Learning*. 1st ed. Information Science and Statistics. Springer-Verlag New York, 17th Aug. 2006. Chapter 8, pages 359–362. URL: http://www.ebook.de/de/product/5324937/christopher_m_bishop_pattern_recognition_and_machine_learning.html.

[78]   A. Pavone et al. »Bayesian uncertainty calculation in neural network inference of ion and electron temperature profiles at W7-X«. In: *Review of Scientific Instruments*, Vol. 89.10, Art. 10K102 (6th Aug. 2018). DOI: 10.1063/1.5039286.

[79] Version svn checkout -r13781. URL: https://svn.minerva-centra
      l.net/minerva/software/minerva/minervaIII/applicati
      ons/minerva-ece-w7x-profile-app.

[80] A. Tarantola. *Inverse Problem Theory and Methods for Model Para-
      meter Estimation*. Other Titles in Applied Mathematics. Society for In-
      dustrial and Applied Mathematics, 2005. DOI: 10.1137/1.97808987
      17921.

[81] J. Svensson, A. Werner and JET-EFDA Contributors. »Current
      tomography for axisymmetric plasmas«. In: *Plasma Physics and Con-
      trolled Fusion*, Vol. 50.8, Art. 085002 (29[th] May 2008). DOI: 10.1088/
      0741-3335/50/8/085002.

[82] O. P. Ford. *Tokamak Plasma Analysis through Bayesian Diagnostic Mod-
      elling*. PhD thesis. University of London, 8[th] Nov. 2010. 179 pages. URL:
      http://www.oliford.co.uk/phys/thesis-final.pdf (visited
      on 24[th] Feb. 2019).

[83] G. Taylor et al. »Update on the status of the ITER ECE diagnostic
      design«. In: *European Physical Journal Web of Conferences*, Vol. 147,
      Art. 02003 (24[th] July 2017). Edited by B. K. Shukla, H. B. Pandya
      and S. L. Rao. DOI: 10.1051/epjconf/201714702003.

[84] X. Han et al. »Design and characterization of a 32-channel heterodyne
      radiometer for electron cyclotron emission measurements on experi-
      mental advanced superconducting tokamak«. In: *Review of Scientific
      Instruments*, Vol. 85.7, Art. 073506 (July 2014). DOI: 10.1063/1.
      4891040.

[85] T. Stange et al. »Advanced electron cyclotron heating and current
      drive experiments on the stellarator Wendelstein 7-X«. In: *European
      Physical Journal Web of Conferences*, Vol. 157, Art. 02008 (23[rd] Oct.
      2017). DOI: 10.1051/epjconf/201715702008.

[86] H. P. Laqua et al. »Overview of W7-X ECRH Results«. In: *EPJ Web of
      Conferences*, Vol. 203 (25[th] Mar. 2019). Edited by E. Poli, H. Laqua and J.
      Oosterbeek, page 02002. DOI: 10.1051/epjconf/201920302002.

[87] K.-J. Brunner et al. »Real-time dispersion interferometry for density feedback in fusion devices«. In: *Journal of Instrumentation*, Vol. 13.09, Art. P09002 (3rd Sept. 2018). doi: 10 . 1088 / 1748 – 0221 / 13 / 09 / P09002.

[88] J. van den Berg et al. »Fast analysis of collective Thomson scattering spectra on Wendelstein 7-X«. In: *Review of Scientific Instruments*, Vol. 89.8, Art. 083507 (17th Aug. 2018). doi: 10 . 1063 / 1 . 5035416.

[89] R. C. Wolf and the Wendelstein 7-X team. »A stellarator reactor based on the optimization criteria of Wendelstein 7-X«. In: *Fusion Engineering and Design*, Vol. 83.7-9 (Dec. 2008), pages 990–996. doi: 10 . 1016/j.fusengdes.2008.05.008.

[90] R. C. Wolf et al. »Electron-cyclotron-resonance heating in Wendelstein 7-X: A versatile heating and current-drive method and a tool for in-depth physics studies«. In: *Plasma Physics and Controlled Fusion*, Vol. 61.1, Art. 014037 (27th Nov. 2018). doi: 10 . 1088 / 1361 – 6587/aaeab2.

[91] S. Schmuck et al. »Bayesian derivation of electron temperature profile using JET ECE diagnostics«. European Physical Society, 38th Conference on Plasma Physics. Strasbourg, France, 2011. url: http://ocs.ciem at.es/EPS2011PAP/pdf/P5.046.pdf (visited on 23rd Apr. 2018).

[92] S. Schmuck et al. »Electron Temperature and Density Inferred from JET ECE Diagnostics«. European Physical Society, 41st Conference on Plasma Physics. Berlin, Germany, 2014. url: http://ocs.ciemat. es/EPS2014PAP/pdf/P1.025.pdf (visited on 23rd Apr. 2018).

[93] S. Schmuck et al. »Bayesian Inference Using JETs Microwave Diagnostic System«. In: *Nuclear Fusion*, (2019). Submitted.

[94] S. K. Rathgeber et al. »Estimation of edge electron temperature profiles via forward modelling of the electron cyclotron radiation transport at ASDEX Upgrade«. In: *Plasma Physics and Controlled Fusion*, Vol. 55.2, Art. 025004 (Feb. 2013). doi: 10 . 1088/0741–3335/55/2/025004.

[95] S. S. Denk et al. »Radiation transport modelling for the interpretation of oblique ECE measurements«. In: *European Physical Journal Web of Conferences*, Vol. 147, Art. 02002 (24th July 2017). Edited by B. K.

Shukla, H. B. Pandya and S. L. Rao. doi: 10.1051/epjconf/201714702002.

[96] S. S. Denk et al. »Analysis of electron cyclotron emission with extended electron cyclotron forward modeling«. In: *Plasma Physics and Controlled Fusion*, Vol. 60.10, Art. 105010 (7th Sept. 2018). doi: 10.1088/1361-6587/aadb2f.

[97] D. Farina et al. »SPECE: a code for Electron Cyclotron Emission in tokamaks«. In: *AIP Conference Proceedings*, Vol. 988.1 (2008), pages 128–131. doi: 10.1063/1.2905053.

[98] K. Rahbarnia et al. »Diamagnetic energy measurement during the first operational phase at the Wendelstein 7-X stellarator«. In: *Nuclear Fusion*, Vol. 58.9, Art. 096010 (4th July 2018). doi: 10.1088/1741-4326/aacab0.

[99] M. Grahl et al. »Archive WEB API: A web service for the experiment data archive of Wendelstein 7-X«. In: *Fusion Engineering and Design*, Vol. 123 (Nov. 2017), pages 1015–1019. doi: 10.1016/j.fusengdes.2017.02.047.

[100] A. Schlaich. *Time-dependent spectrum analysis of high power gyrotrons.* PhD thesis. Karlsruhe: Karlsruhe Institute of Technology, 6th Aug. 2015. 286 pages. doi: 10.5445/ksp/1000046919.

[101] M. Peters, G. Gorini and P. Mantica. »Optical thickness corrections to transient ECE temperature measurements in tokamak and stellarator plasmas«. In: *Nuclear Fusion*, Vol. 35.7 (1995), page 873. doi: 10.1088/0029-5515/35/7/I10.

[102] S. K. Rathgeber. *Electron Temperature and Pressure at the Edge of ASDEX Upgrade Plasmas. Estimation via electron cyclotron radiation and investigations on the effect of magnetic perturbations.* PhD thesis. Ludwig-Maximilian-Universität München, 28th Feb. 2013. 104 pages. URN: urn:nbn:de:bvb:19-157614.

[103] M. Dibon et al. »Blower Gun pellet injection system for W7-X«. In: *Fusion Engineering and Design*, Vol. 98-99 (2015): *Proceedings of the 28th Symposium On Fusion Technology (SOFT-28)*, pages 1759–1762. doi: 10.1016/j.fusengdes.2015.01.050.

[104] T. Möller and B. Trumbore. »Fast, Minimum Storage Ray-Triangle Intersection«. In: *Journal of Graphics Tools*, Vol. 2.1 (1997), pages 21–28. DOI: 10.1080/10867651.1997.10487468.

[105] M. Hirsch et al. »Microwave and Interferometer Diagnostics prepared for first Plasma Operation of Wendelstein 7-X«. In: 1st EPS conference on Plasma Diagnostics (Frascati, Italy, 14th–17th Apr. 2015). Vol. 240. Proceedings of Science. 26th Oct. 2016. DOI: 10.22323/1.240.0111.

[106] A. Pavone et al. »Neural network approximation of Bayesian models for the inference of ion and electron temperature profiles at W7-X«. In: *Plasma Physics and Controlled Fusion*, Vol. 61.7, Art. 075012 (30th May 2019). DOI: 10.1088/1361-6587/ab1d26.

[107] C. M. Bishop and J. W. Connor. »Heat-pulse propagation in tokamaks and the role of density perturbations«. In: *Plasma Physics and Controlled Fusion*, Vol. 32.3 (Mar. 1990), pages 203–211. DOI: 10.1088/0741-3335/32/3/005.

[108] K. K. Kirov et al. »Analysis of ECRH switch on/off events in ASDEX Upgrade«. In: *Plasma Physics and Controlled Fusion*, Vol. 48.2 (Jan. 2006), pages 245–262. DOI: 10.1088/0741-3335/48/2/006.

[109] G. M. Weir. *Thermal Transport by Heat Pulse Propagation on HSX. Motivation, Simulation, and Experimental Preparation*. Master's thesis. University of Wisconsin-Madison, 2011. 60 pages.

[110] Version svn checkout -r13803. URL: https://svn.minerva-central.net/minerva/software/minerva/minervaIII/applications/minerva-ece-w7x-apps.

[111] Version svn checkout -r13862. URL: https://svn.minerva-central.net/minerva/software/minerva/minervaIII/applications/minerva-ece-w7x-profile-app.