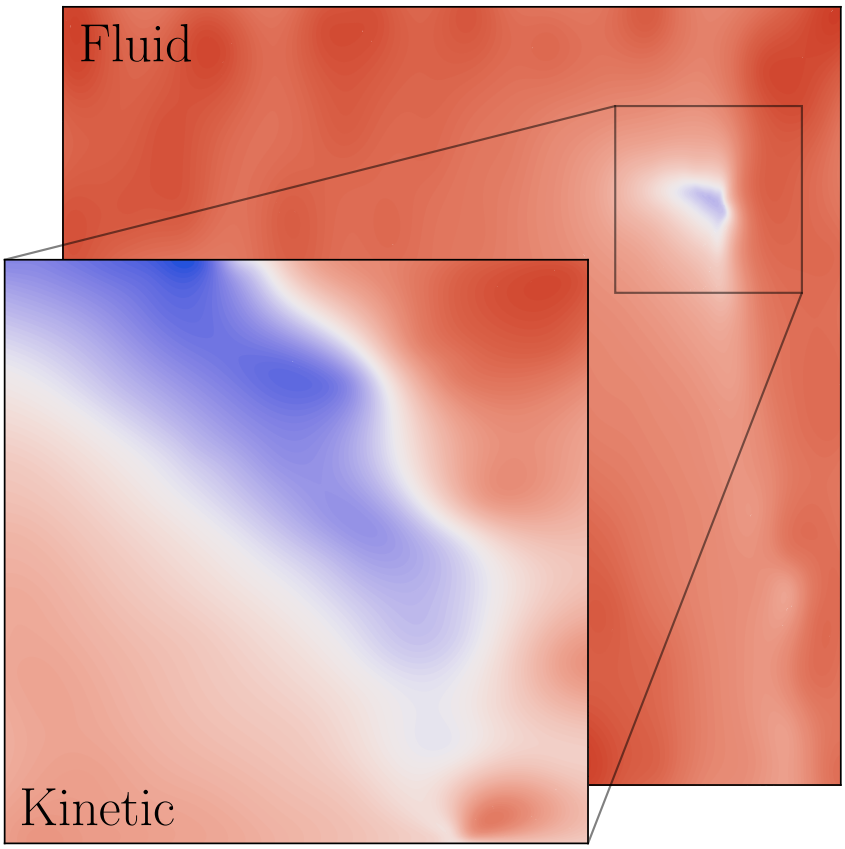

Multi-objective Bayesian optimization of tokamak disruptions using fluid and kinetic models

Thesis for the degree of Master of Science in Complex Adaptive Systems

IDA EKMARK

DEPARTMENT OF PHYSICS

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2023
www.chalmers.se

Master's thesis 2023

Multi-objective Bayesian optimization of tokamak disruptions using fluid and kinetic models

IDA EKMARK

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Physics
Division of Subatomic, High Energy and Plasma Physics
Plasma Theory Research Group
Chalmers University of Technology
Gothenburg, Sweden 2023

Multi-objective Bayesian optimization of tokamak disruptions using fluid and kinetic models
IDA EKMARK

Supervisor: István Pusztai, Department of Physics
Examiner: Tünde Fülöp, Department of Physics

Master's Thesis 2023
Chalmers University of Technology
Department of Physics
Division of Subatomic, High Energy and Plasma Physics
Plasma Theory Research Group
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Gaussian process approximations of the cost function from the optimization using both fluid and kinetic plasma models, as a function of injected deuterium (horizontal axis) and neon (vertical axis) densities. Red represents cost function values corresponding to a dangerous disruption scenario, while blue represents safe scenarios. The large contour plot in the background has been generated from the optimization using the fluid model, and only has a very small area of safe values. In the front, the contour plot generated with the kinetic model is presented, showing a significantly larger area of safe values. The black box indicates in which part of the input space of the fluid optimization the kinetic optimization has been performed.

Typeset in LaTeX
Gothenburg, Sweden 2023

Multi-objective Bayesian optimization of tokamak disruptions using fluid and kinetic models

IDA EKMARK
Department of Physics
Chalmers University of Technology

Abstract

The generation of highly energetic runaway electron beams during tokamak disruptions is a major challenge facing tokamak fusion reactors. One of the most studied disruption mitigation schemes is massive material injection. Finding injected material densities, such that the consequences of the resulting disruption – runaway electron impact, localized heat losses and mechanical stresses – are tolerable, is still an open question, and it represents a multi-objective optimization problem. We have used a Bayesian optimization framework to optimize the injected densities of deuterium and neon in a non-activated ITER-like tokamak set up. The cost function was constructed systematically to maximize information gain, combining the maximum runaway current, final ohmic current, current quench time and conducted thermal losses. The simulations of plasma evolution were performed using the disruption modelling tool DREAM. Optimization of the developed cost function was performed in two layers of physics fidelity, using both fluid and kinetic plasma models. The fluid model is computationally less expensive, which is advantageous for exploring a large parameter space. Once promising parameter regions are located using a wide search with fluid models, these are further studied in higher physics fidelity using kinetic simulations. These simulations resolve the energy distribution of the fast electrons allowing us to also account for fast electron impact ionization and energy transfer. Using two layers, the advantages of each model can be utilized resulting in an efficient optimization with a reliable examination of relevant areas. Additionally, a qualitative comparison of the two models was made to illuminate the differences between the two layers. In general, the kinetic model generated more optimistic results for the disruption consequences. More specifically, the kinetic model favoured higher neon densities and slightly lower deuterium densities compared to the fluid model. In both models, the optima are fairly insensitive to the radial distribution of neon as long as there is a higher neon density at the edge. Furthermore, the optima occurred for a moderately core-localized deuterium density. The explanation for the differences between the fluid and kinetic models was concluded to be that the fluid model overestimates the hot-tail runaway generation for certain injected material densities, resulting in larger runaway currents.

Keywords: fusion plasma, disruption mitigation, runaway electron, material injection, Bayesian optimization, fluid kinetic model

Acknowledgements

I would like to express my deepest gratitude to members of the Plasma Theory Research Group at Chalmers, without whom this thesis could not have been completed. Firstly, I would convey my sincerest appreciation to my examiner and mentor Tünde, for introducing me to plasma physics and the field of magnetically confined fusion, for all the opportunities I have been presented and for her endless support and encouragement. I would also like to wholeheartedly thank my supervisor István, for always being there to aid, guide and advice me on this exploration of tokamak disruptions and massive material injection, as well as for the genuine interest he has shown in the evolution of my work and the thorough review of my thesis from start to finish. Additionally, this project could not have been performed without the help from my mentor Mathias, who is always there to discuss and answer questions about plasma physics, programming, life in academia and everything in between, as well as for challenging me to reach higher potentials and sharing great ideas and suggestions for how to proceed forward in my research. I would like to thank Hannes and Peter for sharing their work with me and letting me continue the research they started; Oscar for taking the time to answer my questions and discussing problems that arose, even during lunch and fika breaks; Patrik for the discussions about how to design my objective function; and the Plasma Theory Research group for letting me be part of this amazing and supportive work environment. Furthermore, a special thanks to opponent David for reading my thesis and giving me great suggestions for how to improve it.

Finally, I would like to thank my friends and family for their unlimited support and heartfelt encouragement.

Ida Ekmark, Gothenburg, April 2023

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis, listed in alphabetical order:

CQ	Current quench
EI	Estimated improvement
GP	Gaussian process
MMI	Massive material injection
RE	Runaway electron
TQ	Thermal quench

Nomenclature

Below is a list of the parameters that have been used throughout this thesis, including their typical units.

I_p	Plasma current [A]
I_re	Runaway current [A]
I_Ω	Ohmic current [A]
R_0	Major radius [m]
a	Minor radius [m]
b	Wall radius [m]
B_0	Toroidal magnetic field on axis [T]
$\delta B/B$	Magnetic perturbation normalized to the unperturbed magnetic field [%]
V_p	Plasma volume [m^3]
T	Temperature [eV]
n	Particle density [m^{-3}]
c	Radial density variation parameter
τ_CQ	Current quench time [s]
η_cond	Fraction of initial thermal energy lost from the plasma through radial heat transport [%]
$\mathcal{L}$	Cost function
μ	GP-approximation of cost function (or cost function components)
$\mathbf{E}$	Electric field [V/m]
$\mathbf{B}$	Magnetic field [T]
θ_p	Pitch angle – angle between magnetic field and electron velocity [rad]
P	Conducted power [J/s]

Contents

1

Introduction

As society progresses, the need for energy is constantly increasing, while global warming is putting a significant constraint on which energy sources can be utilized in the future. A possible energy source which does not put a strain on the environment while having the potential of supporting our future energy needs is fusion energy. Nuclear fusion is the energy source powering the sun, and it is the mechanism of two lighter nuclei fusing into a single nuclei under circumstances of high temperature and pressure. In this process, energy is released due to the reaction products having a smaller combined mass than that of the two initial nuclei. Hydrogen isotopes, such as deuterium and tritium, are viable fusion fuels. Since deuterium can be found in abundance in our oceans and tritium can be produced from lithium which can be mined, fusion energy would be able to satisfy our energy needs for the foreseeable future [1]. There is no long-lived nuclear waste or significant green house gas emissions produced since the major by-product is the stable isotope helium-4. Furthermore, it would be a non-intermittent energy source since it is not dependent on any fluctuating and unpredictable factor such as the weather [2]. Consequently, fusion would be an ideal source of energy if it were to be developed.

The difficulty with fusion is to create the conditions for the reaction to happen – to sustain the temperatures and pressures of the sun for a longer time is unattainable with our technology today, and thus alternative approaches are needed for the development of fusion reactors. Fusion energy research is today on the verge of a new era because of the current development of two fusion devices, SPARC [3] and ITER [4], with the purpose of being the first to produce net energy gain, and they are expected to be operational within a couple of years. Both devices are tokamaks, which is a proposed fusion reactor using magnetic confinement of the the fusion fuel, which must be in the form of plasma to be hot enough for the fusion reaction to be possible. As the name suggests, the magnetic confinement uses magnetic fields to confine the fusion plasma. Tokamaks utilize this approach by shaping the plasma into a torus with both poloidal and toroidal magnetic fields, see figure 1.1. This requires driving a high – several megaampere – current through the plasma, which is achieved by an induced electric field. For tokamaks to be able to operate safely, there are still a number of challenges waiting to be solved.

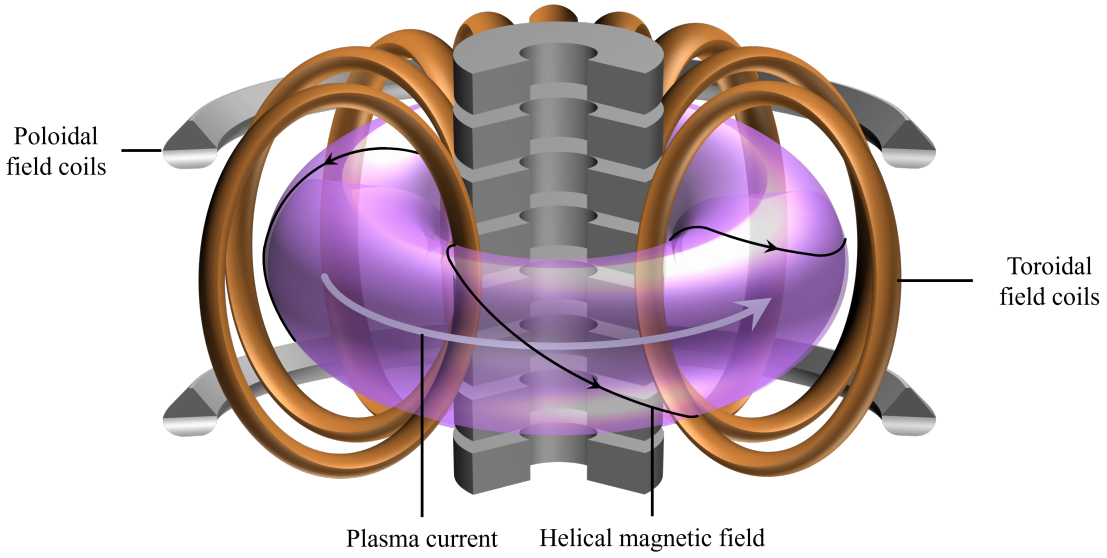

Figure 1.1: Visualization of a tokamak, showing the helical magnetic field and the plasma current.

One of the most important challenges is the occurrence of disruptions. Disruptions are the events when confinement of the plasma energy is lost due to instabilities or system malfunctions. During disruptions there can be rapid and concentrated depositions of heat causing some components of the tokamak to melt or otherwise get damaged, and there are also risks of mechanical stresses due to electromagnetic forces between the plasma and tokamak device [5]. During disruptions, runaway electrons can be generated, which are problematic when they make up a significant fraction of the plasma current in the tokamak. Runaway electrons are electrons which are accelerated by the electric field in the fusion device without restraint, causing them to become relativistic [6]. Such a beam of relativistic electrons is difficult to manage or subdue, and can cause serious damage to the wall of the device as previously mentioned [7]. It is imperative to mitigate or avoid the occurrence of significant runaway beams and other negative effects of the disruption. The most studied approaches for mitigation of disruptions are material injection into the plasma or inducing magnetic perturbations. Massive material injection (MMI) is the injection of cold gases or cryogenic pellets of material to quickly cool down the plasma, ideally causing the thermal energy to be lost due to radiation [8]. Magnetic perturbations cause runaway electrons to be transported into the wall in a more controlled fashion, preferably before they become too energetic or numerous, which avoids the problem of rapid deposition of energy into a concentrated area of the tokamak wall.

The runaway electron generation rate is exponentially sensitive to the initial plasma current, and as such, is expected to be an incomparably larger problem in reactor-scale devices than in today's experiments. With the risks posed to the fusion device in mind, it is necessary to develop viable strategies in numerical rather than physical experiments. There are several models for describing plasma physics; single particle models, kinetic models and fluid models are three examples [9]. Single particle

modelling describes the evolutions of each particle in the plasma based on electromagnetic theory such as the Lorentz force law, which becomes deeply complex and computationally demanding when many particles are involved. Fluid modelling describes each particle species of the plasma as a fluid with local densities, velocities and temperatures, which are evolved by fluid equations (integral moments of the kinetic equation). The most fundamental model, apart from single-particle modelling, is kinetic plasma theory where the distribution function of the plasma is evaluated and evolved in the phase space using kinetic equations.

Returning to runaway electron mitigation, massive material injection can be implemented using several methods and parameter combinations – for example which materials to inject and the amount of each material. As of yet, it is unknown when MMI is successful in mitigating the risks of disruptions, and when it fails. It would be highly beneficial to study massive material injection, for example exploring the injected density parameter space to find regions of safe operation during a tokamak disruption, and which parameters result in the safest disruption scenario. A suitable method, which simultaneously searches for optima while developing an approximation for an objective function, is Bayesian optimization, which is an optimization concept based on Bayes theorem and stochastic processes [10]. Bayesian optimization has previously been used to investigate massive material injection as a runaway mitigation strategy during tokamak disruptions in the master thesis of Hannes Bergström and Peter Halldestam [11] and the subsequent article by István Pusztai et al [12]. In these works, massive material injection of both radially uniform and varying deuterium and neon densities were explored using fluid plasma models. The aim of this thesis is to go beyond these works by doing a more systematic construction of the cost function and performing the optimization in two layers of physics fidelity, using both fluid and kinetic plasma models. Expressly, the objectives of the project are as follows.

- To systematically develop an informative cost function to reliably quantify disruption evolutions of massive material injection scenarios following good optimization practises.

- To find safe operational regions and the optimum of this cost function in the input space spanned by the injected material densities, as well as their radial variation, using fluid plasma models.

- To explore the discovered safe operational regions using both fluid and kinetic plasma models, and perform a qualitative analysis of the differences between the two models.

- To study the electron energy and pitch-angle distribution of the discovered optimum and a runaway-dominated case using the kinetic model.

- To examine the significance of fast electron impact ionization in a high and low runaway case.

1.1 Thesis outline

The theoretical background that this thesis is based on is presented in chapter 2. Firstly, some relevant aspects of physics related to plasmas and tokamaks will be reviewed in section 2.1. Here, the risks associated with disruptions will be explained in deeper detail than in the current chapter, and the concept of runaway electrons together with relevant runaway generation mechanisms will be described. In section 2.2, the key factors of Bayesian optimization as well as the basic logic behind it will be explained.

Before moving to the results, important aspects related to the simulations performed in the thesis will be presented in chapter 3 – here the numerical tool DREAM and the physical disruption model will be covered. Both the results and methods related to the optimization are presented in chapter 4. To begin with, in section 4.1, the derivation of the cost function will be explained. This discussion includes a presentation of the components from which the cost function is constructed, as well as how these are combined. Following this, the optimization algorithm is described in section 4.3, including both the specifications regarding the Gaussian processes, Bayesian optimization and moving between layers of physics fidelity, as well as the black box composition. The results from the optimization of uniformly distributed and radially varying injected densities are presented in section 4.4 and 4.5, respectively.

In chapter 5, the studies of fast electron impact ionization and electron energy and pitch angle distributions are presented for two specific density combinations, together with an analysis of the runaway generation disparities between the fluid and kinetic models. Finally, the conclusions are summarized in chapter 6 together with a review of possibilities for future research.

2

Theoretical background

This thesis utilizes Bayesian optimization for studying massive material injection scenarios during tokamak disruptions, and as such basic knowledge of both the physics related to tokamak disruptions and Bayesian optimization is needed. In this chapter, relevant aspects of plasma and tokamak physics will be reviewed and the principles of Bayesian optimization will be described.

2.1 Plasmas and tokamak physics

This section will cover the plasma physics necessary for this thesis, including some basic plasma concepts as well as different means of modelling plasmas. Additionally, the tokamak concept for magnetic confinement will be reviewed. Finally, runaway electrons will be explained, including how they are generated in a tokamak.

2.1.1 Basic plasma physics

Plasma is the fourth fundamental state of matter – being the state of higher energy than gas. It characterized by being sufficiently ionized for electromagnetic collective effects to dominate its dynamics, rather than atomic collisions. A plasma is defined as being a macroscopically neutral gas containing charged particles that exhibit collective behaviour due to the long range electromagnetic forces of the charged particles [9].

Charged particles will in the presence of electric and magnetic fields experience the Lorentz force

$$\boldsymbol{F} = q\left(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}\right). \tag{2.1}$$

A charged particle in the presence of an electric field will thus be accelerated parallel to the electric field. When a charged particle is in the presence of a homogeneous, static magnetic field, the velocity components parallel and perpendicular to the magnetic field evolve as

$$\frac{\mathrm{d}\boldsymbol{v}_{\parallel}}{\mathrm{d}t} = 0 \tag{2.2}$$

$$\frac{\mathrm{d}\boldsymbol{v}_{\perp}}{\mathrm{d}t} = \frac{q}{m}\left(\boldsymbol{v}_{\perp} \times \boldsymbol{B}\right), \tag{2.3}$$

meaning that the particle will circulate while travelling along the magnetic field with constant speed, resulting in a helical orbit around the magnetic field line. The angle between the velocity and the magnetic field is called the pitch angle θ_p. When a magnetic field is present in combination with an additional force, the force component perpendicular to the magnetic field will cause the particles to drift. Common causes of such drifts include electric fields and spatial inhomogeneities in the magnetic field.

2.1.2 Plasma models

The most accurate approach of modelling a plasma would be to model the dynamics of each particle in the plasma, taking into consideration all of the electromagnetic forces acting on each particle at any given time, including the forces from the electric fields of all other charged particles [9]. However, these single-particle models are computationally intractable. The numerical complexity of modelling interactions of all particle pairs is proportional to (at least) N^2, where N is the number of particles. For a typical fusion plasma $N > 10^{20}$, which would clearly be intractable to model using this method.

There are techniques for modelling plasmas based on the collective behaviour of the charged particles as a group, and the most fundamental of these is the kinetic plasma model. In the kinetic model, the distribution function $f(\boldsymbol{r}, \boldsymbol{v}, t)$ is used to describe the plasma, and as such the kinetic model is a statistical approach. The distribution function is defined such that $f(\boldsymbol{r}, \boldsymbol{v}, t)\mathrm{d}\boldsymbol{r}\mathrm{d}\boldsymbol{v}$ represents the expected number of particles within the six-dimensional volume-element $\mathrm{d}\boldsymbol{r}\mathrm{d}\boldsymbol{v}$ in position and velocity space (or, equivalently, momentum space). The evolution of the distribution function is governed by the Boltzmann equation [9]

$$\frac{\partial f}{\partial t} + \boldsymbol{v} \cdot \boldsymbol{\nabla} f + \frac{\boldsymbol{F}}{m} \cdot \boldsymbol{\nabla}_{\boldsymbol{v}} f = \left(\frac{\partial f}{\partial t}\right)_\mathrm{c}, \tag{2.4}$$

where $(\partial f/\partial t)_\mathrm{c}$ is a collision operator and the force $\boldsymbol{F}$ can be replaced by the Lorentz force in equation (2.1) if it is purely electromagnetic. In kinetic theory, macroscopic quantities of the plasma can be obtained as velocity moments of the distribution function, such as [13]

$$\text{Density, zeroth order: } n(\boldsymbol{r}, t) = \int f(\boldsymbol{r}, \boldsymbol{v}, t)\mathrm{d}\boldsymbol{v} \tag{2.5a}$$

$$\text{Mean velocity, first order: } \boldsymbol{u}(\boldsymbol{r}, t) = \frac{1}{n}\int \boldsymbol{v} f(\boldsymbol{r}, \boldsymbol{v}, t)\mathrm{d}\boldsymbol{v} \tag{2.5b}$$

$$\text{Temperature, second order: } T(\boldsymbol{r}, t) = \frac{2}{3n}\int \frac{m}{2}|\boldsymbol{v} - \boldsymbol{u}|^2 f(\boldsymbol{r}, \boldsymbol{v}, t)\mathrm{d}\boldsymbol{v}. \tag{2.5c}$$

The distribution function describing a plasma in local thermodynamic equilibrium is the Maxwell distribution [9]

$$f_\mathrm{M} \propto \left(\frac{1}{\sqrt{\pi} v_\mathrm{th}}\right)^3 \exp\left(-\frac{v^2}{v_\mathrm{th}^2}\right). \tag{2.6}$$

Here, the so called thermal speed $v_{\text{th}} = (2T/m)^{1/2}$ determines the width of the distribution, where T (measured in units of energy, typically eV) is the temperature of the particle species and m the particle mass.

The Boltzmann collision operator is rather cumbersome to work with. A more tractable form of the Boltzmann equation is obtained by using a Fokker–Planck collision operator in equation (2.4), and the resulting equation is called the Fokker–Planck equation. The Fokker–Planck collision operator is derived by assuming that all collisions occur between charged particles, so called Coulomb collisions, and that they only cause small deflections of the respective velocities, which allows expressing the effect of collisions through the following differential operator in velocity space

$$\left(\frac{\partial f}{\partial t}\right)_{\text{c}} = \sum_k \sum_l \frac{\partial}{\partial v_k}\left[-A_k f + \frac{\partial}{\partial v_l}\left(D_{kl} f\right)\right], \tag{2.7}$$

which is a sum over all phase space coordinates and where A_k is an advection vector and D_{kl} a diffusion tensor [14].

Plasmas can be further approximated by treating the plasma as a fluid. The fluid model simplifies the kinetic theory by averaging over the velocity, taking the velocity moments of the Boltzmann equation, yielding a set of so called fluid equations [9]. As such, plasmas are described by the previously mentioned macroscopic quantities in equation (2.5), obtained via taking the moments of the distribution function. Using a fluid model greatly reduces the complexity found in the kinetic model – the kinetic model describes up to six dimensions while the fluid model considers at most three – but resolution of processes in the velocity space is lost.

2.1.3 Tokamak

One of the most promising fusion reactor concepts is the tokamak, which uses magnetic fields to confine the fusion plasma. The tokamak is a torus-shaped fusion reactor, associated with a toroidal coordinate system and two characteristic geometric parameters – the plasma minor radius a and major radius R_0 – as defined in figure 2.1.

A toroidal magnetic field is applied by external magnetic field coils, as illustrated in figure 1.1, with the purpose of guiding the charged particles around the torus. However, such a magnetic field exhibits spatial gradients, and specifically this toroidal magnetic field varies as $1/R$ [15], causing vertical drifts of the particles and, without any additional measures taken, large particle losses. The (vertical) direction of the drift due to these spatial gradients will only depend on the charge of the particle, and since the charge of the particle will not change, neither will the vertical direction of the drift. Notably, the direction of the vertical drift will be the same on both the upper and lower halves of the plasma torus [13].

Introducing a smaller poloidal magnetic field allows to twist the magnetic field around the torus, which solves the particle loss caused by the spatial magnetic

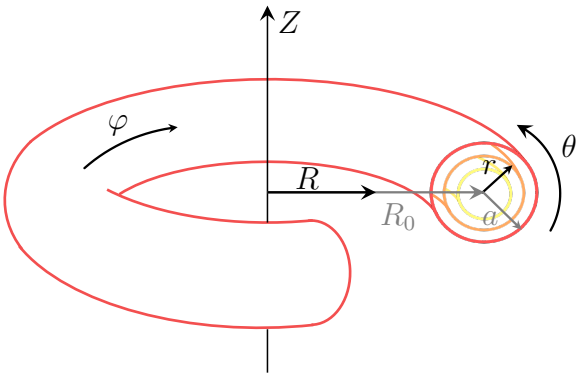

Figure 2.1: Illustration of the tokamak geometry, where the plasma minor radius a and major radius R_0 are labelled together with the radial coordinate r, poloidal coordinate θ and toroidal coordinate φ. Additionally, the toroidal coordinate forms a cylindrical coordinate system together with the radial coordinate R and vertical coordinate Z. The red, orange and yellow torus surfaces represent nested flux surfaces.

field gradients. If the particle drifts away from the plasma column in the upper half of the plasma torus, it will drift towards the plasma column on the lower half and these drifts will cancel out as the particle moves along the twisted magnetic field.

The magnetic field lines trace out closed toroidal surfaces, so called flux surfaces [15], and here we have chosen to define the radial coordinate r to be constant on such a surface. Nested flux surfaces are illustrated as the nested red, orange and yellow toroidal surfaces in figure 2.1. The degenerate flux surface at $r = 0$ is called the magnetic axis. The poloidal magnetic field is produced by driving a high toroidal plasma current through the torus. Additionally, the toroidal current is driven by a toroidal electric field. For example in ITER, the plasma current needs to be on the order of 10 MA.

In practice, a tokamak plasma does not always have a circular cross section. Typically, the cross section shape of the flux surfaces is more triangular and vertically elongated compared to a circle. The actual flux surface geometry may be parametrized by the so called shaping parameters, as illustrated in figure 2.2 together with the plasma minor radius a and wall radius b. Elongation is parametrized by $\kappa(r)$ and triangularity parametrized by $\delta(r)$ – both of these are dimensionless quantities [16]. In this thesis, the Shafranov shift $\Delta(r)$ parametrizes the horizontal displacement of the centre of a flux surface from the magnetic axis [15].

2.1.4 Tokamak disruptions

Plasma confinement is one of the most important tasks of fusion reactors, and the tokamak concept is based on confining the fusion plasma magnetically. Still, there may be off-normal events during tokamak discharges when the plasma confinement is rapidly lost; these events are called disruptions [15]. As mentioned in chapter 1,

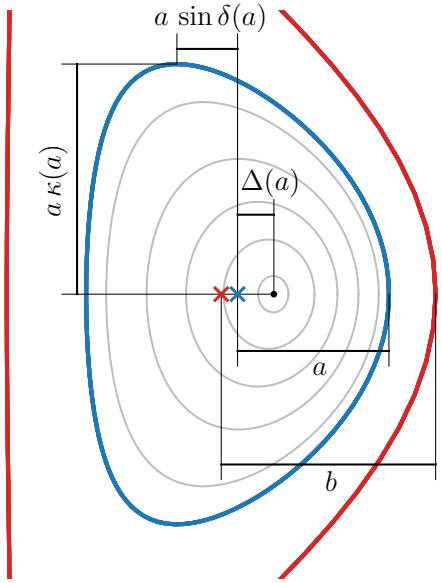

Figure 2.2: Shaping parameters for a tokamak; elongation $\kappa(r)$, triangularity $\delta(r)$ and Shafranov shift $\Delta(r)$ evaluated at the flux surface at $r = a$. The increasingly elongated and triangular circular shapes are nested flux surfaces, and the degenerate flux surface at $r = 0$ is marked with a black dot. The blue cross marks the centre of the outermost, blue flux surface, with radius a (plasma minor radius). The red cross marks the centre of the tokamak wall, with wall radius b.

disruption events constitute a major challenge facing the success of the tokamak as a viable fusion reactor because of the significant mechanical stresses and high heat loads on the machine.

A disruption is often divided into a number of characteristic phases [15], which are visualized in figure 2.3. Initially, instabilities are introduced to the plasma due to some change in the underlying conditions of the discharge. These instabilities cause the temperature of the plasma to rapidly drop by several orders of magnitude due to radiation and heat transport, during the phase called thermal quench (TQ). For ITER, the thermal quench duration is of the order of milliseconds. As the temperature decreases, the resistivity of the plasma increases such that the ohmic plasma current, carried by electrons with thermal velocities, starts to decay. The phase during which the current decays is known as the current quench (CQ), and usually lasts longer than the thermal quench – on the order of 10–100 ms for ITER. As the current decays, an electric field is induced according to Faraday's law. This increase in electric field can enable the generation of a significant current of relativistic electrons, also called runaway electrons. If a significant runaway current is generated, there is an additional phase of the disruption called the runaway plateau. During

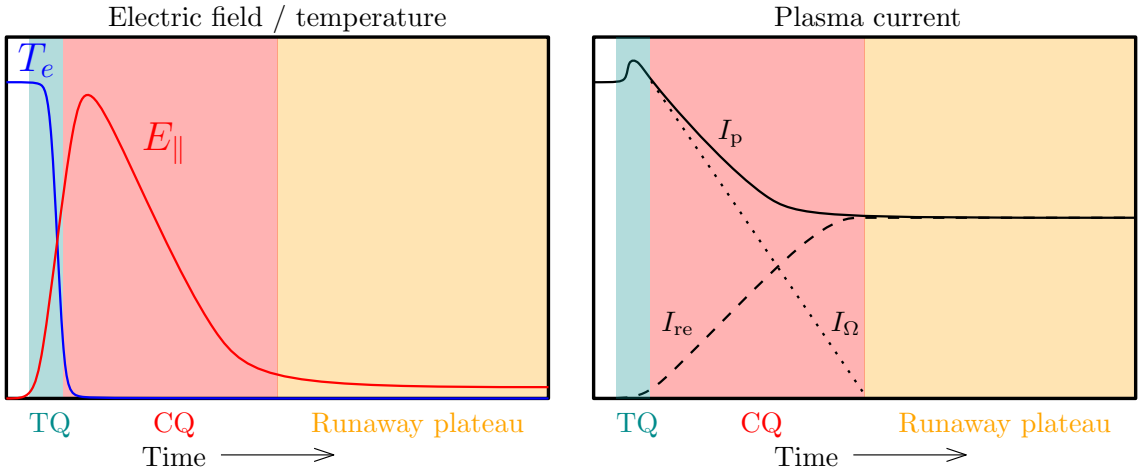

Figure 2.3: Illustration of the three important stages during a disruption: the thermal quench (TQ) during which the temperature decays, current quench (CQ) during which the ohmic current decays, which induces an electric field enabling a significant runaway current to be generated. If a significant runaway current is generated, it might reach a stable level during the stage known as the runaway plateau. Note that the total plasma current consists of the ohmic current I_Ω and runaway current I_re. (Source: M. Hoppe, Runaway electron model development and validation in tokamaks [17].)

the runaway plateau, the runaway current has reached a stable and dominant level, but may still be slowly evolving.

The duration of the current quench is important for the safety of a disruption, and it may be quantified by the CQ time

$$\tau_\mathrm{CQ} = \frac{t_{I_\Omega = 0.2 I_\mathrm{p}^\mathrm{max}} - t_{I_\Omega = 0.8 I_\mathrm{p}^\mathrm{max}}}{0.6}, \tag{2.8}$$

where I_p is the plasma current and I_Ω the ohmic current. If the current quench time is too short, the forces due to the eddy currents (loops of currents within a conducting material due to varying electromagnetic fields [18]) are large enough to damage certain modules and the first wall of the tokamak [5]. However, if the current quench lasts too long, there is sufficient time for the plasma to drift into the tokamak wall resulting in high heat loads and mechanical stresses. For ITER, the CQ must be longer than 50 ms but shorter than 150 ms according to Ref. [5].

Another important factor related to the safety of a disruption is the heat that has been transported out of the plasma through radial transport, and thus into the surrounding machine. One metric to quantify this is the fraction of the initial thermal energy content in the plasma $W_\mathrm{th}(t=0)$ that has been lost from the plasma through radial heat transport, also called the conducted heat load

$$\eta_\mathrm{cond} = \frac{\int_0^{\tau_\mathrm{d}} P(t)\mathrm{d}t}{W_\mathrm{th}(t=0)}, \tag{2.9}$$

where $P(t)$ is the radially transported power at time t and τ_d is the duration of the disruption. In ITER, heat loads larger than 10 % are expected to enable melting of important tokamak components [5].

Apart from the total radial transport of heat, there can be heat related damages on the machine if a large amount of energy is deposited on a small area during a short period of time [8]. These are common features of the impact of runaway current on the tokamak wall, as the runaway current may contain a large fraction of the stored magnetic energy in the form of highly energetic electrons. The risk of significant runaway currents increase exponentially with maximum plasma current, and thus pose a substantial threat to tokamak projects such as ITER. Thus, runaway electron avoidance constitutes a crucial component of disruption mitigation research, especially for large machines such as ITER.

There are several proposed disruption mitigation schemes, and massive material injection (MMI) is the most studied one. As the name suggests, massive material injection consists of injecting large amounts of material into the plasma in order to cool it down by increasing radiation, and thus avoiding large heat loads on the plasma facing components [8]. Furthermore, controlling the plasma composition allows a certain degree of control over the temperature and the plasma current evolution. MMI can be achieved in a number of ways, and two examples are massive gas injection and injection of cryogenic pellets [5].

Common candidates for injection materials are deuterium and the noble gases. As deuterium only contains one electron per atom and gets fully ionized at rather low temperature it will not be subject to the same large radiation losses of heavier atoms, and thus the injected deuterium will mostly function to dilute the plasma without triggering instabilities associated with runaway electron generation [19]. Injection of noble gases, such as neon, will initialize heat loss through radiation and transport. More specifically, a transport event will be triggered at a certain plasma radius after which radial heat and particle transport will be strongly elevated [20]. One notable feature of deuterium pellet injection is that a pure deuterium pellet cloud is expected to drift significantly, but this can be suppressed by introducing a small amount of neon [21].

An additional possible runaway electron mitigation scheme is to perturb the magnetic field by external coils [22–25]. The idea is that a perturbed magnetic field will increase radial transport of runaway electrons. Magnetic perturbations also occur naturally during the TQ phase when the flux surfaces are broken and there are large fluctuations of the magnetic field [17].

2.1.5 Runaway electrons

The motion of an electron inside a tokamak will be affected by the Lorentz force due to the electric field and forces due to Coulomb collisions. Thus, the equation of

motion parallel to the magnetic field can be approximately written as

$$\frac{\mathrm{d}v_\parallel}{\mathrm{d}t} = \frac{e}{m_\mathrm{e}}\left|E_\parallel\right| - \nu_\mathrm{c} v_\parallel, \tag{2.10}$$

where the collision frequency $\nu_\mathrm{c} \propto 1/v^3$ for superthermal velocities. Thus, for high enough velocities, the Lorentz force will become larger than the collisional friction forces, resulting in a continuous acceleration. Electrons experiencing this kind of unhindered acceleration due to the electric field are called runaway electrons and they are accelerated to relativistic speeds. The velocity at which an electron becomes a runaway in a plasma is called the critical velocity. A specific electric field value, the so called critical electric field, is of particular importance for the dynamics of runaway electrons. It is derived by solving equation (2.10) under the condition that the friction force is minimal, which occurs at $v_\parallel \approx c$ [26], yielding

$$E_\mathrm{c} = \frac{e^3 n_\mathrm{e} \ln\Lambda}{4\pi\epsilon_0^2 m_\mathrm{e}^2 c^2}, \tag{2.11}$$

where n_e is the electron density, and $\ln\Lambda$ is the Coulomb logarithm quantifying the maximum impact parameter of Coulomb collisions [9]. This is the minimum electric field at which runaway electrons can be generated – for lower electric fields the collisional forces will always be stronger than the Lorentz force.

In a tokamak, there are several physical mechanisms via which runaway electrons can be generated. The generation process of electrons collisionally diffusing above the critical velocity causing them to run away is called the Dreicer process. The Dreicer generation mechanism is thus highly dependent on the electric field strength. When electrons just above the critical velocity in phase-space are accelerated, collisions tend to restore the Maxwellian distribution and fill in the resulting gap in the distribution function.

The hot-tail mechanism occurs when the plasma cools down sufficiently rapidly such that electrons of high enough velocity, which are not as collisional, do not have time to slow down [27]. In such a scenario, the energy distribution of the electrons, which initially would be approximated by a Maxwellian, develops a "tail" into the higher energy regions. If the electric field then is rapidly increased, which is what happens during the CQ, following the TQ, of a tokamak disruption, the electrons of this so called hot-tail of the distribution will be turned into runaways.

The aforementioned generation mechanisms are primary ones, meaning that they are not dependent on runaway electrons already existing in the plasma. Some generation mechanisms are dependent on the existence of runaway electrons, so called secondary generation mechanisms, and one such is avalanche generation. Avalanche generation occurs when there is a large momentum exchange during a Coulomb collision between a runaway electron and another electron, resulting in both electrons having high enough velocity to be accelerated by the electric field and become runaways [28]. Since this process is proportional to the number of already existing runaway electrons, it leads to an exponential increase of the runaway population.

The opposite may also be true, namely that both electrons have lower velocities than the critical velocity after the collision, and this runaway annihilation mechanism will here be referred to as reverse avalanche.

The last two runaway generation mechanisms of relevance to tokamak disruptions are tritium and Compton generation. Tritium is an unstable isotope which decays through β-decay into a helium-3 nucleus, an electron and an electron antineutrino [29]. If the released electron has a speed larger than the critical speed, it will accelerate and become a runaway. Compton generation occurs when the tokamak walls become activated by the neutrons produced during the fusion reaction, emitting γ-rays which can Compton scatter on the electrons, possibly increasing their energy sufficiently for them to become runaways [29].

This concludes the physics part of the theoretical background necessary for this project. However, the method employed to investigate MMI during tokamak disruptions is heavily dependent on Bayesian optimization. Bayesian optimization will now be presented, since a basic knowledge of how it works is beneficial for understanding and analyzing the results of this thesis.

2.2 Bayesian optimization

Black box optimization refers to optimization methods which do not depend on the structure of the function to be optimized. Instead, they map inputs to outputs and bases their progression on the values of these input-output pairs. Bayesian optimization is a common black box optimization method, which is based on Bayes' theorem.

Bayes' theorem regards conditional probabilities and describes the probability of an event based on prior knowledge related to the event. Let A and B denote two events, and $p(A|B)$ is the probability that A transpires given that B is true. Bayes' theorem states that

$$p(A|B) = \frac{p(A)p(B|A)}{p(B)}, \tag{2.12}$$

where $p(B|A)$ is the probability that B transpires given that A is true [30]. Bayes' theorem is a way of determining the probability of an event based on prior knowledge.

Fundamentally, Bayesian optimization uses Bayesian inference, and Bayes' theorem is applied on probability distributions: "Bayesian inference is a framework for inferring uncertain features of a system of interest from observations grounded in the laws of probability" [10]. This is achieved by assuming that all unknowns of the system to be inferred are random variables. In the context of Bayesian optimization, on a small scale this means inferring the value $\phi = f(x)$ of the *objective* or *cost function* f to be optimized at a certain point x. Bayesian optimization extends this to a large scale, by inferring the objective function on a finite domain.

At small scales, Bayes' theorem for probability distributions of the objective function value at a certain point x is

$$p(\phi|x,y) = \frac{p(\phi|x)p(y|x,\phi)}{p(y|x)}, \tag{2.13}$$

where the variable y is the measurement received when observing the objective function f at x [10]. Here, $p(\phi|x)$ is called the prior distribution (or just prior), and describes our prior knowledge of how plausible different function values ϕ are, before observing any data. The likelihood function $p(y|x,\phi)$ (called likelihood) describes the likelihood of getting the measurement y as a function of $\phi = f(x)$ when making an observation of the objective function at x. These two probability distributions determine the posterior distribution function (or just posterior) $p(\phi|x,y)$, which describes the probability distribution over the objective function value ϕ given x and the measurement y observed at x. The denominator of (2.13) is called evidence and ensures the normalization of the posterior,

$$p(y|x) = \int p(\phi|x)p(y|x,\phi)\mathrm{d}\phi. \tag{2.14}$$

This can be extended to find the posterior distribution function of the entire function $f(x)$

$$p(\phi(x)|\mathcal{D}) = \frac{p(\phi(x)|x)p(\mathcal{D}|x,\phi(x))}{p(\mathcal{D})}, \tag{2.15}$$

where $\phi(x)$ represents the actual function values in a subset to the domain of f. Here, the knowledge we have about the objective function is collected in the data set $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{N}$, which is compromised of N observation pairs – measurements y_i observed at x_i.

In order to use (2.15) to make viable predictions for the objective function, an assumption of the probability distribution of $\phi(x)$ is needed. For this, Bayesian optimization uses stochastic processes, which are infinite collections of random variables [10]. An especially common stochastic process used for Bayesian optimization is the Gaussian process (GP), in which the random variables are distributed according to multivariate Gaussian distributions. A Gaussian process on the objective function $f(x)$ is specified by a mean function $\mu(x) = \mathbb{E}[\phi|x]$, which determines the expected function value $\phi = f(x)$ at any x, and a covariance function, or kernel, $K(x,x')$, which measures the correlation between ϕ for points x and x' [10].

The covariance function $K(x,x')$ can be defined in different ways. Thus, the kernel needs to be specified before a GP can be used. One widely used kernel is the Matérn kernel, with which the desired degree of smoothness for the objective function can be controlled with a parameter, ν [10]. The smoothness parameter determines how smooth the functions drawn from the GP should be, which in turn impacts the smoothness of the mean function determining the expected values of the objective function. Some specific values of ν correspond to specific features of the function, for example $\nu = 3/2$ and $\nu = 5/2$ correspond to being once or twice differentiable functions, respectively.

In Bayesian optimization, the mean and the covariance functions set in the GP are used together with an acquisition function to determine which new point in the domain would yield most information regarding the position of the optimum of the objective function. There are several kinds of acquisition functions, and a commonly used one is the expected improvement acquisition function (EI). Using EI, the new points are selected based on the maximum of the expected improvement

$$\boldsymbol{x}_{\text{new}} = \arg\max_{\boldsymbol{x}} \left[\mathbb{E}[I(\boldsymbol{x}, y^*)]\right] = \arg\max_{\boldsymbol{x}} \left[\int I(\boldsymbol{x}, \phi, y^*) p(\phi|\boldsymbol{x}, \mathcal{D}) \mathrm{d}\phi\right], \tag{2.16}$$

where

$$I(\boldsymbol{x}, \phi, y^*) = \begin{cases} \max(\phi - y^*, 0) & \text{(for maximization)} \\ \max(y^* - \phi, 0) & \text{(for minimization)} \end{cases} \tag{2.17}$$

is a measure of the improvement [31]. Here, y^* is the optimal objective function value encountered so far.

To summarize, the Bayesian optimization method starts with a GP prior, determined by the initial knowledge and assumptions we have made about the objective function. By using the data set $\mathcal{D}$ of previous observations, a likelihood distribution function is obtained which, together with the GP prior, yields a GP posterior. This GP posterior is then used by the acquisition function to determine the most promising new point $\boldsymbol{x}$ to observe in order to improve the current optimum, yielding a new data point $(\boldsymbol{x}_{\text{new}}, y_{\text{new}})$ to $\mathcal{D}$. Finally, the Bayesian optimization relies on an iterative process, where each new observation, based on the acquisition function's proposition, can be added to $\mathcal{D}$ and the method can be repeated with the new $\mathcal{D}$ to obtain the next, best point to observe proposed by the acquisition function.

As with many stochastic optimization algorithms, the concepts of exploration and exploitation are relevant [10]. An optimizer which favours exploitation will prefer to look for the optimum close to the best candidate for an optimum found so far. On the other hand, if exploration is favoured, the optimizer will prefer to look for the optimum in regions far from other observations where the uncertainties are large. For Bayesian optimization, the exploration–exploitation trade-off is modelled in the acquisition function.

3

Simulations

For the Bayesian optimizer, the objective function to be optimized can be viewed as a black box, where its content is irrelevant to the optimization progression – it only connects the inputs given to the black box to the output received. However, the contents of the black box is vital for interpreting and understanding the results of the optimization. The major ingredient of our black box representing MMI in tokamak disruptions is the simulation performed. Since the simulations play such a crucial role for this thesis, this chapter will describe DREAM, which is the numerical tool used to simulate the disruption scenarios, as well as how the disruption scenarios have been modelled.

3.1 Numerical tool DREAM

The simulation code DREAM (Disruption Runaway Electron Analysis Model) used in this thesis was specifically developed to study runaway electrons during tokamak disruptions by members of the Plasma Theory group at the Chalmers University of Technology [32]. It self-consistently evolves the plasma parameters most important to tokamak disruptions. The background parameters, such as the temperature and electric field, are simulated using fluid models, while electrons can be treated with several different options of plasma models – within the DREAM community referred to as fluid, isotropic, superthermal or fully kinetic – ranging from low to high kinetic treatment and computational cost.

Since DREAM was developed to study runaway electrons, the electron evolution of the simulated disruption is important. In the context of disruptions, the electrons can be divided into three distinct populations – cold (thermal) electrons, hot electrons and runaway electrons – based on momentum as shown in figure 3.1. There are several plasma models to choose from in DREAM, corresponding to different degrees of kinetic treatment for the three electron populations. Kinetic modelling means that the distribution of electrons is determined in momentum and pitch angle space. Recall that the fluid model does not determine the distribution of electrons in momentum space, as described in section 2.1.2. In the fully kinetic model, both the cold and hot electrons are treated kinetically, as illustrated in figure 3.1a and evolved using the same distribution function. In the isotropic and superthermal

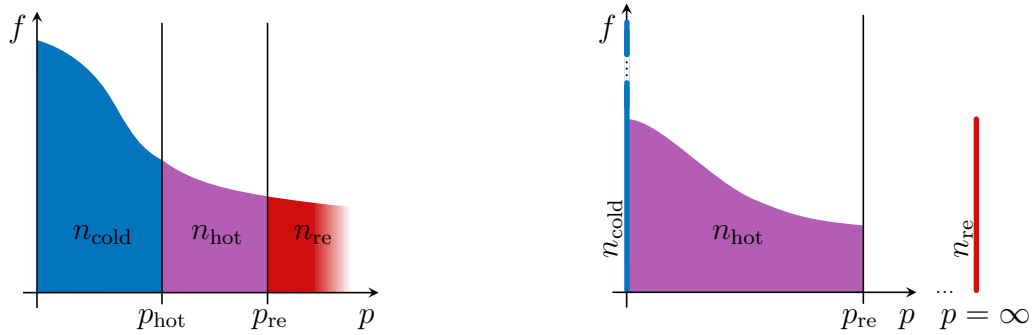

(a) Electron populations in the fully kinetic model with runaway electrons resolution – all electrons are treated kinetically.

(b) Electron populations in the isotropic and superthermal models without runaway electron resolution. The cold and runaway electron populations are treated as fluids and thus not treated kinetically.

Figure 3.1: Illustration of the separation of electron populations in the momentum space in the fully kinetic, isotropic and superthermal models. The fluid model is not as computationally expensive as the isotropic/superthermal models, by not explicitly describing the hot electron population.

models, only the hot electrons are treated kinetically, as shown in figure 3.1b. In the fluid model, only the spatial density of electrons is evolved, and the details about the electron momentum is generally neglected. Additionally, the runaway electrons can be treated either kinetically or as a fluid, regardless of how the cold and hot electrons are treated.

The cold, or thermal, electron population is assumed to be distributed according to a Maxwell distribution. In DREAM, it is represented by the cold electron temperature T_{cold}, the ohmic current density j_Ω, and the cold electron density which is evolved to preserve quasi-neutrality.

In all models of DREAM except the fluid model, the hot electron population is treated explicitly during the simulations. The hot electron population consists of electrons with momentum above p_{hot}, which is the momentum value separating the hot from the cold population. In DREAM, the hot electrons are represented by the distribution function f_{hot}, from which the electron density n_{hot} and current density j_{hot} can be calculated. The distribution function is evolved according to the Fokker–Planck equation, as described in Ref. [32]. In the isotropic and superthermal models of DREAM, the cold electron population is treated as a fluid, and there is a flux of particles from the hot electron population to the cold electron population through a particle sink at $p = 0$.

Finally, there is the runaway electron population represented by the runaway electron density n_{re}, current density j_{re} and optionally the distribution function f_{re}, and is separated from the hot electron population at the critical momentum value p_{re}.

When the runaway electrons are not treated kinetically, they are assumed to move along the magnetic field with the speed of light.

One major difference between the superthermal model and the fully kinetic model is that the plasma can practically have two temperatures in the former, one for the cold population and one for the hot [32]. In the fully kinetic model, the cold electrons are treated kinetically as described above, and they are then described by the hot distribution function f_{hot} alongside the hot electrons. The initial distribution function is therefore described by a modified Maxwell distribution – the Maxwell–Jüttner distribution which takes relativistic effects into account – at temperature T_{cold}. Since the cold and hot electrons are described by the same distribution function with one temperature, the fully kinetic model is similar to the fluid model where all non-runaway electrons are part of the cold electron population with one temperature. In the superthermal model, in which only the hot electron population is treated kinetically, the electron–electron collision frequencies in the collision operator are taken at the $T_{\text{cold}} \to 0$ limit. The cold electron population will still have the temperature T_{cold}, but since f_{hot} no longer describes the cold electron population, the hot electron temperature T_{hot} derived from taking the second moment of f_{hot} would be different from T_{cold}, although T_{hot} is not explicitly calculated as an output in DREAM. This could be particularly suitable for a MMI scenario, where cold electrons are generated from ionization of the cold injected material, while the bulk of the initial plasma is part of the hot electron population with a much higher temperature.

In DREAM, the superthermal and isotropic models are very similar, with the isotropic model being a further simplification of the superthermal model by analytically averaging the Fokker–Planck equation over pitch angle [32]. The analytical derivation is based on the assumption that pitch angle scattering dominates the electron dynamics, meaning that the fastest process is the collisional isotropization of the distribution function (hence the name "isotropic"). The lowest order deviation from isotropy (that gives the hot electron current) is analytically calculated. This makes the isotropic model less complex and computationally expensive, since the resolution dimensionality is reduced.

In the fluid simulations, the runaway electron density is evolved according to

$$\begin{aligned}\frac{\mathrm{d}n_{\text{re}}}{\mathrm{d}t} = \Gamma_{\text{Ava}} n_{\text{re}} + \gamma_{\text{Dreicer}} + \gamma_{\text{hot-tail}} + \gamma_{\text{tritium}} + \gamma_{\text{Compton}} \\ + \frac{1}{V'}\frac{\partial}{\partial r}\left[V'\left(An_{\text{re}} + D\frac{\partial n_{\text{re}}}{\partial r}\right)\right],\end{aligned} \tag{3.1}$$

where Γ_{Ava} is the avalanche growth rate, and the rest, except the final term, are primary runaway rate terms due to the respective mechanisms [32]. The final term of the equation is the advective and diffusive radial transport of the runaway electrons, where A and B can specified, and $V' = V'(r)$ is the spatial Jacobian at radius r. Different models can be used for each of the runaway mechanisms, and those used in this thesis will be described in section 3.2.

In the kinetic models (namely the isotropic and superthermal models used in this thesis) the runaway density is evolved as

$$\frac{\mathrm{d}n_{\mathrm{re}}}{\mathrm{d}t} = \Phi_{\mathrm{hot}}^{(p)} + \Gamma_{\mathrm{Ava}} n_{\mathrm{re}} + \gamma_{\mathrm{tritium}} + \gamma_{\mathrm{Compton}} + \frac{1}{V'}\frac{\partial}{\partial r}\left[V'\left(An_{\mathrm{re}} + D\frac{\partial n_{\mathrm{re}}}{\partial r}\right)\right], \quad (3.2)$$

and the Dreicer and hot-tail runaway generation processes are automatically included as a consequence of the detailed modelling of the plasma collision physics and as such are part of the momentum space flux $\Phi_{\mathrm{hot}}^{(p)}$. In the fluid model, these runaway rates must be added explicitly as there is no flux $\Phi_{\mathrm{hot}}^{(p)}$ of electrons as shown above.

Besides the electrons, DREAM also evolves the background plasma parameters, such as the temperature, electric field and ions, in space and time. The electron temperature is modelled through the evolution of the cold electron energy $W_{\mathrm{cold}} = 3n_{\mathrm{cold}}T_{\mathrm{cold}}/2$. In turn, the cold electron energy depends on ohmic heating due to the energy transferred from the electric field, collisional heat transfer from the other electron populations and ions, energy losses due to inelastic atomic processes (radiation) and radial transport of heat through diffusion and advection. The distribution of ion charge states is evolved through ionization and recombination, which are based on the rate coefficients of the Open-ADAS database [33] for the thermal electron population. The effects of opacity for the hydrogen isotopes to Lyman radiation can be taken into account, relevant at low temperatures, and for this case the the rate coefficients from the AMJUEL database [34] are used instead. Additionally, ionization by the non-thermal electron populations – so called fast electron impact ionization – can be included [32]. The fast electron impact ionization is modelled by a kinetic ionization rate based on the ionization cross-section in Ref. [35]. The current density $j_{\mathrm{tot}} = j_{\Omega} + j_{\mathrm{hot}} + j_{\mathrm{re}}$ and electric field are evolved via Ampère's and Faraday's laws. If the runaway electrons are not treated kinetically, they are all assumed to travel at the speed of light parallel to magnetic field lines so that their contribution to the current density is $j_{\mathrm{re}} = ecn_{\mathrm{re}}$.

Finally, there are some technical aspects to review. The tokamak-specific parameters, such as the magnetic field, time-varying shaping parameters as well as major, minor and wall radii need to be set. For time evolution, there is the option of either having a fixed time step length or an adaptive one based on the ionization time scale [36]. The equation system is evolved using the first-order implicit time-stepping. Specifically, due to the non-linear nature of the equation system, a Newton solver is used to obtain the solution for each time step. Alternatively, the system may be linearized in time, yielding a so-called linearly implicit time-stepping method. In this thesis, the iterative Newton's method will be used for converging to this solution, and adjustable tolerances are used to determine whether or not the solver has converged.

3.2 Disruption model

An ITER-like tokamak setup has been used for the simulations, because of its importance towards reaching the goals of the global fusion program. Furthermore, the risk of runaway electrons will be high in ITER due to its large size and high current. To achieve an ITER-like setup, the major radius is set to $R_0 = 6\,\mathrm{m}$, the minor radius is set to $a = 2\,\mathrm{m}$ and the toroidal magnetic field on axis is set to $B_0 = 5.3\,\mathrm{T}$. Additionally, the wall radius, or effective radius of the first toroidally closed conducting structure, is set to $b = 2.833\,\mathrm{m}$ consistent with the value calculated in Ref. [37] by matching the poloidal magnetic energy to that in simulated ITER discharges.

It is worth noting that the toroidicity of the plasma has been taken into account along with realistic shaping, rather than using the cylindrical limit $a/R \to 0$. The radially varying shaping parameters used are the same as the ITER-like, and physically consistent, shaping parameters used in Ref. [38].

During activated operation, the plasma fuel will contain tritium, which is a radioactive isotope. In our model, the plasma fuel is assumed to consist solely of deuterium for the modelling of non-activated operation, while the activated model has 50 % deuterium and 50 % tritium, with all particles being fully ionized at the start of the simulation. In both cases, the initial electron density is $10^{20}\,\mathrm{m}^{-3}$. The hydrogen isotope populations in the plasma are assumed to not be fully transparent to the emitted line radiation. Instead, bound-bound electron transitions involving the hydrogen isotope ground state are neglected to account for the hydrogen isotopes being opaque to Lyman line radiation [39].

As for initial conditions of other parameters evolved in DREAM, the initial plasma current $I_\mathrm{p} = 15\,\mathrm{MA}$ and initial plasma temperature on axis $T(r = 0) = 20\,\mathrm{keV}$, which are based on the plasma current and temperature that a high fusion performance plasma discharge will have in ITER. In our simulations, the initial electron and ion temperatures are equal, but they are evolved as separate parameters in DREAM and as such can be different during the simulations. Furthermore, the initial current density profile is assumed to be

$$\hat{j}(r) = \left[1 - \left(\frac{r}{a}\right)^2\right]^{0.41} \tag{3.3}$$

and the initial temperature profile

$$\hat{T}(r) = 1 - 0.99\left(\frac{r}{a}\right)^2, \tag{3.4}$$

which are based on Ref. [40] but takes into consideration that $\hat{T}(r = a) \neq 0$. The current and the temperature profiles are shown in figure 3.2a and 3.2b, respectively. It is worth noting that current density relaxation expected at the onset of an ITER disruption is neglected during the TQ in our disruption model; possible effects of such relaxation are discussed in Ref. [38].

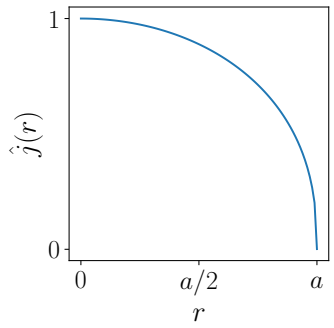

(a) Illustration of the current density profile $\hat{j}(r)$.

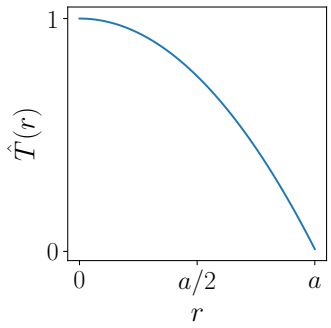

(b) Illustration of the temperature profile $\hat{T}(r)$.

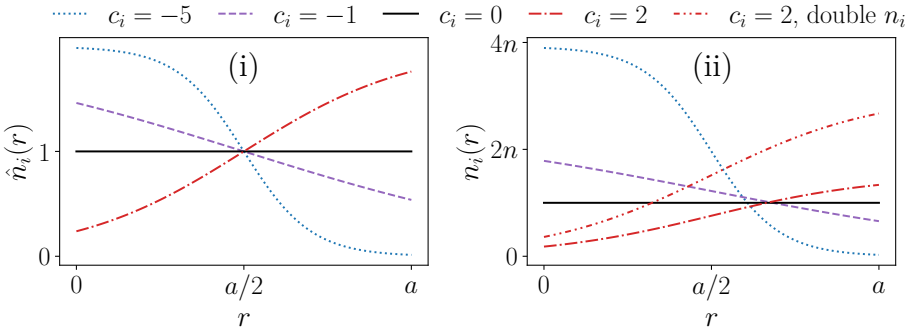

(c) Illustration of the (i) density profile $\hat{n}_i(r)$ and (ii) injected density distribution $n_i(r)$ for different values of the inputs n_i and c_i, and $i \in \{\mathrm{D}, \mathrm{Ne}\}$. For all curves of (ii) except the red dash-double-dotted, the injected density parameter $n_i = n$, while for this dash-double-dotted curve $n_i = 2n$. Note that the circular cross section of the plasma favours higher peak densities for negative c_i, as shown for $c_i = -1$ (dashed purple line) compared to $c_i = 2$ (dash-dotted red line).

Figure 3.2: Illustrations of initial profiles of important plasma parameters.

The runaway dynamics is governed by equation (3.1) in the fluid model and (3.2) in the kinetic model, where all possible runaway generation mechanisms are listed, but all might not be present in a given disruption scenario. For the non-activated disruption model, Dreicer, avalanche and hot-tail runaway generation processes are included. Additionally, Compton and tritium generation rates are included for the activated model. The avalanche generation rate is of the form

$$\Gamma_{\mathrm{Ava}} \propto \frac{n_{\mathrm{tot}}}{n_{\mathrm{cold}}} \left(E_{\parallel} - E_{\mathrm{c}}^{\mathrm{eff}} \right), \tag{3.5}$$

where $E_{\mathrm{c}}^{\mathrm{eff}}$ is the effective critical electric field and n_{tot} is the total electron density, including bound electrons, as derived in Ref. [41]. In the fluid model, the Dreicer generation rate is determined with a neural network established in Ref. [42] and the hot-tail generation is evaluated based on the critical momentum found using an angle-averaged kinetic equation in the limit of high charge numbers, as in Ref. [43].

Furthermore, the runaway electrons can be lost from the plasma through transport due to magnetic perturbations, as described below.

In our simulations, MMI is modelled as a material already distributed in the plasma. The injected material is assumed to consist of cold atom populations distributed instantaneously, either spatially uniformly or with the radial profile

$$\hat{n}_i(r; c_i) = 1 + \tanh\left[c_i\left(\frac{r}{a} - 0.5\right)\right], \tag{3.6}$$

with the index $i \in \{\mathrm{D}, \mathrm{Ne}\}$ referring to the atom species. Note that the uniform density distribution corresponds to $c_i = 0$. The radial density distribution as a function of the radius r becomes

$$n_i(r; n_i, c_i) = n_i \hat{n}_i(r; c_i) \frac{V_\mathrm{p}}{\int_{V_\mathrm{p}} \hat{n}_i(r; c_i) \mathrm{d}V}, \tag{3.7}$$

where V_p is the plasma volume, and n_i without argument is a scalar parameter that is used in the optimization. Both the density profile and density distribution are visualized in figure 3.2c.

The magnetic perturbations in our disruption model are only employed during the TQ, and they are chosen to be constant both radially and in time. Magnetic perturbations affect the runaway electron and heat transport, and consequently the plasma evolution. This thesis uses the Rechester–Rosenbluth model [44] which assumes that the magnetic field is fully ergodic and models the transport as a diffusion process with the diffusion coefficient

$$D = \pi R_0 \left|v_\parallel\right| \left(\frac{\delta B}{B}\right)^2, \tag{3.8}$$

where $\delta B/B$ is the magnetic perturbation normalized to the unperturbed magnetic field. Here, $v_\parallel$ is the velocity of the electron along the magnetic field line, which is, to good approximation, the speed of light c for runaway electrons. When calculating the electron heat diffusivity for the cold electron population using the same model, the parallel velocity is integrated over the bulk velocity distribution such that $v_\parallel \sim v_\mathrm{th}$ [32]. For the hot population, included in the isotropic and superthermal models, the parallel velocity is based on the momentum and pitch angle. The parallel correlation length of the magnetic perturbations should appear in place of the R_0 factor in (3.8), but since it is comparable to the R_0, R_0 has been used instead. Any order unity factors that may appear in reality due to the parallel correlation length are instead absorbed into $(\delta B/B)^2$.

Finally, there are some minor specifications in our disruption model worth mentioning. In the kinetic simulations, the hot electron current will be counted towards the runaway current, to allow fair comparisons between the isotropic and fluid models – the total plasma current is thus divided into the ohmic current and the non-thermal current in both models. The collision frequencies, which are only relevant in the kinetic models, in the isotropic and superthermal models use the $T_\mathrm{cold} \to 0$ limit of

collision frequencies according to equation (34) of Ref. [32]. Partial screening of the nucleus is accounted for during collisions [45]. The Coulomb logarithm is dependent on the electron energy, as defined in equation (18) in Ref. [32]. Bremsstrahlung losses, caused by Coulomb collisions, are taken into account using a mean-force model [32]. The momentum dependent components of the avalanche, Compton and tritium runaway rates are set to not take trapping effects into consideration, as inspired by Ref. [46]. The conductivity is evaluated from a numerical fit of the neoclassical conductivity to Fokker–Planck simulations, as performed in Ref. [47]. Fast electron impact ionization is included in the kinetic models [32].

4

Optimization

In this chapter, the optimization algorithm and its components will be presented. Firstly, the design and motivations of the cost function will be discussed. Thereafter the algorithms and results of the optimizations will be presented. The results will be divided into two sections based on the dimensionality of the optimizations run. The first optimizations are two-dimensional – the input space consists of the densities of the injected materials, which are distributed uniformly. For the four-dimensional optimization, the radial variation of the densities adds two dimensions.

4.1 Cost function

The cost function of an optimization problem is crucial for the optimization to yield meaningful results. A good cost function should reliably quantify the development inside the black box while following good optimization practices. In this section we describe the design of the cost function we have used.

4.1.1 Components of the cost function

The aim of our cost function is to quantify the risks associated with tokamak disruptions. Here it will be constructed from four disruption parameters, with associated safe operational limits, which have been deemed relevant.

As described in Ref. [5], the runaway current should be small for a safe disruption. However, the runaway current output received from DREAM is a time vector while a representative scalar value is required for the cost function. Here we consider two plausible choices for representing the runaway current contribution to the cost – the maximum runaway current $\max_t I_{\mathrm{re}}(t)$ and the runaway current at the time when it makes up 95 % of the plasma current $I_{\mathrm{re}}(t_{I_{\mathrm{re}}=0.95I_{\mathrm{p}}})$, which we will call the 95-percent runaway current. If a runaway current plateau is reached during the disruption, the runaway current will in most cases be slowly but monotonically increasing with time. In this case the maximum runaway current will be dependent on the simulation length, which is impractical.

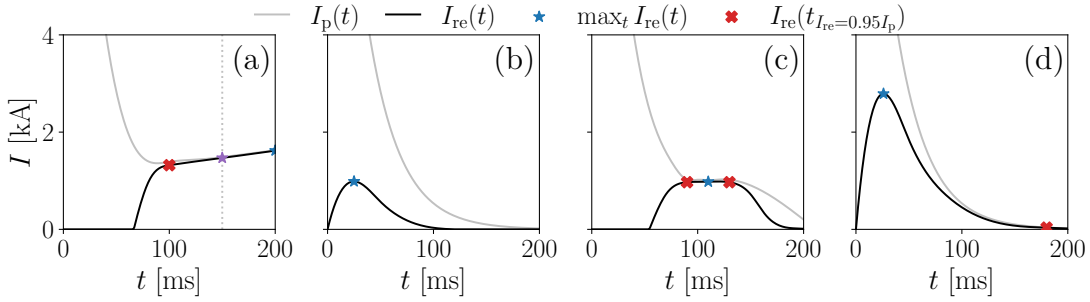

Figure 4.1: Illustration of possible issues with the two representative runaway current definitions. (a) A runaway plateau is reached during the disruption, such that the runaway current is slowly increasing and the maximum runaway current depends on the simulation length. If the simulation is stopped at 200 ms (150 ms) the blue (purple) star marks the maximum runaway current. (b) The runaway current never reaches 95 % of the total plasma current. (c) The runaway current constitutes 95 % of the total plasma current twice, marked by red crosses. (d) The runaway current peaks significantly early during the simulation (blue star), but reaches 95 % of the total plasma current when both currents are negligible (red cross).

Using the 95-percent runaway current would solve this problem, but it gets complicated when $I_{re}(t) < 0.95I_p(t)$ throughout the simulation or when $I_{re}(t) = 0.95I_p(t)$ for several t. Furthermore, for a hypothetical case where the runaway current near the end of the simulation fulfills $I_{re} = 0.95I_p$ while being small, but peaks early during the simulation with an order $\sim 1\,\mathrm{MA}$ without reaching 95 % of the plasma current, $I_{re}(t_{I_{re}=0.95I_p})$ would not be a representative value.

To account for the advantages and disadvantages of both options (illustrated in Fig. 4.1), we chose to use the first occurrence of $I_{re}(t_{I_{re}=0.95I_p})$ unless $\max_t I_{re}(t)$ occurred before $I_{re}(t_{I_{re}=0.95I_p})$ or if $I_{re}(t) < 0.95I_p(t)$ throughout the simulation, and will call it the representative runaway current I_{re}^{repr}.

Additionally, the ohmic current at the end of the simulation has the potential of being converted into runaway current, and signifies an incomplete TQ, and consequently it should also be minimized to reduce the uncertainty with the simulation results. Based on [48], the upper safe operational limit of the runaway current is 150 kA, and due to the correlation between the ohmic and runaway current the two currents will be weighted equally.

The CQ time in the cost function was defined as

$$\tau_{CQ} = \begin{cases} \dfrac{t_{I_\Omega=0.2I_p^{max}} - t_{I_\Omega=0.8I_p^{max}}}{0.6} & \text{if } t_{I_\Omega=0.2I_p^{max}} \in [0, t_{final}] \\ \dfrac{t_{final} - t_{I_\Omega=0.8I_p^{max}}}{0.8 - I_\Omega^{final}/I_p^{max}} & \text{otherwise,} \end{cases} \tag{4.1}$$

and as stated in Ref. [5], its safe operational interval is $[50, 150]$ ms. The last contribution to the cost function is the conducted heat load, defined by (2.9), which

should be minimized and for a safe disruption it should be below 10 %, according to Ref. [5].

4.1.2 Cost function composition

As previously stated, the cost function will consist of four components with associated safe operational limits, which when simultaneously satisfied, correspond to a safe operational region in the parameter space. Three of the components should be minimized while one should be within a certain interval. In the interest of maximizing the amount of information gained from the cost function values, we intend the function to yield function values below one inside the safe operational space and above one outside of this space. In practice, this will approximately be achieved by normalizing the components to their safe operational limits and combining them using the Euclidean norm.

To achieve this boundary value of one between safe and non-safe cases, each component should yield exactly unity at the limit of its safe interval when normalized. Thus, both currents were normalized by 150 kA and the conducted heat by 10 %. In order to get the same characteristics for the CQ time, it was shifted and normalized according to $|\tau_{\mathrm{CQ}} - 100\,\mathrm{ms}|/50\,\mathrm{ms}$. In order for the prompt, namely for safety to be equivalent to the cost function $\mathcal{L} < 1$, to be precisely realized, we would need to use the maximum norm, i.e.

$$\mathcal{L} = \max\left(\frac{I_{\mathrm{re}}^{\mathrm{repr}}}{150\,\mathrm{kA}}, \frac{I_{\Omega}^{\mathrm{final}}}{150\,\mathrm{kA}}, \frac{\eta_{\mathrm{cond}}}{10\,\%}, \frac{|\tau_{\mathrm{CQ}} - 100\,\mathrm{ms}|}{50\,\mathrm{ms}}\right). \tag{4.2}$$

However, this would result in the cost function not being once differentiable (or smooth) and only one component would contribute with information about the disruption scenario to the optimizer in any given point.

A solution to this problem could be combining the normalized components by simply adding them together, i.e.

$$\mathcal{L} = c_{I_{\mathrm{re}}^{\mathrm{repr}}} \frac{I_{\mathrm{re}}^{\mathrm{repr}}}{150\,\mathrm{kA}} + c_{I_{\Omega}^{\mathrm{final}}} \frac{I_{\Omega}^{\mathrm{final}}}{150\,\mathrm{kA}} + c_{\eta_{\mathrm{cond}}} \frac{\eta_{\mathrm{cond}}}{10\,\%} + c_{\tau_{\mathrm{CQ}}} \frac{|\tau_{\mathrm{CQ}} - 100\,\mathrm{ms}|}{50\,\mathrm{ms}} \tag{4.3}$$

where $c_{I_{\mathrm{re}}^{\mathrm{repr}}} + c_{I_{\Omega}^{\mathrm{final}}} + c_{\eta_{\mathrm{cond}}} + c_{\tau_{\mathrm{CQ}}} = 1$. While this would imply $\mathcal{L} \leq 1$ for all four parameters being inside of their safe operational interval, the opposite implication would not be true, namely $\mathcal{L} \leq 1$ implying safety. If there would be a case where only component x is non-zero, this component may rise to $1/c_x$ with $\mathcal{L} \leq 1$ still satisfied.

Both (4.2) and (4.3) are extreme special cases of the p-norm, namely $p = \infty$ and $p = 1$ respectively. The order of the p-norm is thus the trade-off parameter between the largest normalized component dominating the cost function value, and thus loosing information about the other parameters, and having unity as the exact border between safe and unsafe parameter choices. For the prompt to be valid, the

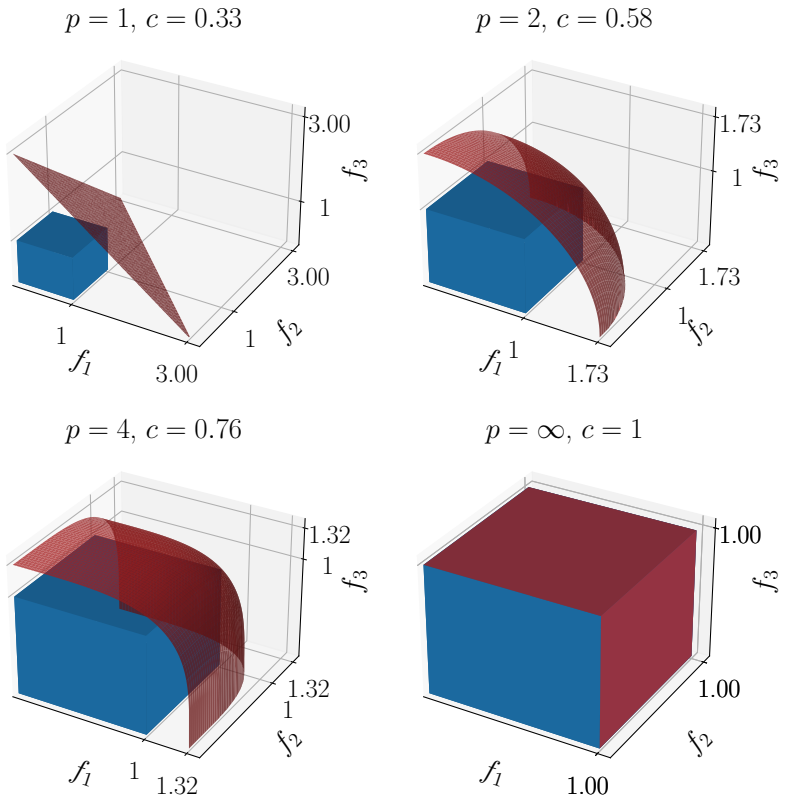

Figure 4.2: Illustration of the relation between p-norm and accuracy of $\mathcal{L} \leq 1$ implying safety. For this example, the cost function consists of three components f_1, f_2 and f_3, weighted equally. The blue box represents the safe operational region and the red surface implies the $\mathcal{L} = 1$ surface. Note that the red surface intersects the axes at $1/c$.

weights must be set as

$$\frac{c_j}{\left(\sum_{i=1}^{n} c_i^p\right)^{1/p}} \rightarrow c_j, \tag{4.4}$$

where c_j for $j = 1, ..., n$ are the components of the cost function. In figure 4.2, the relation between the value of p and the accuracy of $\mathcal{L} \leq 1$ implying safety is illustrated for three arbitrary components f_1, f_2 and f_3. We deemed the Euclidean norm ($p = 2$) to be close enough to $\mathcal{L} \leq 1$ implying safety, without any component being too dominant while large. We chose to use equal weights for all components, specifically $c = 0.5$. With these weights, any component can at most be a factor $1/c = 2$ too large for $\mathcal{L} \leq 1$, and if $\mathcal{L} \leq 0.5$ it is definitely safe.

To summarize the cost function so far, we have four components

$$\begin{aligned} x_{I_{\mathrm{re}}^{\mathrm{repr}}} &= \frac{I_{\mathrm{re}}^{\mathrm{repr}}}{150\,\mathrm{kA}} \\ x_{I_{\Omega}^{\mathrm{final}}} &= \frac{I_{\Omega}^{\mathrm{final}}}{150\,\mathrm{kA}} \\ x_{\eta_{\mathrm{cond}}} &= \frac{\eta_{\mathrm{cond}}}{10\,\%} \\ x_{\tau_{\mathrm{CQ}}} &= \frac{|\tau_{\mathrm{CQ}} - 100\,\mathrm{ms}|}{50\,\mathrm{ms}}, \end{aligned} \tag{4.5}$$

which could be directly combined to yield $\mathcal{L} \lesssim 1$ equivalent with safety, using the weight $c = 0.5$ and the Euclidean norm, accordingly

$$\mathcal{L} = \sqrt{\left(c x_{I_{\mathrm{re}}^{\mathrm{repr}}}\right)^2 + \left(c x_{I_{\Omega}^{\mathrm{final}}}\right)^2 + (c x_{\eta_{\mathrm{cond}}})^2 + \left(c x_{\tau_{\mathrm{CQ}}}\right)^2}. \tag{4.6}$$

This alternative would satisfy all requirements previously stated, but it would give equal significance to minimization of all components (relative to their normalization). For instance, the cost function would take the same value for

$$\begin{aligned} I_{\mathrm{re}}^{\mathrm{repr}} &= 120\,\mathrm{kA} \\ I_{\Omega}^{\mathrm{final}} &= 50\,\mathrm{kA} \\ \eta_{\mathrm{cond}} &= 3\,\% \\ \tau_{\mathrm{CQ}} &= 100\,\mathrm{ms} \end{aligned} \qquad \text{and} \qquad \begin{aligned} I_{\mathrm{re}}^{\mathrm{repr}} &= 0\,\mathrm{kA} \\ I_{\Omega}^{\mathrm{final}} &= 50\,\mathrm{kA} \\ \eta_{\mathrm{cond}} &= 3\,\% \\ \tau_{\mathrm{CQ}} &= 60\,\mathrm{ms}. \end{aligned} \tag{4.7}$$

However, while the runaway current should always be as small as possible, even below its tolerable limit of 150 kA, there is no similar reason for $\tau_{\mathrm{CQ}} = 100\,\mathrm{ms}$ being safer than $\tau_{\mathrm{CQ}} = 60\,\mathrm{ms}$, as they are both within the interval $[50, 150]\,\mathrm{ms}$.

When inside the safe operational space, we would want minimizing the currents to be more important than minimizing the conducted heat, which in turn should be more important than minimizing $|\tau_{\mathrm{CQ}} - 100\,\mathrm{ms}|$. Therefore we would like to introduce the possibility to adjust the sensitivity of the cost to the various components inside their respective safe regions. Outside of the safe operational space, we would not want any term to be too dominating, as previously stated, or for the cost function to blow up – as this could impact the optimizer negatively. Furthermore, it is beneficial for the optimization if the cost function is continuous and differentiable. Therefore, we have instead used

$$\mathcal{L} = \sqrt{\left(c f_{I_{\mathrm{re}}^{\mathrm{repr}}}\right)^2 + \left(c f_{I_{\Omega}^{\mathrm{final}}}\right)^2 + (c f_{\eta_{\mathrm{cond}}})^2 + \left(c f_{\tau_{\mathrm{CQ}}}\right)^2}, \tag{4.8}$$

where

$$f_i = \begin{cases} (x_i)^{k_i}, & \text{if } x_i \leq 1 \\ k_i x_i + 1 - k_i, & \text{if } x_i \geq 1, \end{cases} \tag{4.9}$$

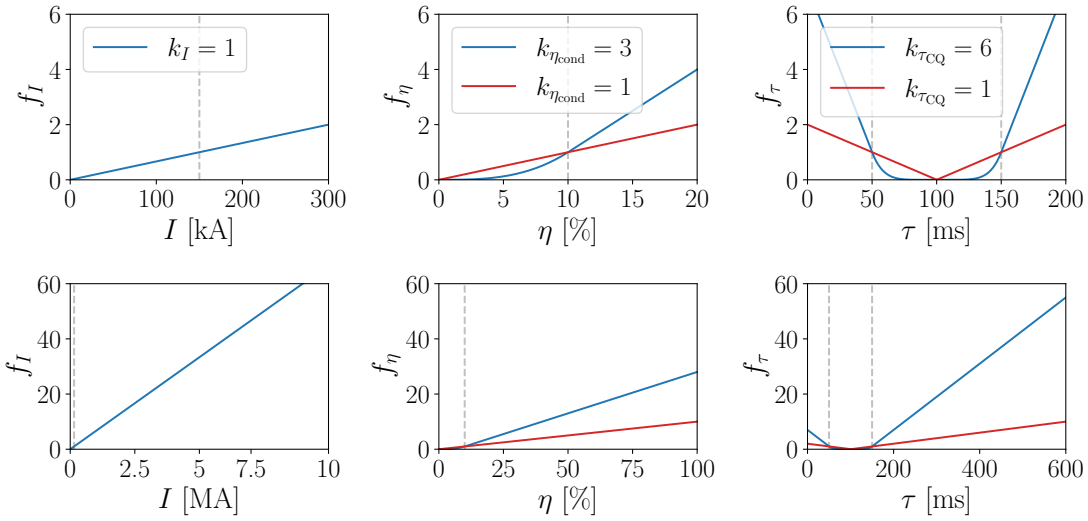

Figure 4.3: Short and long range behaviour of the functions for the currents $I \in \{I_{\text{re}}^{\text{repr}}, I_{\Omega}^{\text{final}}\}$, the conducted heat η_{cond} and CQ time τ_{CQ}. The first row shows the function behaviours close to the safe operational region, while the second row shows the linear behaviours for their full domains. Note that all functions have the same order of magnitude for their long range behaviour.

with $k_{I_{\text{re}}^{\text{repr}}} = k_{I_{\Omega}^{\text{final}}} = 1$, $k_{\eta_{\text{cond}}} = 3$ and $k_{\tau_{\text{CQ}}} = 6$, i.e.

$$f_{I_{\text{re}}^{\text{repr}}} = \frac{I_{\text{re}}^{\text{repr}}}{150\,\text{kA}} \tag{4.10}$$

$$f_{I_{\Omega}^{\text{final}}} = \frac{I_{\Omega}^{\text{final}}}{150\,\text{kA}} \tag{4.11}$$

$$f_{\eta_{\text{cond}}} = \begin{cases} \left(\frac{\eta_{\text{cond}}}{10\,\%}\right)^3, & \text{if } \eta_{\text{cond}} < 10\,\% \\ 3\left(\frac{\eta_{\text{cond}}}{10\,\%}\right) - 2, & \text{if } \eta_{\text{cond}} \geq 10\,\% \end{cases} \tag{4.12}$$

$$f_{\tau_{\text{CQ}}} = \begin{cases} \left(\frac{|\tau_{\text{CQ}} - 100\,\text{ms}|}{50\,\text{ms}}\right)^6, & \text{if } 50\,\text{ms} < \tau_{\text{CQ}} < 150\,\text{ms} \\ 6\left(\frac{|\tau_{\text{CQ}} - 100\,\text{ms}|}{50\,\text{ms}}\right) - 5, & \text{if } \tau_{\text{CQ}} \leq 50\,\text{ms or } \tau_{\text{CQ}} \geq 150\,\text{ms}. \end{cases} \tag{4.13}$$

The behaviour of the cost function components f_I, $f_{\tau_{\text{CQ}}}$ and $f_{\eta_{\text{cond}}}$ close to and far from their safe operational region are plotted in figure 4.3. By implementing these changes inside the safe operational intervals of the CQ time and conducted heat load, they will not be the dominating terms in the cost function if they are relatively far from the interval limits, unless the other cost function terms are negligible as well. This will in turn act to drive the optimization away from the safe interval limits, as intended.

4.2 Initial exploration of optimization space

During the development of the cost function, scans of the activated and non-activated disruption models were made for the optimization using the fluid model input space of logarithmic injected densities,

$$\begin{aligned} \log\left(n_\mathrm{D}/[1\,\mathrm{m}^{-3}]\right) &\in [20, 22.5] \\ \log\left(n_\mathrm{Ne}/[1\,\mathrm{m}^{-3}]\right) &\in [15, 20] \end{aligned} \quad \text{or} \quad \begin{aligned} n_\mathrm{D} &\in [10^{20}, 3.16 \times 10^{22}]\,\mathrm{m}^{-3} \\ n_\mathrm{Ne} &\in [10^{15}, 10^{20}]\,\mathrm{m}^{-3}. \end{aligned} \tag{4.14}$$

The resulting point clouds over the plasma parameters used in the cost function, namely the maximal runaway current $I_\mathrm{re}^\mathrm{max}$ (note that it is not the representative runaway current), final ohmic current I_Ω^final, CQ time τ_CQ and conducted heat load η_cond, are presented in figure 4.4. Each point in the point cloud corresponds to one simulation, where the location is determined by the value of the aforementioned plasma parameters evolved in the simulation. Since the point clouds exist in the four dimensions spanned by the four cost function components, the data points have been projected on the two-dimensional planes. The safe intervals for each component have been plotted as thin, grey lines, and each safe region in two dimensions as black boxes. Furthermore, all data points that are safe – inside the safe interval of each cost function component – are blue while all other data points are red. Since the data clouds are plotted in two dimensions instead of four, some unsafe points are within the bounds of a two-dimensional plot, meaning that they are not within the bounds of one of the other two parameters, which explains why there are red points within the plotted safe boundaries.

There exist safe data points for the non-activated scenario, as shown in figure 4.4a where there is a clear overlap between the data clouds and safe intervals in all projection planes. For the activated case, however, there are no safe data points, as shown in figure 4.4b. Furthermore, the figures suggest it is not possible to achieve any combination of low maximum runaway current, low transported heat fraction and CQ time in the safe interval for the activated case with our MMI model. This is especially clear for the maximum runaway current and conducted heat as all cases with sufficiently low maximum runaway current have $\eta_\mathrm{cond} > 75\,\%$, and all cases with tolerably low conducted heat load have $I_\mathrm{re}^\mathrm{max} \gtrsim 5\,\mathrm{MA}$.

For the non-activated case the $I_\mathrm{re}^\mathrm{max}$–$\tau_\mathrm{CQ}$ and τ_CQ–η_cond plots show a similar behaviour compared to the activated case, but they are slightly deformed such that they overlap with the safe intervals. The $I_\mathrm{re}^\mathrm{max}$–$\eta_\mathrm{cond}$ plot has a completely different appearance. In this thesis we intend to focus on studying the behaviour in the safe regions of the MMI input space, and thus we will only consider the non-activated case for the remainder of this chapter.

Some interesting observations to note from figure 4.4 are the correlations between the cost function components. For both the activated and, especially, the non-activated cases, it seems that just one of I_Ω^final and $I_\mathrm{re}^\mathrm{max}$ can be large at a time – as deduced from the data points lying very close to both axes. Since runaway electron

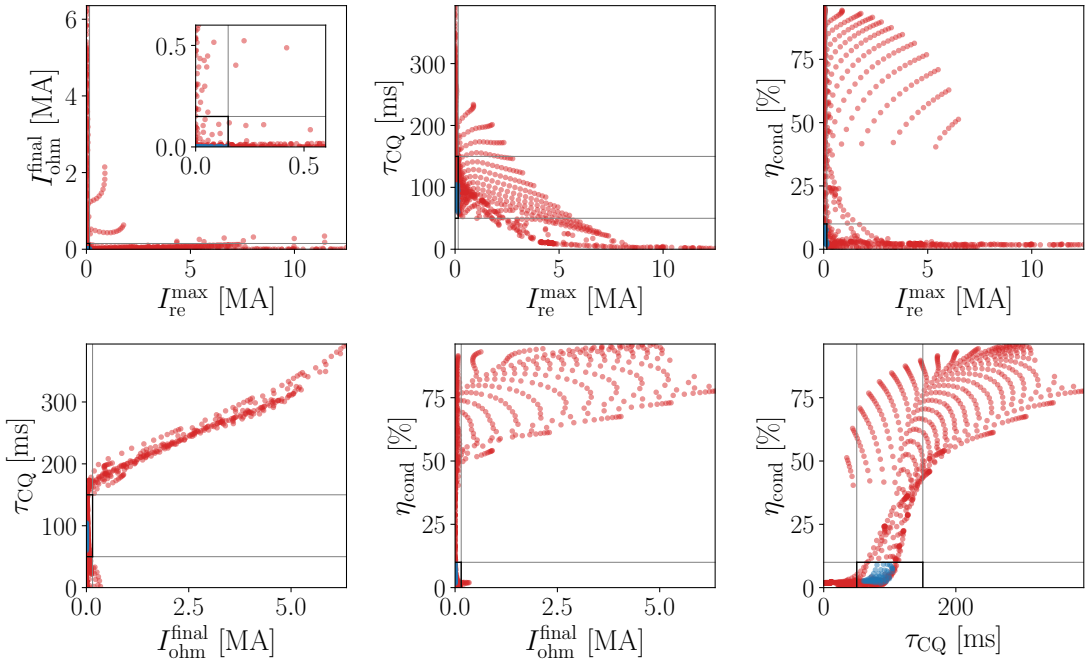

(a) Point cloud for the non-activated disruption model.

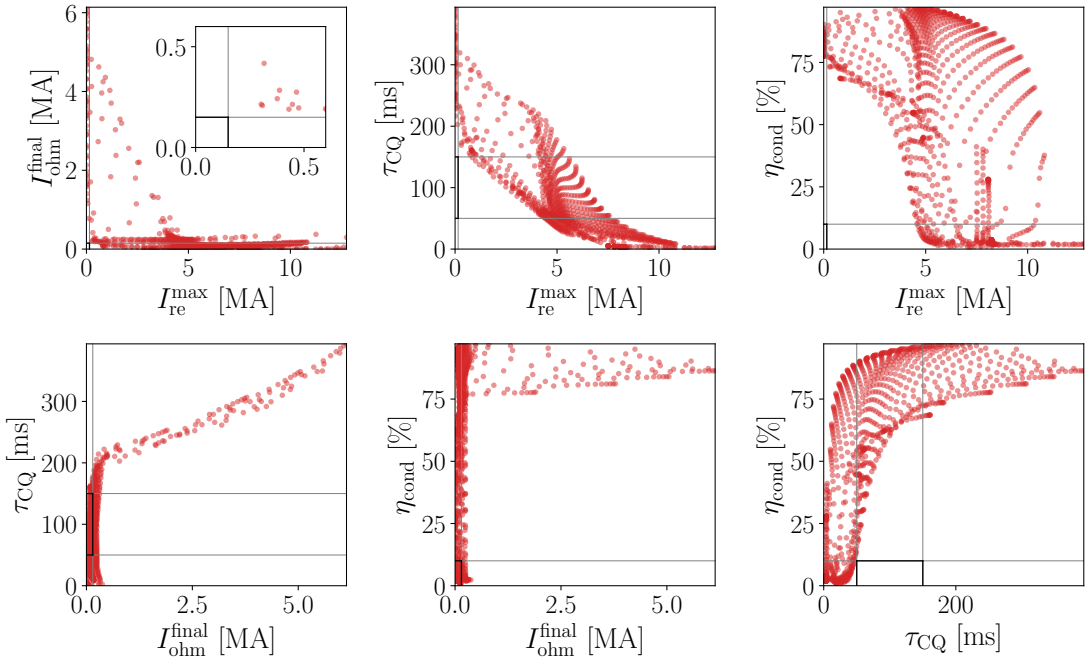

(b) Point cloud for the activated disruption model.

Figure 4.4: A scan over the optimization using the fluid model input space with each point representing a simulation in the scan. With four cost function components, the data cloud becomes a surface in the four dimensional space. Here, it is projected on all possible two-dimensional sub-spaces for the four cost function components; representative runaway current, final ohmic current, CQ time and transported heat fraction.

generation is highly dependent on the electric field, and the electric field is induced by the current drop during the CQ, a large ohmic current would suggest a small induced electric field, which would be unfavourable for runaway electron generation.

Another interesting relation is between the final ohmic current and CQ time, where the current is very low until $\tau_{\mathrm{CQ}} \sim 200\,\mathrm{ms}$, followed by a correlation $\sim 1/(c - x)$. This is expected, as the two components are closely linked through equation (4.1), and suggests that the numerator of (4.1) is approximately constant while $I_{\Omega}^{\mathrm{final}}/I_{\mathrm{p}}$ increases.

4.3 Optimization algorithm

In this section the method used for the optimization will be presented. Firstly, the details and composition of the black box used for the optimization will be described. Secondly, the Bayesian optimization algorithm and the related parameter and model choices will be explained. Both the optimization of uniform and radially varying densities used analogous algorithms for the optimizations. Thus, they will be presented simultaneously and the differences will be discussed.

4.3.1 Black box

In Bayesian optimization, the objective function is treated as a black box, since the optimization process only depends on the inputs and outputs. In this thesis, the black box contains the DREAM simulations described in chapter 3 and the cost function evaluation described in 4.1. Furthermore, our black box has either two or four inputs – the densities n_{D} and n_{Ne} in the two-dimensional optimizations as well as the density profile parameters c_{D} and c_{Ne} in the four-dimensional optimizations. These are MMI parameters and used in connection with the material injection. The injected material is initialized as neutral and at a negligible temperature compared to the plasma temperature ($\sim 20\,\mathrm{keV}$).

The numerical modelling in the black box is divided into several simulations, most importantly separating the TQ and CQ phases. Firstly, a test TQ simulation is performed to determine when the thermal quench is finished. This simulation is performed for 20 ms, adding a margin of order 10 to the anticipated TQ time of 1–2 ms according to Ref. [5] – if the TQ time is longer than this, we can assume that the TQ was not successfully completed and other complications might arise. The condition used to determine when the TQ is completed is that the averaged plasma temperature has dropped below 20 eV, i.e. below 0.1 % of the initial plasma temperature on axis.

After the TQ time has been assessed, the actual TQ simulation is performed with the duration of the simulation having been updated to the time at which the previously mentioned condition was fulfilled in the temporary TQ simulation. Finally, a CQ simulation is performed for 200 ms, as this was deemed long enough to determine if the CQ time is within its safe limits. With the isotropic model, another three

initial simulations are needed, in order to evaluate the initial electric field needed to achieve the correct initial current density. This is needed because the isotropic model requires a consistent steady-state solution of the Fokker–Planck equation for the given electric field, which is not known *a priori*. These initial simulations are not needed with the fluid model, because it has a function to prescribe the initial current density – this only requires Ohm's law to be solved for the electric field, which DREAM can perform in the initialization stage.

The effect of magnetic perturbations, arising from magnetic flux surfaces breaking up during the TQ, are modelled in the simulations by adding runaway electron and thermal energy transport. The magnetic perturbation value is fixed for each optimization, but it was varied between optimizations. We have optimized for magnetic perturbations of $\delta B/B = 0.2\,\%$ and $\delta B/B = 0.5\,\%$. These values of the magnetic perturbations are only used in the TQ simulations. For the CQ simulation, the magnetic perturbations should be turned off, since the magnetic flux surfaces are re-healed, and this is done for the runaway electron transport. A small value of $\delta B/B = 0.04\,\%$ is still used for the thermal energy transport, however, to avoid numerical errors due to non-physical thin current sheets. All electron populations in the simulations account for the particle, and subsequent heat, transport due to the magnetic perturbations – namely the cold and runaway electrons in the fluid model and additionally the hot population in the isotropic model.

When both the TQ and CQ simulations have been performed, the cost function value for an input pair (n_D and n_Ne), or quadruple (n_D, n_Ne, c_D and c_Ne), is evaluated according to section 4.1. However, in some instances, the cost function value cannot be evaluated. In these cases, the cost function returns a large value, which we chose to be 75 because it is larger than, but still comparable to, most encountered cost function values for our input spaces. There are mainly two reasons for why the cost function cannot be evaluated.

In most cases when the cost function cannot be evaluated, the temperature does not reach below 20 eV during the initial TQ simulation, which signifies an incomplete TQ. Since this means that the TQ time is more than an order of magnitude larger than the anticipated TQ time of ITER at 1–2 ms, we assume the TQ will not have completed successfully and does not represent an interesting (safe) scenario, and thus, there is no need to perform the rest of the simulations. It was discovered that incomplete TQs occur in a certain region of the input space for each magnetic perturbation value. A hyperbola was fitted to the boundary of this region in the n_D–n_Ne input space for the fluid simulations, such that the black box returns the large value directly when a n_D–n_Ne input pair was within this region and no simulations were performed. No similar approximation of the boundary of incomplete TQs was performed for the four-dimensional optimizations.

The second case when the cost function cannot be evaluated is when the simulation fails. When using numerical tools, such as DREAM, there is always the risk of the time resolution not being sufficient to find converging solutions, resulting in

the simulations crashing, however there can be other reasons for the simulations failing as well. The black box has the ability to increase the time resolution if such complications occur – but not indefinitely. We set upper limits for the time resolution, and if these are reached the simulations are aborted and the cost function returns the previously mentioned large value. It is worth noting however that these simulation failures that could not be resolved with increased time resolution were rare, and only occurred for less than 1 % of the optimization samples, for both the fluid and isotropic models.

Resolution is as previously mentioned important for the numerical accuracy and reliability of the code, and now the choices regarding resolution in this thesis will be presented. With regards to reliability, higher resolution is always preferred, but higher resolution leads to longer computational times which we want to minimize in this project and all our resolution choices represent a trade-off between these two factors. We decided to use 20 radial grid points, and for the isotropic simulations, 80 momentum grid points were also used, with the upper limit of the momentum space, i.e. the critical momentum value separating the hot and runaway populations, being $p = 2.5 m_e c$.

The time resolution is handled differently for the TQ and CQ simulations. The TQ simulations used the adaptive time stepper function of DREAM, that chooses the time step length on the ionization time scale. The initial time step length was 10^{-11} s in the fluid simulations and 10^{-13} s in the isotropic simulations, while the maximum time step length was 10^{-5} s for both. If the TQ simulations crashed due to insufficient time resolution, the initial time step length could be decreased to at most 10^{-13} s and 10^{-14} s for the fluid and isotropic simulations, respectively. For the CQ simulations, the simulation length was as previously stated 200 ms and initially 10 000 time steps were used for both fluid and isotropic simulations. When numerical complications due to time resolution were encountered, the number of time steps could be increased up to 200 000 in the fluid simulations and up to 50 000 in the isotropic simulations.

4.3.2 Bayesian optimization

For this thesis, we have utilized the Bayesian optimization Python package [49]. The Bayesian optimization functions of this package require a hyper-rectangle in the input space to be used as optimization bounds. Initially, the optimizer explores a specified number of data points chosen randomly before proceeding to using the acquisition function to select new data points. The selection of further data points is then based on the already known data points. Furthermore, the GP kernel and acquisition function type can be specified and the ratio between exploration and exploitation can be controlled.

Since the injected densities for MMI varies with orders of magnitudes, the common logarithm of the injected densities have been used as optimization parameters. The

bounds of the densities have been chosen as

$$\begin{aligned} n_{\mathrm{D}} &\in [10^{20}, 3.16 \times 10^{22}]\,\mathrm{m}^{-3} \\ n_{\mathrm{Ne}} &\in [10^{15}, 10^{20}]\,\mathrm{m}^{-3} \end{aligned} \quad \text{or} \quad \begin{aligned} \log\!\left(n_{\mathrm{D}}/[1\,\mathrm{m}^{-3}]\right) &\in [20, 22.5] \\ \log\!\left(n_{\mathrm{Ne}}/[1\,\mathrm{m}^{-3}]\right) &\in [15, 20] \end{aligned} \tag{4.15}$$

based on [12], where similar bounds were used and an optimum was found within these bounds. More specifically, the upper bounds are chosen by experimental limitations – higher concentrations of these gases cannot be assimilated in the plasma – and the lower limits are chosen because they are on the border of where the injected quantities begin to have an effect on the plasma evolution. For the radially varying densities, the profile parameter bounds have been chosen as

$$\begin{aligned} c_{\mathrm{D}} &\in [-5, 5] \\ c_{\mathrm{Ne}} &\in [-5, 20]. \end{aligned} \tag{4.16}$$

The profile parameter intervals are centred around the optimal values found in Ref. [12].

When choosing the number of sample points to use for the optimization, there is a trade-off between how well the input area is explored and the computational time needed to obtain the objective function value. How "well the input area is explored" refers to how scattered the sample points are in the input space, which determines how probable it is that the optimum discovered is the global optimum (as opposed to a local optimum) and how accurate the mean function of the GP, or the GP approximation of the cost function, is. Furthermore, simulations in DREAM take a considerable amount of time – one fluid simulation takes $\sim 5\,\mathrm{min}$ while one isotropic simulation takes $\sim 30\,\mathrm{min}$. For the two-dimensional optimization using the fluid model, 50 data points were used for the initial random exploration and 200 data points were used in the actual optimization, as this was deemed to yield a sufficient exploration in both interesting (safe) and uninteresting areas. Since the optimization region of the isotropic case is smaller, as described later in this section, the two-dimensional optimization using the isotropic model used 20 initial data points and 80 optimization data points. When the dimensionality is doubled, which is the case for the radially varying density optimization, more data points are needed to explore the now larger optimization space. For this reason the number of data points used were doubled, namely 100 (40) initial data points along with 400 (160) optimization data points for the optimization using the fluid (isotropic) model.

The Matérn Kernel was used for the optimizer GP, with the smoothness parameter $\nu = 1.5$ was chosen because our function is once differentiable (see section 2.2) – at the borders of the safe intervals of τ_{CQ} and η_{cond} the cost function is only once differentiable. For the acquisition function, we chose EI, defined in (2.16) due to it being widely used in Bayesian Optimization for its empirical effectiveness [50]. In Ref. [49], the EI has been modified by introducing an exploration-exploitation trade-off parameter, ξ, to (2.16) according to

$$I(\boldsymbol{x}, \phi, y^*) = \max(\phi - y^*,\ 0) \to \max(\phi - y^* - \xi,\ 0), \tag{4.17}$$

where exploration increases with ξ. For both the optimizations using the fluid and isotropic models, we wanted a well-resolved map of the cost function in the basin of safe values containing the global optimum. If ξ was chosen too small, most of the samples taken by the optimizer would be placed in a very small vicinity of the optimum, which would not be suitable to produce this map. Thus, we chose $\xi = 0.05$ for the optimizations using the fluid model and $\xi = 0.1$ for the optimizations using the isotropic model, taking into the consideration the size of the optimization region compared to the interesting region for each plasma model.

To visualize the cost function after an optimization, a GP was used to fit an approximation μ to the cost function based on the optimization data points. The GP used for visualizing the results was identical to the optimizer GP with $\nu = 1.5$. One difference from the optimizer GP however was that the data points corresponding to a failed simulation, or a simulation when the cost function could not be evaluated and instead 75 was returned, were disregarded. Since 75 did not quantitatively represent the actual cost function value in these data points – the high value was returned because we did not know the actual value – they did not give relevant information for the GP regarding approximating the cost function. In reality, they would have obstructed the cost function approximation because they were not continuous with other data points in the vicinity, making it difficult for the GP to find a valid and continuous approximation.

Finally, the remaining part of the method left to be described is how the bounds used for the optimization of the isotropic case were determined based on the optimization of the fluid case. Concisely, the idea of the method is to evaluate the approximated cost function μ on a hyper-rectangle of increasing size until all cost function values on the sides of the hyper-rectangle are above one – and this hyper-rectangle is then used as the bounds for the optimization using the isotropic model. In practice, only a two-dimensional algorithm was implemented. This will now be presented and we will subsequently explain how this was used in higher dimensions.

A GP-generated cost function approximation was obtained from the optimization data of the optimization using the fluid model, as described above. Using this approximation, a rectangle was centred around the minimum of μ, with the length of the sides being defined from some minimum length of the bound interval. If all values of the approximation were above one on the rectangle, or the side has reached the bounds of the optimization space when using the fluid model, these will be used as the new optimization bounds, see figure 4.5b. Otherwise, each side of the rectangle is moved one small increment, connected to the resolution of the GP approximation, outwards sequentially until all values of the approximation on a side are above one, as in figure 4.5a. Thus, suitable optimization bounds are found when all sides have values above one, or some of the sides of the rectangle reaches the bounds of the optimization space when using the fluid model. In the case when there is no region of safe values below one in the input space, the rectangle of given minimum side lengths centred around the optimum will be used as new bounds.

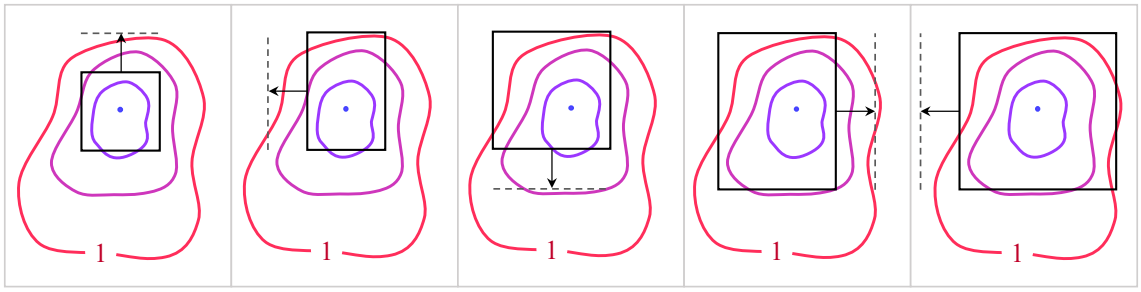

(a) A rectangle is centred around the optimum, and each side is extended sequentially until all values of the cost function on each side of the rectangle are above one. Note that each side will be fixed independently from the others.

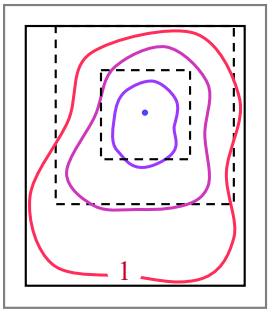

(b) The dashed rectangles represent cases where there are still values on the rectangle that are below one. The solid black box is the smallest rectangle where all values are above one on the box. The grey box represents how margins can be added by the algorithm.

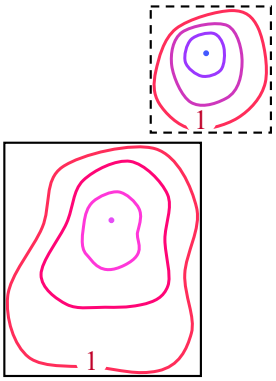

(c) For the case when there are several regions with values below one, the algorithm can choose bounds that maximizes the rectangle area. In this illustration, the global minimum is within the smaller, dashed rectangle, but since the other rectangle is larger (corresponding to a more robust minimum) it will be used for the new optimization bounds.

Figure 4.5: Visualization of the algorithm which is used to find the bounds for the optimization using the isotropic model. The coloured curves are level curves encircling a local minimum. The colour scale of the level curves and optimum goes from blue for low values to red for higher values

As larger regions of safe values are more robust, a feature of our algorithm was to return the bounding rectangle of the largest safe region in the input domain. Thus the same procedure was performed around all local minima. If another safe region corresponds to a larger new bound rectangle, this rectangle will be used for the bounds instead, as illustrated in figure 4.5c. Lastly, a margin can be added to the final bounding rectangle, see figure 4.5b.

The computational resources required by this method increase exponentially with dimensionality, which results in it not being feasible to adapt this method in two dimensions to four dimensions within the bounds of this thesis. Regarding the

optimization of radially varying densities, the two-dimensional method was used for every pair of input parameters while keeping the other two input parameters fixed at their optimal value. For each input parameter, the lowest lower bound and highest upper bound found were used. It is worth noting that since this method does not consider the whole of the four-dimensional space, it can cut off regions with safe operational values that could be of interest to study using the isotropic model. This risk was deemed to have small enough impact on the results to be disregarded in this project. The exploration of other safe operational regions in four dimensions become problematic with the two-dimensional method, since two separate regions in two dimensions can be connected in higher dimensions. Because of this, the two-dimensional boundary adjustment method was not implemented for the four-dimensional case.

This method for finding new bounds has a couple of free parameters – the size of the margins and resolution of the GP approximation. The optimization of uniform density used margins of 2 % of the optimization bound length when using the fluid model and a resolution of 400 data points in each direction when evaluating the approximation of the cost function, i.e. resulting in a grid of $400^2 = 160\,000$ cells. To mitigate the risk of cutting of regions of safe values when using the method of finding new bounds in two dimensions for the four dimensional optimization, a larger margin of 5 % was used in the radially varying density optimization to increase the probability of enclosing the entirety of the safe operational region in the hyper-rectangle. The minimum bounds interval size for each input parameter, i.e. the minimal side of the hyper-rectangle in each input parameter direction, was selected as 0.5 for $\log n_{\mathrm{D}}$, 1 for $\log n_{\mathrm{Ne}}$, and 1 for both c_{D} c_{Ne}.

4.4 Optimization with uniform densities

In this section the results of the optimization of MMI with radially uniform densities will be presented and discussed.

4.4.1 Results

The two-layer optimization was performed for radially uniform deuterium and neon densities and magnetic perturbations $\delta B/B = 0.2\,\%$ and $\delta B/B = 0.5\,\%$, and the results can be found in figure 4.6. In table 4.1, the minima of different optimizations are presented together with another isotropic optimization that will be presented below.

For the lower magnetic perturbation case, shown in figure 4.6a, there is a fairly large region of safe values for both models and the shape and size of both of these regions are somewhat similar. The two models predict different locations for the minimum, with values given in the first two rows of table 4.1. Close to the optimum discovered using the fluid model, the two models give similar cost function values, but with the isotropic model yielding $\mu = 0.025$ being somewhat lower than the fluid model yielding $\mu = 0.037$. The difference between the two models is most significant in the

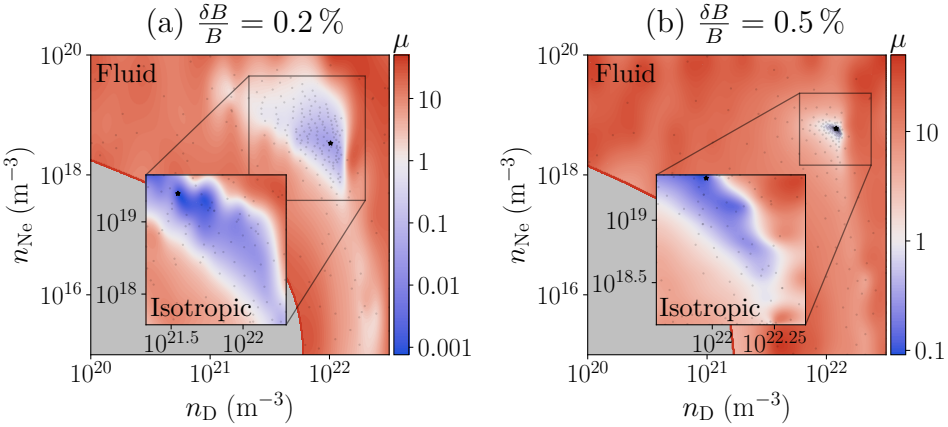

Figure 4.6: Contour plots of the approximated cost functions found using the two-layer, two-dimensional, optimization for two different magnetic perturbations. The large plots in the background are the results generated by the fluid model, and the smaller plots in the lower left corners are obtained by the isotropic model with the inset representing the explored region. The black star represents the minimum of the optimization and the dots are the samples taken by the optimizer. Note that the colour mapping is adapted such that blue shades are below one and red values are above one, and white corresponds to one. The grey areas in the lower left corners are the regions where the TQ was not successful.

upper left quadrant of the input space explored by the isotropic model, around the minimum discovered using the isotropic model. Here, the values of the cost function approximation for the optimization using the fluid model are in the vicinity of one, while the approximation based on the isotropic model is on the order of 0.001.

For the high magnetic perturbation, shown in figure 4.6b, again, safe regions exist with both the fluid and isotropic models, but there are now large disparities between them. The safe region with the fluid model is small enough to use the minimal bound lengths (0.5 for $\log n_{\mathrm{D}}$ and 1 for $\log n_{\mathrm{Ne}}$) to be used for the new bounds instead of the actual bounds of the safe region. On the other hand, the safe region in the isotropic case is reaching outside of its optimization region. In fact, the minimum discovered for the isotropic case might not be the actual minimum of this region since it lies very close to the upper optimization boundary of the neon density. As with $\delta B/B = 0.2\,\%$, the behaviour is somewhat similar between the two models close to the optimum discovered using the fluid model, but while the cost function increases with the fluid model in the upper left direction, it decreases with the isotropic model. However, the cost function does not decrease by several orders of magnitude with the higher magnetic perturbation, as it did for the lower perturbation case, when increasing the neon density and lowering the deuterium density.

In order to further explore the safe region of the isotropic model with the higher magnetic perturbations, and find the minimum with better confidence, another optimization using the isotropic model was performed but with manually determined

Table 4.1: Table listing the minima found using different kinds of optimizations. The first two columns specify the physics model used. The following four columns describe the optimal deuterium and neon densities. Finally, the last column lists the cost function value of the minimum.
* Second isotropic optimization with manually defined bounds.

Model	$\delta B/B$ [%]	$\log n_\mathrm{D}$	n_D [m^{-3}]	$\log n_\mathrm{Ne}$	n_Ne [m^{-3}]	$\min \mathcal{L}$
Fluid	0.2	22.01	$1.02 \cdot 10^{22}$	18.53	$3.37 \cdot 10^{18}$	0.04
Isotropic	0.2	21.55	$3.52 \cdot 10^{21}$	19.39	$2.47 \cdot 10^{19}$	0.001
Fluid	0.5	22.08	$1.21 \cdot 10^{22}$	18.77	$5.93 \cdot 10^{18}$	0.4
Isotropic	0.5	21.98	$9.46 \cdot 10^{21}$	19.34	$2.17 \cdot 10^{19}$	0.1
Isotropic*	0.5	21.95	$8.83 \cdot 10^{21}$	19.38	$2.40 \cdot 10^{19}$	0.1

boundaries. This time the density boundaries used were

$$\begin{aligned} n_\mathrm{D} &\in [3.16 \times 10^{21}, 1.78 \times 10^{22}]\,\mathrm{m}^{-3} \\ n_\mathrm{Ne} &\in [3.16 \times 10^{15}, 5.62 \times 10^{20}]\,\mathrm{m}^{-3} \end{aligned} \quad \text{or} \quad \begin{aligned} \log\Big(n_\mathrm{D}/[1\,\mathrm{m}^{-3}]\Big) &\in [21.5, 22.25] \\ \log\Big(n_\mathrm{Ne}/[1\,\mathrm{m}^{-3}]\Big) &\in [18.5, 19.75], \end{aligned} \tag{4.18}$$

as these were assumed to encompass the safe region discovered using the isotropic model better. The results of this new optimization using the isotropic model are presented in figure 4.7. With these new optimization boundaries, the safe operational region discovered using the isotropic model is much better encompassed. As this new optimization shows, the safe region discovered using the isotropic model is much larger than in figure 4.6, but the minimum did not move significantly; compare the last two rows of table 4.1.

To further explore the differences between the fluid and isotropic models, each sample taken during the optimization using the isotropic model with manually defined boundaries was repeated using the fluid model. In figure 4.8, a GP approximation of the components of the cost function, namely the representative runaway current, final ohmic current and conducted heat load, for both the fluid and isotropic models are presented.

In figure 4.8, we see that the runaway current deviates the most in the two models. The normalized representative runaway current, shown in figure 4.8a, is significantly smaller in the isotropic case – and there are several orders of magnitude of difference between the two models for the majority of the explored region. Whereas the fluid model has a normalized runaway current below one in a fraction of the lower half of the explored space, the isotropic model predicts negligibly small values in almost three quarters of the input space – and several optimization samples have no runaway current.

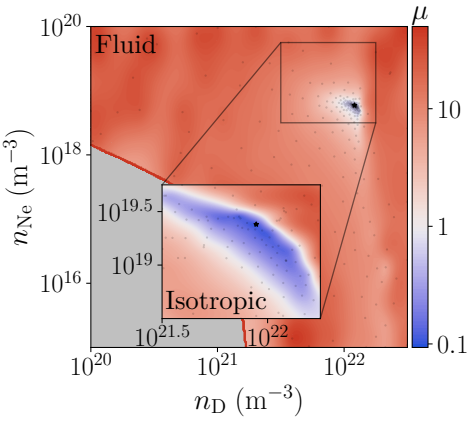

Figure 4.7: Contour plots of the approximated cost functions for $\delta B/B = 0.5\,\%$ using the fluid and isotropic models. The bounds of the optimization using the isotropic model have been set manually. Note that the plot with the fluid model is the same as in figure 4.6b, and is displayed to demonstrate the new optimization space and simplify the comparison.

There is also a fairly large difference between the final ohmic currents in the two models, as shown in figure 4.8b, with several orders of magnitude of difference in the upper left half of the region. With both models, most parts of the regions are below one, representing safe ohmic currents, but the isotropic model, again, gives more optimistic predictions.

The transported heat fraction is quantitatively and qualitatively very similar, as shown in figure 4.8c, but it is the only component for which the fluid model gives better results than the isotropic model. However, this difference is very small – small enough to not impact the overall improvement of the isotropic model over the fluid one. Furthermore, there are only insignificant differences in CQ time between the models.

To summarize, the isotropic model favours higher neon densities and lower deuterium densities compared to the fluid model. In the lower magnetic perturbation case, the shapes and sizes of the safe regions in the two models are comparable, but inside the safe regions the cost function values can differ by orders of magnitude. In the higher magnetic perturbation cases, the magnitude of the cost function in its respective safe regions did not vary as much, but the shape and size of the safe regions are profoundly different. The main reason for the differences between the two models is the currents, and especially the runaway current.

4.4.2 Discussion

The main reason for the difference in location of the minimum in the two models is the current components of the cost functions. These decrease with the isotropic model and increase with the fluid model when increasing the neon density and

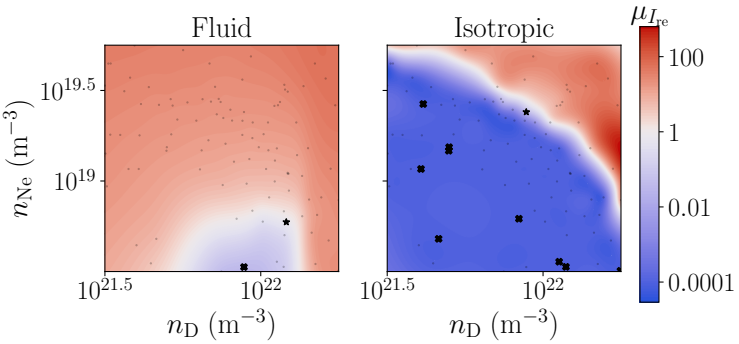

(a) Contour plot of the representative runaway current. There are several black crosses, signifying the minimum runaway current, because all of these samples resulted in no runaway current and are global minima.

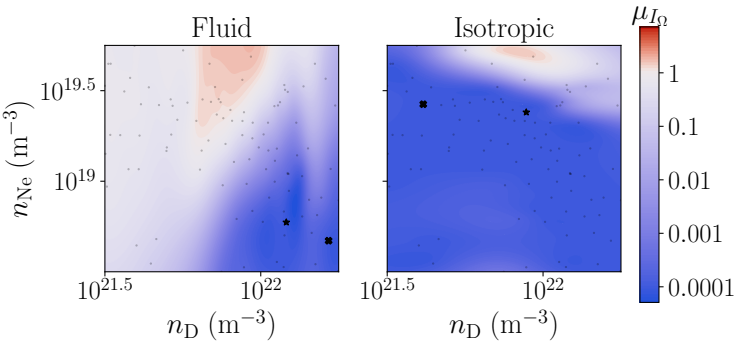

(b) Contour plot of the final ohmic current component.

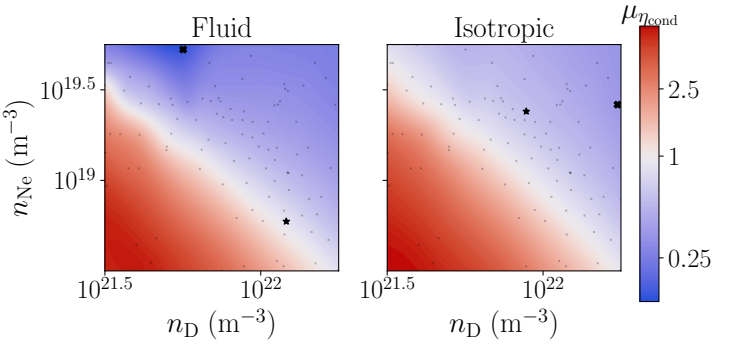

(c) Contour plot of the conducted heat load component.

Figure 4.8: Contour plots of approximations to the normalized cost function components, as functions of deuterium and neon densities, for both the fluid and isotropic models. The black star represents the minimum of the optimization, the black cross the minimum of the plotted component and the dots the optimization samples.

decreasing the deuterium densities from the optimum discovered using the fluid model, while the CQ time and conducted heat load do not vary significantly.

Generally, the optimum is determined by the trade-off between the runaway current – favouring low injected densities – and conducted heat load – favouring high injected densities. As figures 4.8a and 4.8c show, the minimum discovered using the fluid model lies on the very edge of the safe region for both the conducted heat and representative runaway current and there is no point with better trade-off. For the isotropic case, the optimum still lies fairly close to the safe boundary for the representative runaway current, but further away from the boundary for the conducted heat. The main reason for this is the fast decrease of the representative runaway current when moving away from its safe border. Because such a minimum is very sensitive to small positive deviations of the injected densities, this optimum could be a bad alternative for MMI in a real tokamak, owing to the sensitivity to the safe operational limits of the runaway current and conducted heat. However, the sensitivity of the optimum should be studied further, since sufficient accuracies in injection may perhaps be achievable, noting that the densities are logarithmic in the plots.

For the higher magnetic perturbation case, the cost function value of the minimum discovered using the isotropic model is still of the same order of magnitude as the optimum discovered using the fluid model – this is because the conducted heat load is large enough to dominate the cost function when the currents are minimized. We can be fairly certain that the CQ time component does not dominate here because of the high order of $f_{\tau_{\mathrm{CQ}}}$ in equation (4.13) – the normalized CQ time of the explored region is $x_{\tau_{\mathrm{CQ}}} \sim 0.5$ which would yield $f_{\tau_{\mathrm{CQ}}} \sim 0.01$. Furthermore, both the normalized runaway and ohmic current should contribute negligibly to the cost function in the isotropic case. The normalized conducted heat, with the isotropic model, in the upper left half of the explored region in figure 4.8c is ~ 0.8 which is equivalent to $\eta_{\mathrm{cond}} = 8\,\%$ and would yield $f_{\eta_{\mathrm{cond}}} = 0.8^3 \approx 0.5$, which would dominate the other components. However, the lower magnetic perturbation value must have a significantly lower transported heat fraction than the higher perturbation case for the difference in the cost function values to be over several orders of magnitude, as observed here. This is not unexpected as the heat transport increases with magnetic perturbation value.

The main reason for the representative runaway current being larger in the fluid case might be because the Dreicer and/or hot-tail rate formulae used in the fluid model overestimate the actual rates observed with the isotropic model, which will be further studied in chapter 5. This is possibly not the entire reason for the difference in runaway currents between the models, as the CQ is more successful for not only the representative runaway current, but final ohmic current and CQ time as well, in the isotropic case.

4.5 Optimization with radially varying densities

In this section the results of the optimization of MMI with radially varying densities will be presented and discussed.

4.5.1 Results

Only the higher magnetic perturbation case $\delta B/B = 0.5\,\%$ was considered when the optimization was performed for radially varying deuterium and neon densities. The results of the optimization using the fluid model are presented in figure 4.9 together with the new optimization bounds used for the optimization using the isotropic model – since the optimization space is four dimensional, the figures show contour plots of two dimensional cross sections that include the optimum point.

Compared to the fluid case for the uniform densities, we now obtain a much lower optimum value and a larger safe operational area in the n_{D}–n_{Ne} plane, at the optimal profile parameter values. This region is still slightly smaller than the minimum bound length in the density directions, just like in the two-dimensional case, but in the profile parameter directions it is large enough. Interestingly, there is little variation in the neon profile parameter direction resulting in the new bounds for the neon profile parameter being the same with the isotropic case as in the fluid case.

In figure 4.10 the result from the optimization using the isotropic model is presented together with the corresponding result using the fluid model evaluated at the samples from the optimization using the isotropic model. As in the uniform density case, the regions are larger with the isotropic model compared to the fluid model, but the difference between the models is smaller than for the uniform density case. Interestingly, in the 4D optimization the fluid case yields a better result than the isotropic case by a factor of 10, whereas the opposite was true for the two-dimensional case. However, both cost function values, 0.05 (fluid) and 0.005 (isotropic), are small enough to be considered sufficiently safe in a real tokamak. The isotropic model has mostly moved the optimum in the deuterium density and profile parameter directions.

The specifics of the minima discovered in the optimizations using the fluid and isotropic models are presented in table 4.2 together with the fluid model specifics in the optimum discovered using the isotropic model. As indicated by the lower limit of the colour bars in figure 4.10 as well, the minimum is one order of magnitude smaller in the optimum discovered using the fluid model compared to the optimum discovered using the isotropic model. Interestingly, the cost function value from the fluid model in the optimum discovered using the isotropic model is smaller than the minimum discovered during the optimization using the fluid model, suggesting that this optimum was missed during the optimization using the fluid model. This indicates that the qualitative behaviours of the models when using radially varying injection densities are more similar between the fluid and isotropic models than in the uniform case.

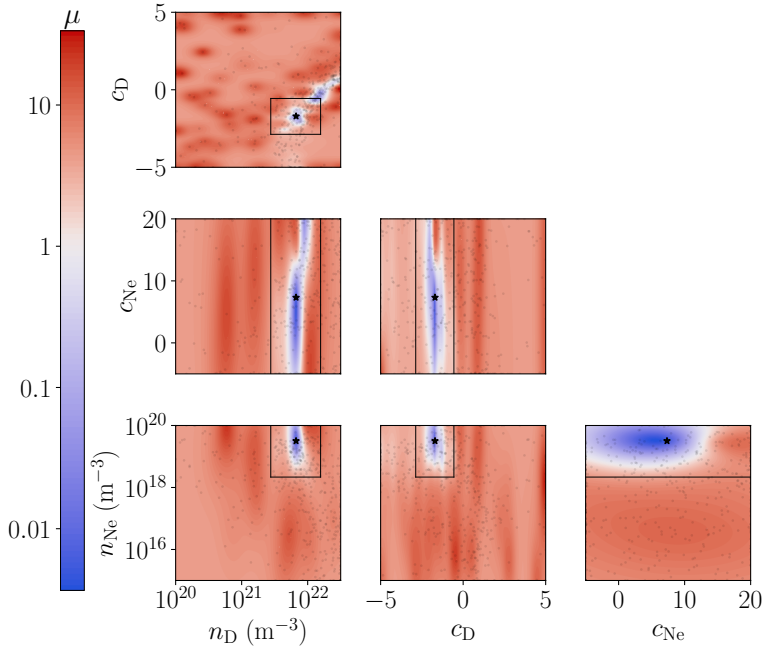

Figure 4.9: Contour plot representation of the GP approximated cost function for the four-dimensional input space of radially varying deuterium and neon densities. The domain of the GP approximations is four-dimensional, which is impossible to visualize directly. Instead, the contour plots are cross sections in the optimal planes, i.e. for each contour plot, two parameters are fixed at their optimal value. For example for the n_D–n_{Ne} plot in the lower left corner, $c_D = -1.7$ and $c_{Ne} = 7.3$. The black star represents the minimum of the optimization and the dots the optimization samples, which have been projected on the planes, and the box represents the new bounds used for the optimization using the isotropic model.

The main reason for the cost function being smaller with the fluid model is the value of the conducted heat load, which dominates in all three cases. Because of the cost function's dependence on the conducted heat load is cubic for $\eta_{cond} < 10\,\%$, a reduction by a factor of 2.5 for the conducted heat load η_{cond} will decrease the conducted heat term of the cost function $f_{\eta_{cond}}$ with a factor of 15.

Besides the conducted heat, the final ohmic current is the only other component of significant value to impact the cost function, and it is only significant for the isotropic case. In general though, all components are well within the safe operational limits and such a case would pose a minimal risk in our ITER-like scenario.

4.5.2 Discussion

From a practical point of view it is advantageous that the safe operational space has little dependence on the radial density profile parameter for neon, based on the large size of the optimization bounds for c_D in figure 4.9. This indicates that it mainly matters that the neon is present and peaked at the edge of the plasma column, but

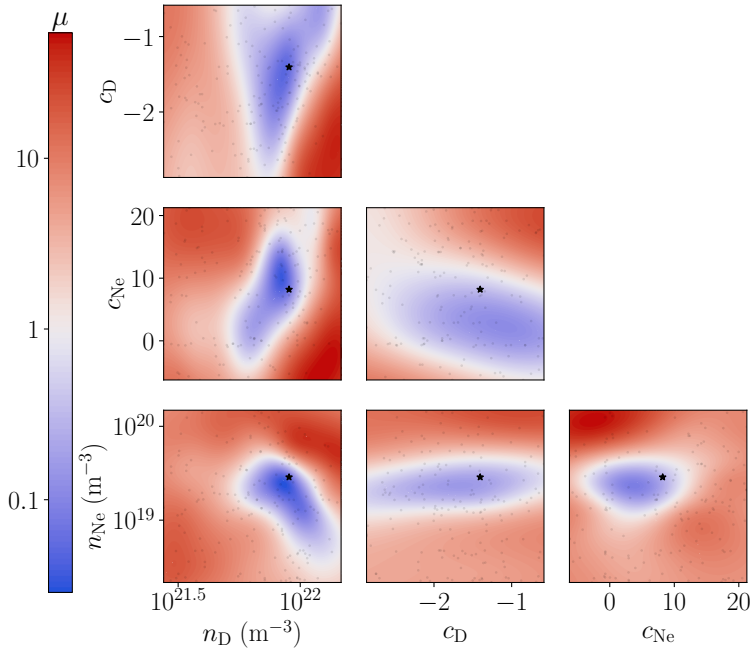

(a) Contour plot representations found using the isotropic model.

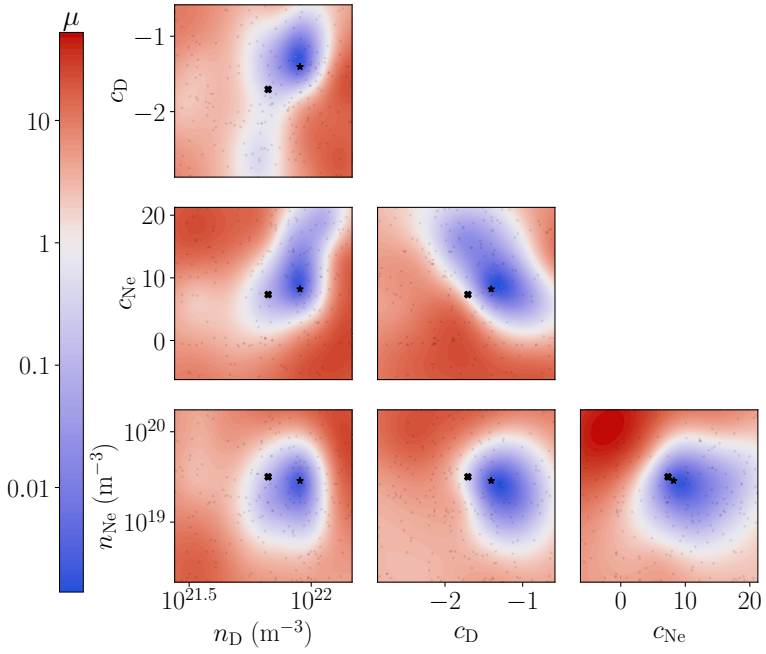

(b) Contour plot representations of the cost function using the fluid model when evaluated at the samples of the optimization using the isotropic model. Note that this is not the same GP approximation as in figure 4.9.

Figure 4.10: Cross sections of the GP approximated cost function that include the optimum point found using the isotropic model. The black star represents the minimum of the optimization using the isotropic model, the black cross the optimum of the optimization using the fluid model (in figure 4.9) and the dots the optimization samples, which have been projected on the planes. Note that the colour scales differ between the two figures.

Table 4.2: Table listing the minima found during the radially varying density optimization. The columns specify the physics model used. The first six rows describe the optimal density combinations and the last five rows describe the value of the cost function and its components at the minimum.
* Fluid repetition of the data points of the isotropic optimization.

	Fluid model	Isotropic model	Fluid model*
$\log n_{\mathrm{D}}$	21.82	21.95	21.95
n_{D} [m^{-3}]	$6.66 \cdot 10^{21}$	$9.00 \cdot 10^{21}$	$9.00 \cdot 10^{21}$
$\log n_{\mathrm{Ne}}$	19.50	19.46	19.46
n_{Ne} [m^{-3}]	$3.17 \cdot 10^{19}$	$2.87 \cdot 10^{19}$	$2.87 \cdot 10^{19}$
c_{D}	-1.7	-1.4	-1.4
c_{Ne}	7.3	8.2	8.2
$\mathcal{L}$	0.005	0.03	0.002
$I_{\mathrm{re}}^{\mathrm{repr}}$ [kA]	0.018	$7.7 \cdot 10^{-5}$	0.17
I_{Ω} [kA]	0.082	3.8	0.40
τ_{CQ} [ms]	78	120	85
η_{cond} [%]	2.0	3.6	1.5

not so much how strong the peaking. In a tokamak it would be difficult to control the radial distributions of the injected material densities because of the diffusion of material in a plasma, and our results indicate that for neon this would not be a critical aspect of the material injection. An edge-localized neon density probably generates the best results as this would establish a region of high radiation and heat transport at the edge of the plasma, while minimizing the avalanche runaway generation in the centre due to the neon density.

For deuterium, there is a stronger correlation between the profile parameter and safety. The most probable reason for a centre-localized injected deuterium density to be favoured is that the deuterium injection would most efficiently dilute the plasma at the centre, where the temperature is the highest.

That the optimal value for the deuterium profile parameter is negative while it is positive for neon could indicate that the deuterium should be injected before the neon giving it more time to migrate into the centre of the plasma column, but due to the complicated dynamics of MMI it probably would not be this simple. One difficulty with MMI during disruptions is that the transport event triggered by the neon injection will not only transport heat in the plasma, but ions as well – in our simulations ion transport is not present. The ion transport significantly complicates control of the radial variation of the injected material, so that achieving the optimal radial density variations found in this thesis might not be impossible.

To fully utilize the results found in this thesis, an in-depth analysis of MMI-schemes should be performed to determine if and how it is possible to achieve a centre-peaked deuterium density and edge localized neon density. One possible MMI scheme to examine would be an initial injection of a deuterium pellet containing a smaller amount of neon impurities. To start with, a deuterium pellet should ensure that the deuterium is mostly deposited at the centre, as this is where the plasma temperature and fuel densities are the largest. The neon impurities in the deuterium pellet would have two purposes – reducing the drift deuterium pellets experience as well as triggering a transport event. When the transport event has been triggered, a massive gas injection of neon could be performed, which would move slower than a pellet, and thus hopefully localize the neon density at the edge. However, it might not be necessary for the optimal profiles to remain throughout the disruption – for example, an edge localized neon that radiates away the heat is only important until the radiative collapse is complete.

The fact that the lowest fluid cost function value evaluated on the sample set from the optimization using the isotropic model is lower than the lowest value discovered by the optimization using the fluid model – that is, that the optimization failed to find a better optimum that clearly exists – indicate that the optimization using the fluid model did not perform flawlessly. The most probable reason is that a minimum closer to the minimum discovered using the isotropic model would have been discovered with the optimization using the fluid model if more samples had been taken during the optimization, and that this might indicate that too few optimization samples were used to explore the four-dimensional input space fully.

That too few optimization samples have been used is further supported by the black cross of figure 4.10b, corresponding to the optimum of the fluid optimization, not even being close to the deepest blue regions. This is also confirmed when the approximation using the fluid model on the samples from the optimization using the isotropic model were plotted in the optimal plane discovered through the optimization using the fluid model. The insufficient amount of optimization samples is also indicated by the fact that 4.10b looks different from the highlighted rectangles of figure 4.9. That the four-dimensional data need to be projected on two-dimensional planes to be visualized obstructs the possibility of clearly seeing which regions have high resolution. In turn, this makes it difficult to determine where the GP approximation is reliable, which makes it difficult to draw accurate conclusions.

It is worth noting that a significantly large region of sufficiently low cost function values were discovered nevertheless, and that the difference in location of the minimum would not impact the safety or outcome of such a disruption scenario. In fact, in an experimental scenario with uncontrollable uncertainties, it is probably better to use a parameter combination in the safe region that is furthest away from the boundary of the safe region, in input space parameters, than the one that has the lowest cost function value – especially if the optimum is close to the boundary.

5

Plasma model analysis

Understanding the reasons for the differences between the fluid and isotropic models, which were discovered in chapter 4, is important for interpreting and using the results found. In this chapter, the physics differences of the models will be analyzed in the high magnetic perturbation case, $\delta B/B = 0.5\,\%$. Firstly, the runaway electron dynamics will be studied and the fluid model will be adapted to discern the reason for the difference in runaway current for the minimum discovered using the isotropic model. Secondly, the impact of the fast electron impact ionization will be explored for cases with high and low runaway currents. Finally, both of these cases will be explored at even higher physics fidelity using the superthermal model of DREAM. The differences between the isotropic and superthermal models will be examined and the energy and pitch angle distributions of the superthermal simulations will be studied.

5.1 Runaway electron dynamics

In the optimal case discovered using the isotropic model, the most significant difference compared to the fluid model occurred for the representative runaway current. With regards to the runaway dynamics, the two models treat the Dreicer and hot-tail runaway processes differently – they are automatically included in the isotropic model but are implemented as source terms to the runaway density n_{re} in the fluid model. It is therefore likely that one or both of these runaway generation processes can explain the differences, at least partly.

To explore this, simulations in the optimal case were performed with both models, as well as one simulation using the fluid model with the Dreicer runaway rate disabled and one with the hot-tail runaway generation disabled. In figure 5.1a, important plasma parameters are presented from all four simulations. Clearly, there are large differences between the isotropic and the standard fluid model, just as anticipated. Furthermore, removing the Dreicer rate does not have a significant impact on the simulations, but removing the hot-tail rate does.

There is a significant improvement on all parameters when disabling the hot-tail runaway rate, meaning that they behave similarly to their isotropic counterparts.

Table 5.1: Figures of merit for the found optimum for using isotropic and fluid models, including results from the fluid model without hot-tail and Dreicer runaway rates.

Model	$I_{\mathrm{re}}^{\mathrm{repr}}$ [kA]	$I_{\Omega}^{\mathrm{final}}$ [kA]	τ_{CQ} [ms]	η_{cond} [%]
Isotropic	3.5	0.0046	66	5.8
Fluid	2700	110	48	3.8
Fluid, w/o Dreicer	2700	100	48	3.8
Fluid, w/o Hot-tail	160	0.012	64	3.8

The only apparent difference between the isotropic model and the fluid model without hot-tail runaway rate is that with the latter the runaway current has a small, but visible, peak for a short period of time. It is thus evident that the fluid hot-tail runaway rate overestimates the amount of runaway electrons generated from the hot-tail mechanism.

This is further supported by comparing this optimal case with and without runaway generation mechanisms relevant for activated operation (not shown here), as the difference between the fluid and isotropic models were much smaller in the activated case. With the activated sources, the runaway current was significantly larger overall, which can be explained by the additions of the tritium and Compton runaway rates. With more runaway generation overall, the impact of the hot-tail runaway rate would be smaller and thus there would be smaller differences between the fluid and isotropic models.

Another interesting conclusion that can be drawn from the plot of the transported power in figure 5.1b is that the inconsistency in conducted heat load is a result of the isotropic model having a large amount of heat transport from the hot electron population during the TQ. The reason for this transport of hot electron energy is that there practically exists two temperatures with the isotropic model, one for the cold and one for the hot electron population, while there only exists one with the fluid model. In the fluid model, the electrons of the initial deuterium plasma, with initial temperature $T = 20\,\mathrm{keV}$, will instantaneously equilibrate with the cold injected material, causing the temperature of the cold electron to instantaneously drop. In the isotropic model however, the temperature of the hot electrons does not drop as rapidly since the electrons of the injected material will be part of the cold population. The hot electron temperature will equilibrate with the cold electron temperature on a longer time scale, during which there will be a larger transported heat from the hot electron population.

In table 5.1, the various figures of merit of the disruption are listed. As shown in the table, the fluid model without the hot-tail runaway rate still produces a runaway current larger than the maximum safe current. This suggests that there are still

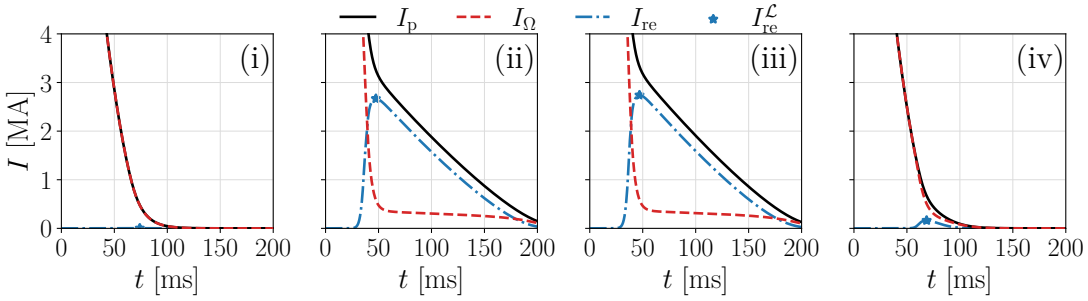

(a) Time evolutions of the currents in the optimal case using the (i) isotropic model, (ii) standard fluid model, as well as the fluid model with the (iii) Dreicer and (iv) hot-tail seed generation mechanisms disabled. The range is cropped at 4 MA to focus on the differences. A blue star marks the representative runaway current.

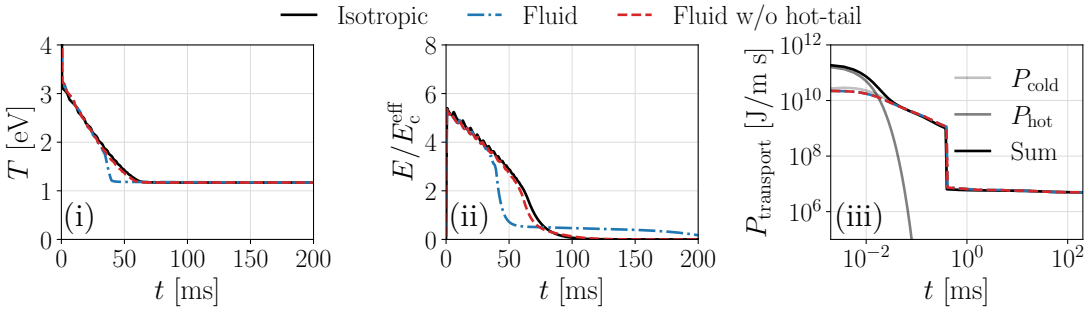

(b) Time evolutions of (i) the volume averaged plasma temperature, (ii) the volume averaged normalized electric field and (iii) the power transported through the edge (note the logarithmic time axis). With the isotropic model, the total transported power (black) is a sum of the contribution from the cold electrons (light grey) and the hot electrons (dark grey), while the total power in the fluid model only comes from the cold electrons.

Figure 5.1: Time evolutions of important plasma parameters in the optimal case, comparing the results between the isotropic and fluid models.

some impactful differences between the two models that we have not considered yet. However, it is worth pointing out that removing a mechanism such as the hot-tail runaway rate from the simulations can result in obscure consequences with regards to the plasma evolution, due to the complicated correlations of the dynamics.

Since the only runaway seed generation mechanism left when the hot-tail generation is removed is Dreicer, some aspect of the fluid model most likely causes the Dreicer runaway generation to be overestimated as well. Another explanation could be that the omission of the hot-tail generation rate might implicitly impact other plasma parameters such that the fluid Dreicer generation is increased. This would explain why removing the Dreicer generation from the fluid model did not seem to have an impact on the runaway current, suggesting that the Dreicer runaway rate is, for this case, negligible. However, when comparing the Dreicer generation rate between the fluid model with and without the hot-tail generation rate enabled it was found

that they were similar. This suggests that some aspect of the fluid model causes the Dreicer runaway rates to be overestimated as well, even though to a lesser degree than the overestimation of the fluid hot-tail runaway rate.

One could therefore suggest scale down the hot-tail and Dreicer generation rates of the fluid model. This would cause the fluid model to yield more similar results to the isotropic model for this simulation of the optimum discovered using the isotropic model – but it is definitely not guaranteed to yield more accurate results for all disruption scenarios. The fluid and isotropic models produce very similar results near the minimum discovered using the fluid model, demonstrating that the differences are parameter dependent, and as such cannot be bypassed by simply eliminating or down-scaling any runaway rate.

To conclude, the hot-tail runaway rate is substantially overestimated in the fluid model for some regions of the two dimensional density space explored in chapter 4. The Dreicer runaway rate in the fluid model is overestimated as well for the minimum discovered using the isotropic model, but not to the same degree. However, there might be a few additional inconsistencies between the fluid and isotropic models – concerning the electron heat and electric field feedbacks, as well as the fast electron impact ionization to be studied in the next section – but these are most likely not as significant as the overestimated hot-tail runaway rate.

5.2 Fast electron impact ionization

To study the effect of the fast electron impact ionization on the simulation results generated with the isotropic model, two cases with $\delta B/B = 0.5\,\%$ were considered – the optimum discovered using the isotropic model and a case when a significant runaway current was generated while the CQ time was within the safe interval $[50, 150]$ ms.

It was apparent that the fast electron impact ionization only impacts the simulations to a small degree based on the plasma parameters we have plotted – qualitatively, the only clear difference is in the runaway currents. Table 5.2 clarifies that only the currents exhibit non-negligible differences. For instance, the runaway current is almost a factor of 10 larger in the optimal case, but both of these runaway currents are very small compared to the safe operational limit of 150 kA. The same is true for the final ohmic current of the high runaway current case. Only the runaway current for the high runaway case is of a significant size while still exhibiting a notable, but not too large, difference of about 17 %.

The explanation is that in both of the studied cases the hot electron population is a relatively small fraction of the total electron population, and as such it does not enhance ionization much compared to the large cold electron population. In the high runaway current case, there might have been a significant difference if the runaway region was also kinetically resolved, as this population would be significantly larger

Table 5.2: Figures of merit for study of impact of kinetic ionization on simulations using the isotropic model.

Case	Kinetic ionization	$I_{\mathrm{re}}^{\mathrm{repr}}$ [kA]	$I_{\Omega}^{\mathrm{final}}$ [kA]	τ_{CQ} [ms]	η_{cond} [%]
Minimum	Enabled	3.5	0.0046	66	5.8
Minimum	Disabled	30	0.0060	66	5.6
High $I_{\mathrm{re}}^{\mathrm{repr}}$	Enabled	1800	35	63	5.0
High $I_{\mathrm{re}}^{\mathrm{repr}}$	Disabled	2100	52	60	5.0

(and contain much more kinetic energy) than the hot population, but this has not been studied.

As mentioned above, there is a higher runaway current when the fast electron impact ionization is disabled in the high runaway case. This can be explained by the avalanche generation, as it is inversely proportional to the number of free electrons in the plasma according to (3.5). When fast electron impact ionization is disabled, there should be a slightly smaller number of free electrons then in the case with the fast electron impact ionization enabled, and this will in turn result in a larger amount of runaway electrons being generated from avalanche generation.

5.3 Superthermal plasma model

In this thesis, one major theme has been the use of different plasma models to simulate the fusion plasmas and how they compare. The fluid and isotropic plasma models in DREAM have been utilized because of their relatively low computational cost, which is highly beneficial when making a large number of simulations. However, DREAM has several more plasma models of higher physics fidelity. The superthermal mode of DREAM is very similar to the isotropic model, except that it evolves also the pitch angle part of the hot electron distribution function.

Using the superthermal model as part of the optimizations would not be possible within the time constraints of this project, as the computational cost is much larger than that of the isotropic model – even when the radial resolution was lowered by more than 50 % one simulation took several hours. As the impact of plasma model is still very relevant for the purposes of this thesis, the superthermal model was used to simulate the same two cases of section 5.2 – the minimum discovered using the isotropic model case and a case of high runaway current and long CQ time. In order to reduce the computational cost of the superthermal simulations, a radial resolution of only seven points was used, in contrast to the 20 points that were used in the isotropic model. Seven radial points were also used in the isotropic simulations performed for this section, in order for a fair comparison of the models.

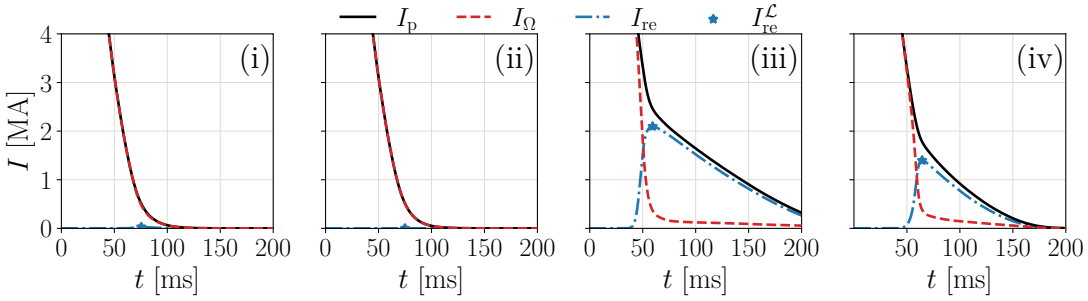

(a) Time evolutions of the currents in the optimal case/high runaway case with the (i)/(iii) isotropic and (ii)/(iv) superthermal model. A blue star marks the representative runaway current.

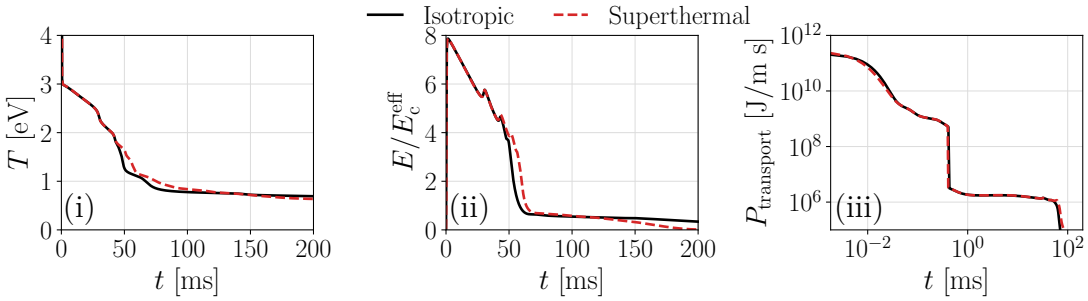

(b) Time evolutions of (i) the volume averaged plasma temperature, (ii) the volume averaged normalized electric field and (iii) the transported power for the high runaway current case.

Figure 5.2: Time evolutions of important plasma parameters for comparison between isotropic and superthermal models.

In figure 5.2 the same plasma parameters as in the previous sections are presented for the isotropic and superthermal simulations of the optimal and high runaway cases. In the optimal case, all plasma parameters evolved identically according to the plots (which is why they are not shown here), and this suggests that the isotropic model is accurate in this case. The plasma parameter evolutions are very similar in the high runaway current case as well, but there are small differences as shown in both figure 5.2a and 5.2b. The major difference is that the runaway current is 33 % lower with the superthermal model in the high runaway case.

In the electric field evolution, there are fewer, but larger, spikes, when comparing the isotropic cases in this section to those of the previous sections in this chapter. These spikes are thus most likely caused by a finite radial resolution, which suggests that even using only 20 points for the radial resolution is less than ideal.

An interesting aspect of the superthermal mode is the energy and pitch angle distribution of the hot electron population, which we can see in figure 5.3. In the optimal case discovered using the isotropic model, shown in figure 5.3a, there is a small but non-negligible hot electron population at $t = 75\,\text{ms}$ compared to the end

of the simulation, when all hot electrons have either cooled down and joined the cold population or become accelerated into the runaway population. The hot electron population is much larger in the high runaway current case, as shown in figure 5.3b.

Perhaps more unexpectedly, there is still a significant hot electron distribution at the end of the simulation in the high runaway case. After 200 ms, it could be expected that the electron distribution is distinctly divided into a bulk (cold) and a runaway population with very few electrons in between. This is because of the flux of hot electrons into the cold population through the lower momentum boundary for the hot electron population, and the runaway generation through momentum flux through the critical momentum boundary. Recall that the lower momentum boundary for the hot electron population is at $p = 0$ since the cold electron population is not resolved in momentum space, as shown in figure 3.1b. The reason that there is still a significant hot electron population might be that the electric field is non-negligible until about 180 ms, which could reduce the flux of electrons from the hot to the cold populations, however, this should increase the runaway generation as well.

Furthermore, the distributions are shifted towards lower pitch angles, which is expected since the electrons have smaller average pitch angles because of the electric field acceleration. However, since the hot electron population is noticeably beamed in the forward direction, it is not approximately isotropic. This shift towards lower pitch angles could offer an explanation as to why there is a larger runaway current in the isotropic simulations. If the assumption about collisional isotropization is wrong, the isotropic model can be overestimating the runaway electron densities for higher pitch angles, resulting in a higher runaway current than when the pitch angle dependence is resolved.

In this thesis, the hot electron current is counted towards the runaway current, but the hot electrons are not part of the runaway electron population. This is significant because the runaway electrons can only become part of the thermal electron population through reverse avalanche, while there is a constant flux of electrons into the cold electron population from the hot population. As shown in figure 5.3a, there is a significant hot electron population in the simulation with the superthermal model at $t = 64$ ms. However, most of the hot electrons are located close to the cold electron population near $p = 0$, while few electrons sit near the runaway boundary, indicating that there is no significant runaway generation due to momentum flux from the hot population. In the fluid model, where there is no hot electron population, these electrons would be already defined as runaways at $t = 64$ ms. Because these electrons are part of the runaway population in the fluid model, they will further power the runaway electron generation through avalanche. Since the only losses from the runaway population are from transport and reverse avalanche, the runaway electrons can only return to the cold electron population under certain circumstances, while there is a constant flux of hot electrons to the cold population at $p = 0$. Thus, the inability of the fluid model to resolve the momentum-space dynamics of the hot-tail generation and subsequent thermalization and acceleration of the superthermal population impacts the fluid results significantly. As such, the

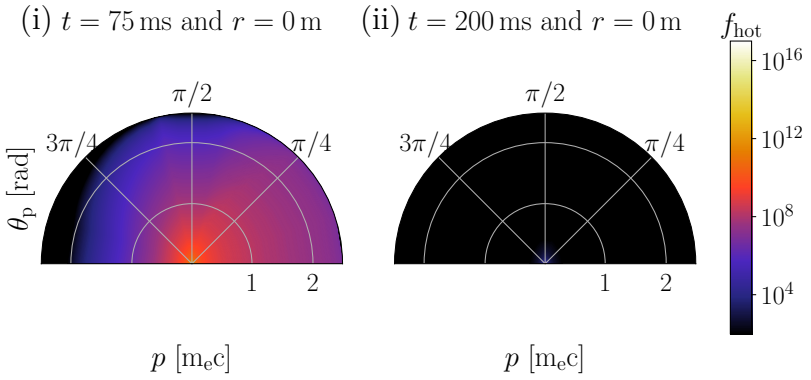

(a) Energy and pitch angle distributions from the superthermal model in the minimum discovered using the isotropic model.

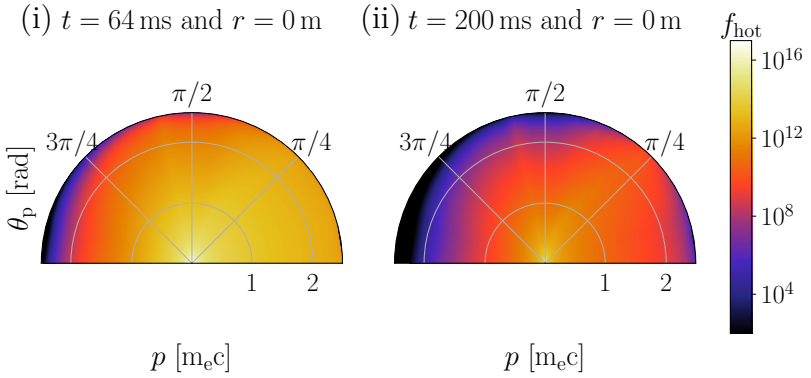

(b) Energy and pitch angle distributions from the superthermal model in the case with high runaway current and fairly long CQ time.

Figure 5.3: Hot electron distribution as a function of the momentum and pitch angle, which has been obtained from superthermal simulations. The distribution functions at $r = 0$ are presented, and (i) at the time when the representative runaway current is taken and (ii) the final time step of the simulation.

energy and pitch angle distributions from the superthermal simulations further supports the conclusion of section 5.1 about the hot-tail runaway rate being the major reason for the inconsistencies between the fluid and isotropic models.

6

Concluding remarks

Realizing fusion as an energy source would solve many of the environmental problems our society faces today. One possible path for achieving fusion is through the tokamak concept, but it still faces many challenges before it can become a viable source of energy. One of the major challenges is the occurrence of disruptions that can cause mechanical stresses and heat related damages to the machines. Disruption mitigation is thus an important subject of research. In this thesis, massive material injection of deuterium and neon for the mitigation of disruptions have been studied using both fluid and kinetic plasma models. In chapter 4, optimization of the injected densities was performed, both with uniform and radially varying density distributions. A cost function was systematically developed, combining the relevant disruption parameters – runaway current, final ohmic current, current quench time and conducted heat loads – in order to reliably quantify the disruption evolution and maximize the amount of information gain. Bayesian optimization was used to find the optimum of this cost function, as well as to explore regions of the parameter space exhibiting tolerable disruption outcomes. The optimization was initially done in a larger region using the computationally efficient fluid plasma model, and then in a smaller, promising region utilizing the higher physics fidelity of the kinetic plasma model. Furthermore, the differences between the models in the kinetic optimization domain were analyzed. Finally, specific disruption scenarios were studied more thoroughly in chapter 5 to determine the reasons for the differences between the fluid and kinetic models, as well as determining the impact of fast electron impact ionization and studying the energy and pitch angle distributions obtained from the kinetic models. Here we will summarize the conclusions related to the objectives of chapter 1, followed by a review of possible future research directions.

6.1 Conclusions

The specific conclusions from this thesis are now presented together with the corresponding aims specified in chapter 1.

To systematically develop an informative cost function to reliably quantify disruption evolutions of massive material injection scenarios following good optimization practises.

The cost function was designed to yield values lower than one for disruption scenarios when all disruption parameters were within their safe operational limits, and higher than one otherwise. Outside the safe operational limits, all disruption parameters contributed linearly to the cost function. Inside the limits, the disruption parameters were processed differently, to reflect the order of importance of the various figures of merit. Namely, the cost function favoured lower currents over lower heat loads, which in turn would be more favourable than current quench times close to the middle of its safe operational interval.

To find safe operational regions and the optimum of this cost function in the input space spanned by the injected material densities, as well as their radial variation, using fluid plasma models.

For uniformly distributed densities, safe regions of the cost function were discovered for both higher and lower magnetic perturbation levels mostly centred around the discovered optimum. For the lower perturbation case, the safe region was fairly large, but this was not true for the higher perturbation case. The optima were found around $n_{\mathrm{D}} = 10^{22}\,\mathrm{m}^{-3}$ and $n_{\mathrm{Ne}} = 10^{19}\,\mathrm{m}^{-3}$, but the optimum moved to a lower neon density for the case with the lower magnetic perturbation level. When optimization was performed with radially varying injected densities, a radial profile of the injected deuterium that is slightly peaked close to the plasma centre was favoured. The neon density was less sensitive to the radial distribution, but the optimum corresponded to edge localized neon densities.

To explore the discovered safe operational regions using both fluid and kinetic plasma models, and perform a qualitative analysis of the differences between the two models.

For radially uniform densities, the kinetic model yielded more optimistic results – yielding lower values of the cost function for the lower magnetic perturbation, and a larger region of safe values for the higher perturbation. The optimum is found at higher neon densities and slightly lower deuterium densities using the kinetic model, compared to the fluid model. Of the cost function components, the runaway current was the parameter that attained the most disparate values between the fluid and kinetic models. This was explained by the hot-tail runaway generation being overestimated in the fluid model. The only parameter, for which the kinetic model yielded more pessimistic results, was the conducted heat load, however the differences between the models was one order of magnitude smaller than the limit of safe operation. For radially varying densities, the differences were less pronounced. Fairly large regions of safe values were discovered for all four-dimensional optimizations, and all optima were orders of magnitude lower than one. However, it was difficult to draw reliable conclusions regarding the cost function behaviour in the four dimensional parameter space, as the results indicated that too few optimization points had been used, resulting in somewhat unreliable GP approximation.

To study the electron energy and pitch-angle distribution of the discovered optimum and a runaway-dominated case using the kinetic model.

A non-negligible amount of mildly superthermal electrons were present during the simulations in the optimal case, but they had disappeared (thermalized) by the end of the simulation. For the runaway dominated case, there were large amounts of runaways both in the middle and at the end of the simulation. In both cases the distributions were shifted in the momentum and pitch angle space towards lower pitch angles.

To examine the significance of fast electron impact ionization in a high and low runaway case.

The fast electron impact ionization had negligible effect on both the optimal and runaway dominated cases. We suspect this tentative conclusion could be changed in simulations that kinetically resolve also the runaway region, but that is outside the scope of this thesis.

6.2 Outlook

There are numerous challenges connected with tokamak disruptions and runaway electrons, and this thesis only touched on a fraction of these. Nevertheless, proceeding from this thesis, there are many possible avenues of research investigations that can be explored.

One of the more straightforward ideas would be to do a sensitivity scan around the optima discovered using the isotropic model – it would be advantageous to know how large perturbations of the input parameters would be needed to have a large effect on any of the cost function components. This would certainly be needed if the results of this thesis were to be directly implemented for an ITER-like scenario, as it would inform about the reliability and stability of the optima.

One natural extension of this thesis would be to add even more layers of higher physics fidelity to our optimization algorithm. Especially the runaway current was decreased when going from the fluid to the isotropic model, and the high runaway scenario indicated that the same could be true when going from the isotropic model to the superthermal model. However, there is a substantial increase in computational cost when using plasma models of higher physics fidelity – and thus, this might not be practical. However higher fidelity checks of a finite number of interesting cases would definitely be possible and interesting. Regarding optimizations, this would only be relevant if the computational cost could be lowered. The total computational time might be lowered by lowering the resolution or computationally using practices such as parallelization.

Another approach related to plasma models would be to optimize an activated disruption scenario using the kinetic models of this thesis. Only non-activated discharges were thoroughly studied in this thesis, while activated cases are highly relevant with regards to fusion development. In Ref. [12], a similar optimization was performed, but using only the fluid model, and no safe scenarios were discovered.

Exploring an activated disruption scenario more thoroughly, also using kinetic models and possibly with other MMI compositions, would be very relevant and it would be very advantageous for the future of ITER-like tokamaks if safe scenarios could be constructed. As of now, DREAM only has the ability of considering fluid-generated tritium and Compton electrons. Since the momentum-space dynamics of these electrons could potentially play a role, it could be interesting to treat these source terms kinetically as well.

We have only studied a very simple model of MMI in this thesis, where the injected material is instantaneously distributed according to radial distributions of very strict forms and only two injected materials have been used. Further optimization studies could be made with other injection materials, such as argon, other optimization parameters, such as when the different materials are injected, as well as changing MMI model itself, such as how the injected material is distributed in the plasma over time.

Finally, to fully utilize the results of the radially varying optimization, an in-depth analysis of MMI schemes would have to be performed to determine if and how it is possible to achieve the optimal radial density distributions discovered during optimization. One approach, combining the injection of a deuterium pellet with small amounts of neon impurities followed by massive gas injection of neon, was suggested in chapter 4, but it is uncertain whether it would work. Furthermore there are many possibilities to explore, since there are many realizations of MMI, for example various forms of pellet injection and massive gas injection, and even more ways of designing and combining these MMI methods.

Bibliography

[1] *Fuelling the fusion reaction*. [Online]. Available: http://www.iter.org/sci/fusionfuels (visited on 03/03/2023).

[2] *Advantages of fusion*. [Online]. Available: http://www.iter.org/sci/fusion (visited on 02/10/2023).

[3] *SPARC*. [Online]. Available: https://www.psfc.mit.edu/sparc (visited on 03/03/2023).

[4] *The ITER tokamak*. [Online]. Available: https://www.iter.org/mach (visited on 03/03/2023).

[5] E. M. Hollmann *et al.*, "Status of research toward the ITER disruption mitigation system," *Physics of Plasmas*, vol. 22, no. 2, p. 021 802, Nov. 2015. DOI: 10.1063/1.4901251.

[6] B. N. Breizman, P. Aleynikov, E. M. Hollmann, and M. Lehnen, "Physics of runaway electrons in tokamaks," *Nuclear Fusion*, vol. 59, no. 8, p. 083 001, Jun. 2019. DOI: 10.1088/1741-4326/ab1822.

[7] C. Reux *et al.*, "Runaway electron beam generation and mitigation during disruptions at JET-ILW," *Nuclear Fusion*, vol. 55, no. 9, p. 093 013, Aug. 2015. DOI: 10.1088/0029-5515/55/9/093013.

[8] M. Lehnen *et al.*, "Disruptions in ITER and strategies for their control and mitigation," *Journal of Nuclear Materials*, vol. 463, pp. 39–48, Aug. 2015, ISSN: 0022-3115. DOI: 10.1016/j.jnucmat.2014.10.075.

[9] F. Chen, *Introduction to Plasma Physics and Controlled Fusion*, 3rd ed. Switzerland: Springer, 2018, pp. 19–24, 51, 169, 211–212, 217, 222.

[10] R. Garnett, *Bayesian Optimization*. Padstow, UK: Cambridge University Press, 2023, pp. 6–11, 16–17, 51–52. DOI: 10.1017/9781108348973.

[11] H. Bergström and P. Halldestam, "Optimization of tokamak disruption scenarios," *Chalmers Open Digital Repository*, Jun. 2022. [Online]. Available: https://hdl.handle.net/20.500.12380/305257 (visited on 03/03/2023).

[12] I. Pusztai, I. Ekmark, H. Bergström, P. Halldestam, P. Jansson, M. Hoppe, O. Vallhagen, and T. Fülöp, "Bayesian optimization of massive material injection for disruption mitigation in tokamaks," *Journal of Plasma Physics*, vol. 89, no. 2, p. 905 890 204, 2023. DOI: 10.1017/S0022377823000193.

[13] D. Anderson, M. Lisak, P. Johannisson, and M. Marklund, *Basic Plasma Physics*. 2003, pp. 25, 37–38, 58.

[14] P. Helander and D. J. Sigmar, *Collisional Transport in Magnetized Plasmas*. United States: Cambridge University Press, 2005, pp. 22–29.

[15] J. Wesson, *Tokamaks*, 4th ed. United States: Oxford University Press, 2011, pp. 13, 102, 104, 120, 115, 378–380.

[16] R. L. Miller, M. S. Chu, J. M. Greene, Y. R. Lin-Liu, and R. E. Waltz, "Noncircular, finite aspect ratio, local equilibrium model," *Physics of Plasmas*, vol. 5, no. 4, pp. 973–978, Oct. 1998. DOI: 10.1063/1.872666.

[17] M. Hoppe, "Runaway-electron model development and validation in tokamaks," Ph.D. dissertation, Chalmers University of Technology, 2022.

[18] *Eddy current | electronics*, Feb. 2023. [Online]. Available: https://www.britannica.com/science/eddy-current (visited on 03/15/2023).

[19] E. Nardon, D. Hu, M. Hoelzl, D. Bonfiglio, and the JOREK team, "Fast plasma dilution in ITER with pure deuterium shattered pellet injection," *Nuclear Fusion*, vol. 60, no. 12, p. 126 040, Oct. 2020. DOI: 10.1088/1741-4326/abb749.

[20] D. Hu, E. Nardon, M. Hoelzl, F. Wieschollek, M. Lehnen, G. Huijsmans, D. C. van Vugt, S.-H. Kim, J. contributors, and J. team, "Radiation asymmetry and MHD destabilization during the thermal quench after impurity shattered pellet injection," *Nuclear Fusion*, vol. 61, no. 2, p. 026 015, Jan. 2021. DOI: 10.1088/1741-4326/abcbcb.

[21] A. Matsuyama, "Neutral gas and plasma shielding (NGPS) model and cross-field motion of ablated material for hydrogen–neon mixed pellet injection," *Physics of Plasmas*, vol. 29, no. 4, p. 042 501, Apr. 2022. DOI: 10.1063/5.0084586.

[22] G. Papp, M. Drevlak, T. Fülöp, P. Helander, and G. I. Pokol, "Runaway electron losses caused by resonant magnetic perturbations in ITER," *Plasma Physics and Controlled Fusion*, vol. 53, no. 9, p. 095 004, Jul. 2011. DOI: 10.1088/0741-3335/53/9/095004.

[23] P. Svensson, O. Embreus, S. L. Newton, K. Särkimäki, O. Vallhagen, and T. Fülöp, "Effects of magnetic perturbations and radiation on the runaway avalanche," *Journal of Plasma Physics*, vol. 87, no. 2, p. 905 870 207, 2021. DOI: 10.1017/S0022377820001592.

[24] R. Tinguely *et al.*, "Modeling the complete prevention of disruption-generated runaway electron beam formation with a passive 3D coil in SPARC," *Nuclear Fusion*, vol. 61, p. 124 003, 12 2021. DOI: 10.1088/1741-4326/ac31d7.

[25] V. Izzo, I. Pusztai, K. Särkimäki, A. Sundström, D. Garnier, D. Weisberg, R. Tinguely, C. Paz-Soldan, R. Granetz, and R. Sweeney, "Runaway electron deconfinement in SPARC and DIII-D by a passive 3D coil," *Nuclear Fusion*, vol. 62, p. 096 029, 9 Aug. 2022. DOI: 10.1088/1741-4326/ac83d8.

[26] J. Connor and R. Hastie, "Relativistic limitations on runaway electrons," *Nuclear Fusion*, vol. 15, no. 3, p. 415, Jun. 1975. DOI: 10.1088/0029-5515/15/3/007.

[27] P. Helander, H. Smith, T. Fülöp, and L.-G. Eriksson, "Electron kinetics in a cooling plasma," *Physics of Plasmas*, vol. 11, no. 12, pp. 5704–5709, Nov. 2004. DOI: 10.1063/1.1812759.

[28] M. Rosenbluth and S. Putvinski, "Theory for avalanche of runaway electrons in tokamaks," *Nuclear Fusion*, vol. 37, no. 10, p. 1355, Oct. 1997. DOI: 10.1088/0029-5515/37/10/I03.

[29] J. Martín-Solís, A. Loarte, and M. Lehnen, "Formation and termination of runaway beams in ITER disruptions," *Nuclear Fusion*, vol. 57, no. 6, p. 066 025, Apr. 2017. DOI: 10.1088/1741-4326/aa6939.

[30] J. Joyce, *Bayes' theorem*, Sep. 2003. [Online]. Available: https://plato.stanford.edu/entries/bayes-theorem/ (visited on 03/16/2023).

[31] G. De Ath, R. M. Everson, A. A. M. Rahat, and J. E. Fieldsend, "Greed is good: Exploration and exploitation trade-offs in Bayesian optimisation," *ACM Transactions on Evolutionary Learning and Optimization*, vol. 1, no. 1, pp. 1–22, Apr. 2021. DOI: 10.1145/3425501.

[32] M. Hoppe, O. Embreus, and T. Fülöp, "DREAM: A fluid-kinetic framework for tokamak disruption runaway electron simulations," *Computer Physics Communications*, vol. 268, p. 108 098, Jul. 2021, ISSN: 0010-4655. DOI: 10.1016/j.cpc.2021.108098.

[33] *OPEN-ADAS*. [Online]. Available: https://open.adas.ac.uk/ (visited on 03/10/2023).

[34] D. Reiter, "The data file AMJUEL: Additional atomic and molecular data for EIRENE," Jan. 2000.

[35] N. A. Garland, H.-K. Chung, C. J. Fontes, M. C. Zammit, J. Colgan, T. Elder, C. J. McDevitt, T. M. Wildey, and X.-Z. Tang, "Impact of a minority relativistic electron tail interacting with a thermal plasma containing high-atomic-number impurities," *Physics of Plasmas*, vol. 27, no. 4, p. 040 702, Apr. 2020. DOI: 10.1063/5.0003638.

[36] M. Hoppe, *Radialgrid — DREAM documentation*, 2021. [Online]. Available: https://ft.nephy.chalmers.se/dream/frontend/api/Settings/timestep.html (visited on 02/24/2023).

[37] I. Pusztai, O. Vallhagen, L. Hanebring, and T. Fülöp, "Report for task D3.1: Simulations of SPI mitigated ITER disruptions," ITER Organisation, Tech. Rep. IO/IA/21/4300002402, 2022.

[38] I. Pusztai, M. Hoppe, and O. Vallhagen, "Runaway dynamics in tokamak disruptions with current relaxation," *Journal of Plasma Physics*, vol. 88, no. 4, p. 905 880 409, Aug. 2022. DOI: 10.1017/S0022377822000733.

[39] O. Vallhagen, I. Pusztai, M. Hoppe, S. Newton, and T. Fülöp, "Effect of two-stage shattered pellet injection on tokamak disruptions," *Nuclear Fusion*, vol. 62, no. 11, p. 112 004, Sep. 2022. DOI: 10.1088/1741-4326/ac667e.

[40] O. Vallhagen, O. Embreus, I. Pusztai, L. Hesslow, and T. Fülöp, "Runaway dynamics in the DT phase of ITER operations in the presence of massive material injection," *Journal of Plasma Physics*, vol. 86, no. 4, p. 475 860 401, Aug. 2020. DOI: 10.1017/S0022377820000859.

[41] L. Hesslow, O. Embréus, O. Vallhagen, and T. Fülöp, "Influence of massive material injection on avalanche runaway generation during tokamak disruptions," *Nuclear Fusion*, vol. 59, no. 8, p. 084 004, Jun. 2019. DOI: 10.1088/1741-4326/ab26c2.

[42] L. Hesslow, L. Unnerfelt, O. Vallhagen, O. Embreus, M. Hoppe, G. Papp, and T. Fülöp, "Evaluation of the Dreicer runaway generation rate in the presence of high-Z impurities using a neural network," *Journal of Plasma Physics*, vol. 85, no. 6, p. 475 850 601, Dec. 2019. DOI: 10.1017/S0022377819000874.

[43] I. Svenningsson, O. Embreus, M. Hoppe, S. L. Newton, and T. Fülöp, "Hot-tail runaway seed landscape during the thermal quench in tokamaks," *Phys. Rev. Lett.*, vol. 127, p. 035 001, 3 Jul. 2021. DOI: 10.1103/PhysRevLett.127.035001.

[44] A. B. Rechester and M. N. Rosenbluth, "Electron heat transport in a tokamak with destroyed magnetic surfaces," *Phys. Rev. Lett.*, vol. 40, pp. 38–41, 1 Jan. 1978. DOI: 10.1103/PhysRevLett.40.38.

[45] L. Hesslow, O. Embréus, A. Stahl, T. C. DuBois, G. Papp, S. L. Newton, and T. Fülöp, "Effect of partially screened nuclei on fast-electron dynamics," *Physical Review Letters*, vol. 118, p. 255 001, 25 Jun. 2017. DOI: 10.1103/PhysRevLett.118.255001.

[46] C. J. McDevitt and X.-Z. Tang, "Runaway electron generation in axisymmetric tokamak geometry," *Europhysics Letters*, vol. 127, no. 4, p. 45 001, Sep. 2019. DOI: 10.1209/0295-5075/127/45001.

[47] O. Sauter, C. Angioni, and Y. R. Lin-Liu, "Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime," *Physics of Plasmas*, vol. 6, no. 7, pp. 2834–2839, Jun. 1999. DOI: 10.1063/1.873240.

[48] M. Lehnen and the ITER DMS task force, *The ITER disruption mitigation system - design progress and design validation*, Presented at Theory and Simulation of Disruptions Workshop, PPPL, Jul. 2021.

[49] F. Nogueira, *Bayesian optimization: Open source constrained global optimization tool for Python*, 2014–. [Online]. Available: https://github.com/fmfn/BayesianOptimization.

[50] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian optimization of machine learning algorithms," vol. 25, 2012. [Online]. Available: https://proceedings.neurips.cc/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf (visited on 02/22/2023).

DEPARTMENT OF PHYSICS
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY