

Fluid and kinetic studies of tokamak disruptions using Bayesian optimization

I. Ekmark^{1,†}, M. Hoppe², T. Fülöp¹, P. Jansson³, L. Antonsson¹,
O. Vallhagen¹ and I. Pusztai¹

¹Department of Physics, Chalmers University of Technology, Göteborg SE-41296, Sweden

²Department of Electrical Engineering, KTH Royal Institute of Technology, Stockholm SE-11428, Sweden

³Department of Computer Science and Engineering, Chalmers University of Technology and University of Gothenburg, Göteborg SE-41296, Sweden

(Received 8 February 2024; revised 19 April 2024; accepted 22 April 2024)

When simulating runaway electron dynamics in tokamak disruptions, fluid models with lower numerical cost are often preferred to more accurate kinetic models. The aim of this work is to compare fluid and kinetic simulations of a large variety of different disruption scenarios in ITER. We consider both non-activated and activated scenarios; for the latter, we derive and implement kinetic sources for the Compton scattering and tritium beta decay runaway electron generation mechanisms in our simulation tool DREAM (Hoppe *et al.*, *Comput. Phys. Commun.*, vol. 268, 2021, 108098). To achieve a diverse set of disruption scenarios, Bayesian optimization is used to explore a range of massive material injection densities for deuterium and neon. The cost function is designed to distinguish between successful and unsuccessful disruption mitigation based on the runaway current, current quench time and transported fraction of the heat loss. In the non-activated scenarios, we find that fluid and kinetic disruption simulations can have significantly different runaway electron dynamics, due to an overestimation of the runaway seed by the fluid model. The primary cause of this is that the fluid hot-tail generation model neglects superthermal electron transport losses during the thermal quench. In the activated scenarios, the fluid and kinetic models give similar predictions, which can be explained by the significant influence of the activated sources on the runaway dynamics and the seed.

Key words: fusion plasma, runaway electrons, plasma simulation

1. Introduction

Plasma-terminating disruptions are one of the main challenges facing tokamak fusion reactors. Disruptions are off-normal events, leading to a sudden cooling of the plasma – a thermal quench (TQ) – associated with an increase in the plasma resistivity, causing the plasma current to decay over the longer current quench (CQ). The toroidal current cannot change significantly on the short TQ time scale, and therefore an inductive

† Email address for correspondence: ida.ekmark@chalmers.se

electric field is produced, which can accelerate electrons to relativistic energies (Helander, Eriksson & Andersson 2002; Breizman *et al.* 2019). There are several problems associated with disruptions, including heat loads on the plasma facing components and mechanical stresses due to electromagnetic forces (Hollmann *et al.* 2015). One of the most significant problems, however, is the generation of highly energetic runaway electron (RE) beams.

Runaway generation is expected to be a major problem in the next generation of machines with high plasma currents, such as ITER (Hender *et al.* 2007). The reason is that avalanche multiplication of seed REs is exponentially sensitive to the initial plasma current (Rosenbluth & Putvinski 1997). In addition, during deuterium-tritium operation, REs can also be generated by Compton scattering of γ photons and tritium beta decay. These activated generation sources will represent two irreducible sources of energetic electrons, and are therefore expected to be the dominant RE seed generation mechanisms in future reactor-scale tokamaks when hot-tail and Dreicer generation have been successfully mitigated. An uncontrolled, localized loss of a relativistic runaway beam may result in damage to the first wall. One of the most prominent mitigation methods is massive material injection (MMI) of a combination of radiating impurities (e.g. neon) and hydrogen isotopes, either in gaseous or solid state. Shattered pellet injection is the reference concept for the disruption mitigation system in ITER. However, as of yet, the problem of designing a successful disruption mitigation strategy is still unsolved.

Several sophisticated simulation codes have been developed to study disruptions. One such code is DREAM (Hoppe, Embréus & Fülöp 2021a) that is primarily developed for the study of REs. DREAM has several options for simulating electron dynamics – ranging from fluid to fully kinetic models. DREAM has already been used to study several aspects of tokamak disruption mitigation – including finding parameter spaces without REs for SPARC (Tinguely *et al.* 2021; Izzo *et al.* 2022; Tinguely *et al.* 2023a), STEP (Berger *et al.* 2022) and ITER (Pusztai *et al.* 2023). These studies used RE fluid models, i.e. the momentum-space runaway dynamics is replaced by formulae for the steady-state runaway generation rates as functions of the background parameters.

Kinetic plasma models are physically more rigorous and can, in certain scenarios, yield significantly different results compared with those obtained from fluid models. In kinetic models, the Dreicer and hot-tail RE generation mechanisms are naturally included because the dynamics in the electron velocity space is modelled by the Fokker–Planck collision operator. In certain cases, there may be substantial discrepancies between the kinetic and fluid generation rates. For example, in DREAM, the hot-tail generation mechanism is implemented as a generation rate based on a simplified model (Smith & Verwichte 2008), which is derived under the assumption that there are no magnetic perturbations. This model can significantly overestimate the RE generation rate when magnetic perturbations are present.

Pusztai *et al.* (2023) used the fluid RE model in DREAM to explore the injected density space of MMI, with Bayesian optimization, to investigate scenarios for disruption mitigation in ITER. However, such a comprehensive exploration of the parameter space has not been performed using kinetic modelling. Even though isolated studies using kinetic models have been done (Svenningsson *et al.* 2021; Vallhagen *et al.* 2022), a thorough exploration of the MMI injected density space would not be feasible for a fully kinetic model, due to the many simulations needed combined with the high computational cost of fully kinetic simulations. Using an isotropic kinetic model for such exploration is however feasible, i.e. the model which evolves a distribution function which is analytically averaged over pitch angle. Such a procedure is justified for the mildly superthermal region where the complicated hot-tail dynamics takes place, while the relativistic electrons can again be

treated in a simplified manner. The isotropic kinetic model in DREAM is several orders of magnitude faster than the corresponding fully kinetic model (Hoppe *et al.* 2021a).

The aim of this article is to compare fluid and kinetic disruption simulations for an ITER-like disruption scenario. We explore the injected density space of MMI, using Bayesian optimization for sample selection. We start by deriving kinetic expressions for the RE generation sources from tritium beta decay and Compton scattering (§ 2). Then we proceed by constructing an informative cost function (§ 3.2) for efficient sample selection during the Bayesian optimization (§ 3).

Our results show that in the non-activated scenarios, the fluid and kinetic models yield significantly different optima (§ 4). This can be explained mainly by the generation rate of the runaway seed being overestimated when using the fluid model. In the activated scenarios, the fluid and kinetic models yield more similar results. In these cases, the Compton generation mechanism is dominant, which is well represented in the fluid model.

2. Kinetic runaway electron generation sources

Runaway electron generation is an inherently kinetic process, as it depends on details in the velocity distribution function that are determined by a balance between collisional friction, electric-field acceleration and radiation processes. With the kinetic models in DREAM, the generation of REs from the Dreicer and hot-tail mechanisms is included in the dynamics modelled by the Fokker–Planck collision operator, while REs generated via avalanche multiplication are accounted for using a semi-analytical growth rate by Hesslow *et al.* (2019a) which has been benchmarked against kinetic simulations.

The activated generation processes, Compton scattering and tritium beta decay, provide energetic electrons that are continually generated over a range of electron velocities. Under the assumption that these generation mechanisms are isotropic in velocity space, kinetic sources can be derived using the same reasoning as for their fluid counterparts (Martín-Solís, Loarte & Lehnen 2017). For tritium beta decay, this assumption is justified since the emission angle of electrons during this process is random and uniformly distributed. For Compton scattering, gamma photons will be emitted from the walls and enter the plasma from all sides in a nearly isotropic manner, even traversing walls, because the plasma is optically thin to the photons. Therefore, the resulting electron distribution can be assumed to be approximately uncorrelated with the pitch angle.

2.1. Tritium beta decay

The kinetic source rate due to tritium decay rate $(df/dt)_T$ can be derived using the expression for the energy spectrum of electrons emitted through beta decay (Krane 1988)

$$f_\beta(W) \propto F(p, 2)p\mathcal{W}(W_{\max} - W)^2 \quad \text{for } W \leq W_{\max}, \text{ and } f_\beta(W) = 0 \text{ above } W_{\max}. \quad (2.1)$$

Here $\mathcal{W} = m_e c^2 \gamma$ is the total energy, $W = m_e c^2 (\gamma - 1)$ is the kinetic energy, $\gamma = \sqrt{p^2 + 1}$ is the Lorentz factor and p is the momentum normalized to $m_e c$. Since the maximum energy of the emitted electrons is $W_{\max} = 18.6 \text{ keV} \ll m_e c^2$, we may take the non-relativistic limit of the Fermi function (Fermi 1934),

$$F(p, Z_f) = \frac{2\pi\alpha Z_f/\beta}{1 - \exp(-2\pi\alpha Z_f/\beta)}, \quad (2.2)$$

with the normalized speed $\beta = p/\gamma$, the charge of the final state nucleus Z_f (for tritium decay $Z_f = 2$) and the fine structure constant $\alpha \approx 1/137$.

Since the source is isotropic, we may write

$$4\pi p^2 dp \left(\frac{df}{dt} \right)_T = dW f_\beta(W). \quad (2.3)$$

Using $dW/dp = m_e c^2 \beta$, we find

$$\left(\frac{df}{dt} \right)_T = \frac{1}{4\pi p^2} f_\beta(W) m_e c^2 \beta \propto \frac{1}{p^2} \frac{p\gamma(\gamma_{\max} - \gamma)^2}{1 - \exp(-4\pi\alpha/\beta)}. \quad (2.4)$$

The requirement that the source integrated over the entire momentum space yields the tritium decay rate, $(\ln 2) n_T / \tau_T$ – where $\tau_T \approx 4500$ days is the half-life for tritium – determines the proportionality factor in (2.4) such that

$$\left(\frac{df}{dt} \right)_T \approx C \frac{\ln 2}{4\pi} \frac{n_T}{\tau_T} \frac{1}{p^2} \frac{p\gamma(\gamma_{\max} - \gamma)^2}{1 - \exp(-4\pi\alpha/\beta)}, \quad (2.5)$$

with the constant

$$C = 1 \int_0^{p_{\max}} \frac{1}{p^2} \frac{p\gamma(\gamma_{\max} - \gamma)^2}{1 - \exp(-4\pi\alpha/\beta)} p^2 dp \approx 31\,800, \quad (2.6)$$

where p_{\max} is the momentum corresponding to W_{\max} .

2.2. Compton scattering rate

The kinetic Compton scattering rate can be calculated using the Compton cross-section and the gamma energy flux spectrum emitted by a given wall material. When using the kinematic relation between the photon energy W_γ , and the energy W and deflection angle θ of the scattered electron, we can write

$$\cos \theta = 1 - \frac{m_e c^2}{W_\gamma} \frac{W}{W'_\gamma}, \quad (2.7)$$

where $W'_\gamma = W_\gamma - W$ is the scattered photon energy. The Compton scattering process is described by the Klein–Nishina differential cross-section

$$\frac{d\sigma}{d\Omega} = \frac{r_e^2}{2} \frac{W_\gamma^2}{W_\gamma^2} \left[\frac{W_\gamma}{W'_\gamma} + \frac{W'_\gamma}{W_\gamma} - \sin^2(\theta) \right], \quad (2.8)$$

with $r_e = e^2/4\pi\epsilon_0 m_e c^2$ being the classical electron radius.

Thus, in a plasma with the total electron density $n_{e,\text{tot}}$, i.e. the electron population of both free and bound electrons, and isotropic photon distribution, we can write

$$4\pi p^2 dp \left(\frac{df}{dt} \right)_C = n_{e,\text{tot}} \int_{W_{\gamma 0}}^{\infty} dW_\gamma \left[\Gamma_\gamma(W_\gamma) \left(2\pi \sin \theta \frac{d\sigma}{d\Omega} \right) d\theta(W, W_\gamma) \right], \quad (2.9)$$

where $\Gamma_\gamma(W_\gamma)$ is the gamma energy flux spectrum. The lower integration limit $W_{\gamma 0}$ can be derived from (2.7) as $W_{\gamma 0} = (p + \gamma - 1)/2$.

Martín-Solís *et al.* (2017) estimated the photon flux energy spectrum in ITER by $\Gamma_\gamma(W_\gamma) = \Gamma_0 \exp[-\exp(-z) - z + 1]$, where $z = [\ln(W_\gamma/10^6) + 1.2]/0.8$ and

$$\Gamma_0 = \Gamma_{\text{flux}} \int_0^\infty \exp[-\exp(-z) - z + 1] dW_\gamma = \Gamma_{\text{flux}}/5.8844, \quad (2.10)$$

with the total flux $\Gamma_{\text{flux}} = 10^{18} \text{ (m}^2 \text{ s)}^{-1}$. This applies for an H-mode discharge at 15 MA and 500 MW fusion power. For an L-mode case, when the plasma energy is lower, the total flux would be approximately 25 % of that of the H-mode (Martín-Solís *et al.* 2017). Part of the flux is caused by the time-integrated activation of the walls, and the typical ratio between the prompt gamma flux from fusion neutrons and the gamma flux after the fusion reactions cease is 1000. These calculations were performed for a beryllium first wall. ITER currently plans to have a first wall made of tungsten, which might change the photon flux and spectrum.

The runaway production rate from Compton scattering can be obtained as

$$\left(\frac{df}{dt}\right)_C = \frac{n_{e,\text{tot}}}{2} \frac{1}{p^2} \int_{W_{\gamma 0}}^\infty dW_\gamma \Gamma_\gamma(W_\gamma) \sin\theta \frac{d\sigma}{d\Omega} \frac{d\theta}{dp} \quad (2.11a)$$

$$= \frac{n_{e,\text{tot}}}{2} \frac{1}{p^2} \int_{W_{\gamma 0}}^\infty dW_\gamma \Gamma_\gamma(W_\gamma) \frac{d\sigma}{d\Omega} \frac{\beta}{\left(\frac{W_\gamma}{m_e c^2} + 1 - \gamma\right)^2}, \quad (2.11b)$$

where we have used

$$\frac{d\theta}{dp} = \frac{d}{dp} \arccos \left(1 - \frac{\frac{m_e c^2 / W_\gamma}{\frac{W_\gamma}{m_e c^2 (\gamma - 1)} - 1}} \right) = \frac{p/\gamma}{\sqrt{1 - \cos^2\theta} (\gamma - 1)^2 \left(\frac{W_\gamma/m_e c^2}{\gamma - 1} - 1\right)^2}. \quad (2.12)$$

Both the kinetic tritium and Compton sources are consistent with the fluid runaway generation rate (Martín-Solís *et al.* 2017)

$$\left(\frac{dn_{\text{re}}}{dt}\right)_{T/C} = \int_{p > p_c} d^3p \left(\frac{df}{dt}\right)_{T/C}, \quad (2.13)$$

when the runaway region is set to be above the critical momentum for RE generation p_c .

3. Bayesian optimization of simulated disruptions

In this paper, we will compare disruption simulations using an isotropic reduced kinetic model, described in appendix B.2 of Hoppe *et al.* (2021a), and a computationally less expensive fluid model. The comparison of fluid and kinetic simulations for a wide variety of disruption scenarios was facilitated using Bayesian optimization. The disruption model and simulation details used are described in § 3.1, while the Bayesian optimization is detailed in § 3.3. In § 3.2, the design and motivation of the cost function is presented.

3.1. Simulation set-up

The simulations are performed using the numerical tool DREAM (Disruption Runaway Electron Analysis Model). DREAM simulates the plasma evolution of tokamak disruptions self-consistently, including the evolution of the plasma temperature, current densities,

electron and ion densities, ion charge states, and electric field. The simulations are performed with a multi-fluid plasma model, with the possibility of simulating some, or all, subsets of the total electron population kinetically. DREAM divides the electrons into up to three populations.

Both in the fluid and the kinetic simulations performed here, the bulk and RE populations are treated as distinct fluids. The bulk electrons are characterized by their density n_e , temperature T_e and the Ohmic current density j_Ω . The REs are described by their density n_{re} , which also yields the runaway current density $j_{re} = ecn_{re}$, because they are assumed to move with the speed of light parallel to the magnetic field. In reactor-scale tokamak disruptions, this approximation, i.e. $v_{re} = c$, is typically valid, as supported by Buchholz (2023). The total current density $j_\Omega + j_{re}$ is governed by the evolution of the poloidal flux $\psi(r)$ (Hoppe *et al.* 2021a). For the electrical conductivity σ which enters in this evolution, we use the model described by Redl *et al.* (2021) that is valid for arbitrary plasma shaping and collisionality, and takes into account the effects of trapping. Neutral collisions – which are only relevant at very low post-TQ temperatures that yield strongly unfavourable outcomes – are not accounted for in the resistivity, as they are not relevant in the parameter ranges of interest (Vallhagen *et al.* 2020).

In the kinetic simulations, there is an additional population of superthermal electrons that are represented by the distribution function f_{hot} , which is evolved using a pitch-angle averaged Fokker–Planck equation (Hoppe *et al.* 2021a). This division of the hot and cold populations is especially suitable when there are two distinct Maxwellian electron populations present in the plasma, as is the case when cold materials are injected into a hot plasma during massive material injection.

The temperature of the cold population is determined by a time dependent energy balance equation, accounting for Ohmic heating, line and recombination radiation, and bremsstrahlung, as well as a radial heat transport, according to (43) of Hoppe *et al.* (2021a). The ionization, recombination and radiation rates are taken from the ADAS database for neon and the AMJUEL database¹ for deuterium. AMJUEL contains coefficients which account for opacity to Lyman radiation, which is important at high deuterium densities (Vallhagen *et al.* 2022).² The cold electron density is constrained by the condition of quasi-neutrality.

In the kinetic model, the Dreicer and hot-tail generation mechanisms are naturally present as velocity space particle fluxes, while in the fluid model, they are both implemented as generation rates. The Dreicer generation rate is computed by a neural network (Hesslow *et al.* 2019b) if the input parameters are within the training interval. However, for low electric fields, outside the training interval ($E/E_D < 0.01$), the values given by the neural network are inaccurate. In such cases, an analytic extrapolation is performed to recover a similar asymptotic behaviour as to the Connor–Hastie formula (Connor & Hastie 1975), i.e. $dn_{re}/dt \propto \exp(-A(E_D/E)^b) \rightarrow 0$ as $E_D/E \rightarrow \infty$. Here, A and b are positive parameters fitted to the dependence of the neural network on the electric field for $E/E_D > 0.01$, with the density and temperature fixed at the values of the current time step. They are fitted under the constraint that $(dn_{re}/dt)_{Dreicer}(E/E_D)$ should be once differentiable for all values of E/E_D .

The hot-tail generation rate is evaluated from the model distribution function of Smith & Verwichte (2008) and using the formula derived by Svenningsson (2020) – it is presented

¹Documentation for the AMJUEL database: www.eirene.de/Documentation/amjuel.pdf.

²The effective ionization loss rate, including the line radiation emitted during the bound-bound transitions leading up to the ionization, was adjusted to enforce that it only takes values larger than the ionization potential for each ionization event. This adjustment was only active at electron densities $\sim 10^{22} \text{ m}^{-3}$ and temperatures $\lesssim 1 \text{ eV}$, close to the limits of the validity range for the provided polynomial fit.

in appendix C.4 of Hoppe *et al.* (2021a). In both fluid and kinetic models, the avalanche generation is determined by the model by Hesslow *et al.* (2019a).

In the fluid simulations, the activated generation sources from tritium decay and Compton scattering are evaluated according to Vallhagen *et al.* (2020); in the kinetic simulations, we use the sources derived in § 2. During the TQ, the photon flux is set to $\Gamma_{\text{flux}} = 10^{18} \text{ (m}^2 \text{ s)}^{-1}$ to be consistent with the prediction of the photon flux at ITER for an H-mode discharge at 15 MA and 500 MW fusion power (Martín-Solís *et al.* 2017). However, with the drop in plasma temperature, the fusion power ceases. For this reason, the photon flux after the TQ is decreased by a factor of 10^3 , as was done by Vallhagen *et al.* (2024).

The disruption simulations are performed in an ITER-like tokamak scenario with either a deuterium plasma of density 10^{20} m^{-3} for the non-activated scenarios or a deuterium-tritium plasma with the same total ion density and 50–50% isotope concentrations. In both scenarios, the fuel density is spatially uniform. Discharges without tritium are simulated for 200 ms, while activated discharges are simulated for 400 ms, to capture the interesting runaway current dynamics in both scenarios.

The magnetic field $B_0 = 5.3 \text{ T}$ on axis, the resistive wall time $\tau_w = 0.5 \text{ s}$, the major radius $R_0 = 6 \text{ m}$ and the plasma minor radius $a = 2 \text{ m}$. Pusztai *et al.* (2023) evaluated the effective radius of the first toroidally closed conducting structure by matching the poloidal magnetic energy to that in simulated ITER discharges, yielding $b = 2.833 \text{ m}$. However, during disruptions, the plasma can be vertically displaced, resulting in a continuous contraction of the confined region and a corresponding reduction of the poloidal magnetic field energy reservoir that contributes to runaway generation. In this case, a more tightly fit wall ($b = 2.15 \text{ m}$) might be more appropriate. Simulations are therefore performed with both of these values of the conducting wall radius to bracket a plausible range of magnetic to kinetic energy conversion. Furthermore, the geometry is determined by realistic and radially varying shaping parameters congruent with the Miller model equilibrium (Miller *et al.* 1998), used also by Pusztai *et al.* (2023).

At $t = 0$, the plasma current $I_p = 15 \text{ MA}$ with current density profile $\hat{j}(r) = [1 - (r/a)^2]^{0.41}$, the plasma temperature is parabolic with $T(r = 0) = 20 \text{ keV}$ and the plasma is fully ionized. The deuterium and neon of the MMI are assumed to consist of cold atom populations, deposited instantaneously and homogeneously, as was done by Pusztai *et al.* (2023).

To account for the effect of the stochastic magnetic field during the TQ, we employ spatially and temporally constant magnetic perturbations of $\delta B/B = 0.5 \%$. Such a value of $\delta B/B$ is consistent with representative transport levels in MHD simulations of ITER disruptions (Hu *et al.* 2021). Although we do not consider different values of $\delta B/B$ here, we note that lower values tend to reduce conducted heat losses as well as the runaway seed, thereby increasing the region in parameter space that yields tolerable disruptions, while the location of the optimum is less affected, consistent with the findings of Pusztai *et al.* (2023). Following the Rechester–Rosenbluth model (1978), this magnetic perturbation leads to a transport of hot electrons, REs and electron heat, with a diffusion coefficient $D \propto |v_{\parallel}| R_0 (\delta B/B)^2$. Here, $|v_{\parallel}|$ is replaced by $v_{te} = \sqrt{2T_e/m_e}$ (local electron thermal speed) for the heat transport and by the speed of light c for the RE transport, where we assume a parallel streaming of the REs along the perturbed field lines. As noted by Pusztai *et al.* (2023), this approach gives an upper bound on the effect of the RE transport as it does not account for the energy and angular dependence of the RE distribution, as well as finite Larmor radius and orbit width effects. In the kinetic simulations, the parallel velocity dependence for the hot electrons is determined by f_{hot} .

The transport of superthermal and runaway electrons is switched off during the current quench to signify that the magnetic flux surfaces are reformed. We still use a small but finite heat transport corresponding to $\delta B/B = 0.04\%$ to avoid the development of non-physical narrow hot Ohmic channels (Putvinski *et al.* 1997; Fehér *et al.* 2011). As long as such channels do not try to form, this low transport level is subdominant to radiative heat losses at post-TQ temperatures. In the absence of a clear, widely accepted definition of the end of the TQ, in the simulations, we define the end of the thermal quench by the mean temperature falling below 20 eV. We discard simulations where the TQ condition has not been fulfilled within 20 ms, which is a factor of ~ 10 greater than the anticipated TQ time of 1–2 ms according to Hollmann *et al.* (2015). In such cases, we can assume that the TQ was not successfully completed and other complications might arise, such as reheating and unsuccessful current quench.

We note that the criterion used to switch the energetic electron transport off and reduce the heat transport is a source of uncertainty, as the time it takes for flux surfaces to reform might be delayed somewhat after the temperature has reached its post-TQ value. Clearly, longer windows of runaway loss are expected to reduce the final RE current. However, the presence of a relatively small reformed region may be sufficient to support significant RE currents (Tinguely *et al.* 2023b).

3.2. Cost function

The cost function of an optimization problem is critical for the optimization to yield meaningful results and, here, the cost function should reliably quantify the risks associated with tokamak disruptions, in the form of a single scalar. First, the runaway current should be accounted for in the cost function. We will consider two plausible choices for representing the runaway current contribution to the cost function – the maximum runaway current $I_{\text{re}}^{\text{max}} = \max_t I_{\text{re}}(t)$ and the runaway current at the time when it makes up 95% of the plasma current $I_{\text{re}}^{95\%} = I_{\text{re}}(t_{I_{\text{re}}=0.95I_p})$, which we will call the 95% runaway current. To account for the advantages and disadvantages of both options (discussed in Appendix A), we choose to use the first occurrence of $I_{\text{re}}^{95\%}$ unless $I_{\text{re}}^{\text{max}}$ occurs before $I_{\text{re}}^{95\%}$ or if $I_{\text{re}}(t) < 0.95I_p(t)$ throughout the simulation. This will be called the representative runaway current $I_{\text{re}}^{\text{repr}}$.

Additionally, any remaining Ohmic current at the end of the simulation has the potential of being converted into runaway current, as well as indicating if the termination of the discharge was successful. Consequently, it should also be included in the cost function. Based on Lehnen & the ITER DMS task force (2021), the upper safe operational limit of the runaway current is 150 kA, which will also be used as the upper safe operational limit for the Ohmic current.

Another indication of the success of the disruption mitigation system is the duration of the CQ (CQ time). In this work, the CQ time is evaluated through extrapolation, as $\tau_{\text{CQ}} = (t_{I_{\Omega}=3\text{MA}} - t_{I_{\Omega}=12\text{MA}})/0.6$ if the Ohmic current drops below 3 MA during the simulation and, otherwise, $\tau_{\text{CQ}} = (t_{\text{final}} - t_{I_{\Omega}=12\text{MA}})/(0.8 - I_{\Omega}^{\text{final}}/I_p^{t=0})$. Here $t_{I_{\Omega}=3\text{MA}}$ ($t_{I_{\Omega}=12\text{MA}}$) is the time when I_{Ω} is 20% (80%) of the initial plasma current. As stated by Hollmann *et al.* (2015), its safe operational interval is [50, 150] ms for 15 MA discharges.

Finally, the last figure of merit considered is the transported heat loss fraction η_{cond} , meaning the fraction of the initial plasma kinetic energy which has been lost from the plasma due to energy transport. It is accounted for using the Rechester–Rosenbluth transport in the collisionless limit, which means that the radial heat diffusion is caused by parallel streaming of electrons along perturbed field lines (Rechester & Rosenbluth 1978), as detailed in § 3.1. For a safe disruption, the transported heat load fraction should be below 10%, according to Hollmann *et al.* (2015).

Our cost function $\mathcal{L}(I_{re}^{repr}, I_{\Omega}^{final}, \tau_{CQ}, \eta_{cond})$ is designed to yield a value below one if all components are within their intervals of safe operation. To achieve this, I_{re}^{repr} , I_{Ω}^{final} and η_{cond} are normalized to their respective safe operational limits. The CQ time is shifted and normalized according to $|\tau_{CQ} - 100 \text{ ms}|/50 \text{ ms}$ to achieve the same behaviour for both the lower and upper limit of safe operation. The four figures of merit can be combined using different orders of p -norms and here we have chosen to use the Euclidean norm ($p = 2$). In [Appendix A](#), we discuss the impact of different choices for p .

The final consideration in the design of our cost function is the relative importance of the components. With what we have discussed so far, a disruption with $I_{re}^{repr} = 120 \text{ kA}$ and $\tau_{CQ} = 100 \text{ ms}$ would yield the same value of \mathcal{L} as $I_{re}^{repr} = 0 \text{ kA}$ and $\tau_{CQ} = 60 \text{ ms}$ (if $I_{\Omega}^{final} = 50 \text{ kA}$ and $\eta_{cond} = 3 \%$), whereas the latter would be preferable in an actual disruption. More specifically, the relevant consideration for the CQ time is that it is between 50 ms and 150 ms (and preferably with some margin for uncertainties), not necessarily that it is as close to the middle of this safe interval (i.e. 100 ms) as possible. For the RE current however, it is preferable that it is as close to zero as possible, even if RE currents up to 150 kA are tolerable. Furthermore, minimizing I_{re}^{repr} was, in this work, deemed more important than minimizing η_{cond} . To achieve this order of significance when the figures of merit are within their safety intervals, we introduced nonlinear (but smooth) rescaling functions f_i , which take the form

$$f_i = \begin{cases} (x_i)^{k_i} & \text{if } x_i \leq 1, \\ k_i(x_i - 1) + 1 & \text{if } x_i \geq 1, \end{cases} \tag{3.1}$$

for the currents $k_I = 1$ and $x_I = I/150 \text{ kA}$, for the conducted heat load $k_{\eta_{cond}} = 3$ and $x_{\eta_{cond}} = \eta_{cond}/10 \%$, and for the CQ time $k_{\tau_{CQ}} = 6$ and $x_{\tau_{CQ}} = |\tau_{CQ} - 100 \text{ ms}|/50 \text{ ms}$. Note that this implementation of f_i is differentiable.

The final expression for the cost function can thus be written as

$$\mathcal{L} = \sqrt{(C_{I_{re}^{repr}} f_{I_{re}^{repr}})^2 + (C_{I_{\Omega}^{final}} f_{I_{\Omega}^{final}})^2 + (C_{\eta_{cond}} f_{\eta_{cond}})^2 + (C_{\tau_{CQ}} f_{\tau_{CQ}})^2}. \tag{3.2}$$

Each of the components are weighted equally, with weight $C = 0.5$, to ensure values below one for scenarios of safe operation.

For some simulations, the cost function cannot be evaluated – mainly due to lack of convergence or incomplete thermal quench within 20 ms. In these cases, the cost function was set to a high value ($\mathcal{L} = 75 \gg 1$) to decrease the probability of the optimizer exploring the surrounding area further.

3.3. Optimization set-up

Bayesian optimization is an optimization strategy which uses Bayesian inference to estimate a probability distribution function for the entire objective (or cost) function based on a prior set of samples (input-output pairs). This is achieved by assuming that all unknowns of the system to be inferred are random variables.

To make viable predictions for the objective function, an assumption of the nature of its probability distribution is needed. For this, Bayesian optimization uses stochastic processes, which are infinite collections of random variables (Garnett 2023). A common stochastic process used for Bayesian optimization is the Gaussian process (GP), in which the random variables are distributed according to multivariate Gaussian distributions. In the context of Bayesian optimization, a GP is specified by a mean function $\mu(x)$, which determines the expected objective function value at any point x of the search space, and

a covariance function $\phi(x, x')$, which measures the correlation between points x and x' (Garnett 2023). Consequently, this optimization method not only yields an optimum, but also an estimate of the objective function itself in the form of the mean function $\mu(x)$ as well as its uncertainty.

During the optimization, the optimizer advances by taking samples of the objective function and successively updating its sample set and subsequently its GP estimate of the probability distribution function. The strategy for choosing which point to sample next is defined by the acquisition function, which uses the GP estimate to determine which point has the highest potential of improving the current optimum.

In this work, the Bayesian optimization was performed using the Python package by Nogueira (2014) which uses GPs. The Matérn kernel with smoothness parameter $\nu = 3/2$ was used for the GPs and the expected improvement acquisition function was used as the optimization strategy (Garnett 2023).

4. Disruption mitigation optimization

The goal of this section is to compare fluid and kinetic simulations of MMI scenarios for a range of injected material densities. Since the RE dynamics is different for discharges with and without tritium, we will do this comparison both with and without activated sources in § 4.2 and § 4.1, respectively. While the main focus will be on the differences between the models, we will also discuss the simulated disruption outcomes, including a comparison of our results with those of Pusztai *et al.* (2023). We will furthermore comment on the impact of choice of wall radius on the results.

The code used to produce the results in this section, as well as the optimization data, can be found at Zenodo.³

4.1. Optimization of non-activated discharges

The non-activated scenarios are characterized by only having the Dreicer and hot-tail generation mechanisms as seed generation sources,⁴ which tend to be more susceptible to runaway mitigation measures, as we shall see. For the non-activated scenarios, the optimization was performed on the MMI densities within $\log(n_D/[1 \text{ m}^{-3}]) \in [20, 22]$ and $\log(n_{Ne}/[1 \text{ m}^{-3}]) \in [15, 20]$ for the fluid model. For the kinetic model, the intervals were set as $\log(n_D/[1 \text{ m}^{-3}]) \in [21, 22]$ and $\log(n_{Ne}/[1 \text{ m}^{-3}]) \in [17.5, 20]$. The total injected number of atoms of $\sim 10^{25}$ corresponding to the upper end of the deuterium density ranges, 10^{22} m^{-3} , is an approximate upper bound on material assimilation using shattered pellet injection at ITER (Vallhagen *et al.* 2024). When using the fluid (kinetic) model, the optimization was performed using 250 (200) disruption simulations.

Figures 1(a) and 1(b) show the results from the optimization of the non-activated scenario for both the fluid and kinetic models, respectively. The most prominent difference between the results of the fluid and kinetic models is that there exists a region of safe operation (as indicated by blue shades in the figure) for the simulations done with the kinetic model, but not with the fluid model. With the kinetic model, the band of safe operation reaches from $(n_D, n_{Ne}) = (10^{21.5}, 10^{19.5}) \text{ m}^{-3}$ to $(n_D, n_{Ne}) = (10^{22}, 10^{19}) \text{ m}^{-3}$. Furthermore, the kinetic model yields lower values of the cost function overall compared with the fluid model for $\log(n_D/[1 \text{ m}^{-3}]) \in [21, 22]$ and $\log(n_{Ne}/[1 \text{ m}^{-3}]) \in [17.5, 20]$.

We find that the main difference between the two models is the runaway current. This is illustrated in figure 2, where the safe regions of each of the cost function components

³Simulation data for this article: <https://doi.org/10.5281/zenodo.10998363>.

⁴Note that in the non-activated cases, we assume that the gamma flux from the wall – which might be induced by D-D neutron bombardment – is sufficiently small that the contribution of the corresponding Compton seed to the final runaway current is negligible compared with that of the non-activated sources.




FIGURE 1. Logarithmic contour plots of the cost function estimate μ for the non-activated scenario, generated using (a) the fluid and (b) the kinetic model in DREAM. Note that the colour mapping is adapted such that blue shades represent regions of safe operation. The black star (diamond) indicates the optimal samples found using the fluid (kinetic) model, while the black dots indicate all optimization samples. The grey area covers the region of incomplete TQ.

are indicated by the coloured areas. For the fluid model, the safe region of $I_{\text{re}}^{\text{repr}}$ only covers the lower part of the $\log(n_{\text{D}}/[1 \text{ m}^{-3}]) \in [21, 22]$ and $\log(n_{\text{Ne}}/[1 \text{ m}^{-3}]) \in [17.5, 20]$ square, but for the kinetic simulations, it covers all but the upper right corner. Thus, with the kinetic model, we find that the band of safe operation that can be observed in figure 1(b) corresponds to the overlap of $I_{\text{re}}^{\text{repr}} < 150 \text{ kA}$ and $\eta_{\text{cond}} < 10 \%$, which are the two competing components. Similarly, the optimum found using the fluid model is located at a point between the region of safety for $I_{\text{re}}^{\text{repr}}$ and η_{cond} , since here there is no overlap. It is worth noting, however, that such a region of overlap of $I_{\text{re}}^{\text{repr}} < 150 \text{ kA}$ and $\eta_{\text{cond}} < 10 \%$ might not be possible with a tungsten wall, as the additional high- Z impurity can have a more significant negative impact concerning the runaway current (Walkowiak *et al.* 2024) than the improvement the extra radiative losses represent for η_{cond} .

The shapes of the region of safety for both $I_{\Omega}^{\text{final}}$ and τ_{CQ} are also noticeably different, though not to the same degree as for $I_{\text{re}}^{\text{repr}}$. This is simply a consequence of the difference in runaway current. In cases when less of the Ohmic current is transformed into runaway current, the Ohmic current is higher at the end of the simulation and the CQ time is therefore longer. This causes both of the regions of $I_{\Omega}^{\text{final}} < 150 \text{ kA}$ and $50 \text{ ms} < \tau_{\text{CQ}} < 150 \text{ ms}$ to shift upwards when the $I_{\text{re}}^{\text{repr}} < 150 \text{ kA}$ region expands. The choice of plasma model does not have a significant impact on the conducted heat load – for both models, the safe region covers the upper right corner of the region explored with the kinetic model.

The results shown in figures 1 and 2 correspond to simulations with a wall radius of 2.833 m, but the results are qualitatively similar for a wall radius of 2.15 m – the corresponding landscapes as presented in figure 2 are much the same. The main difference is that with the smaller wall radius, the currents are in general lower and the CQ times are shorter. In figure 3, the plasma and runaway currents are presented for the optimum found using the kinetic model, as well as fluid simulations for the same parameters using $a = 2.833 \text{ m}$ and $a = 2.15 \text{ m}$. With the larger wall radius, $I_{\text{re}}^{\text{repr}} \approx 4 \text{ MA}$, while the smaller wall radius (which also uses a resistive wall time of 0.5 s) yields $I_{\text{re}}^{\text{repr}} \approx 400 \text{ kA}$. This dramatic reduction in the representative runaway current caused by placing the conducting wall near the plasma, thereby reducing the magnetic energy reservoir for runaway acceleration, is consistent with recent findings reported by Pusztai *et al.* (2023)

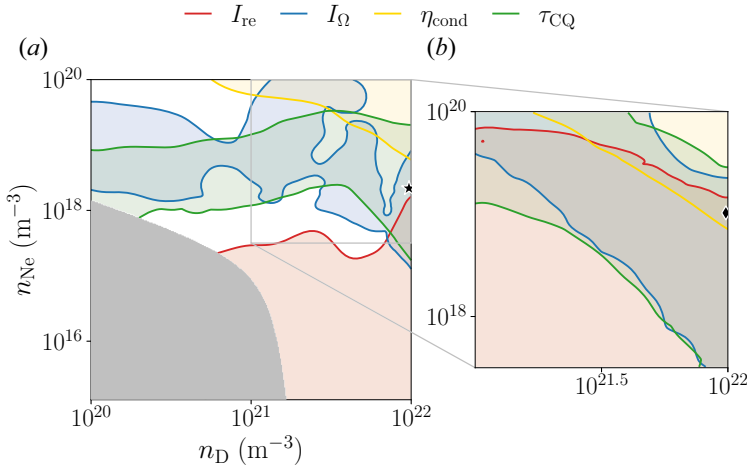


FIGURE 2. Regions of safe operation (shaded) for the non-activated case with regards to I_{re}^{repr} (red), I_{Ω}^{final} (blue), τ_{CQ} (green) and η_{cond} (yellow). Panel (a) uses fluid simulations and panel (b) uses kinetic simulations. The optimal samples are indicated by a star in panel (a) and a diamond in panel (b). Note that only in the kinetic case, there is a finite intersection of all the regions of safe operation.




FIGURE 3. Total plasma current (solid) and runaway current (dashed) of the optimum found using the kinetic model, simulated with (a) the kinetic model and (b) the fluid model and $a = 2.833$ m, as well as (c) the fluid model and $a = 2.15$ m.

and Vallhagen *et al.* (2024). The reason to consider such a reduced wall radius – compared with the nominal one matching the total available magnetic energy content in ITER – is to emulate the reduction of the magnetic energy that is converted to runaway acceleration in a vertical displacement event, as the confined region is gradually scraped off from the plasma channel. This approach can give a lower bound on the representative runaway current.

The difference in runaway current for the two models is dominated by differences in the modelling of the hot-tail RE-generation mechanism, which is highly dependent on the energy distribution of the superthermal electrons. In the kinetic simulations, f_{hot} describes the energy distribution of the superthermal electrons, which enables the hot-tail generation to be included in the velocity flux determined by the Fokker–Planck collision operator. This energy distribution of the superthermal electrons is not evolved in the fluid

simulations and, therefore, a model for the energy distribution based on the simulated plasma parameters is needed. Such a model, with the purpose of being used to determine the hot-tail generation rate for fluid simulations, was derived by Smith & Verwichte (2008). They neglect the electric field and diffusion terms in the Fokker–Planck equation, and instead only retain collisional friction, which gives the solution

$$f_{\text{hot}}^{\text{fl}}(t, p) = \frac{n_0}{\pi^{3/2} p_{\text{Te}}^3} \exp \left[-\frac{(p^3 + 3\tau(t))^{2/3}}{p_{\text{Te}}^2} \right], \quad (4.1)$$

where n_0 is the initial free electron density, $p_{\text{Te}} = \sqrt{2T_0/m_e c^2}$ is the initial thermal momentum of the electrons normalized to $m_e c$ (with T_0 being the initial electron temperature) and $\tau(t)$ is the time-integrated collision frequency. Using this model distribution function for the superthermal electrons, the hot-tail generation rate of the fluid model is derived by Svenningsson (2020) as

$$\frac{\partial n_{\text{re}}}{\partial t} = -4\pi p_c^{\text{fl}2} \frac{\partial p_c^{\text{fl}}}{\partial t} f_{\text{hot}}^{\text{fl}}(t, p_c^{\text{fl}}), \quad (4.2)$$

where the critical momentum p_c^{fl} is defined as the momentum at which the electric field acceleration and collisional friction balance each other; see (C.24) in Hoppe *et al.* (2021a).⁵

The optimum found using the kinetic model is a clear example of the fluid and kinetic models yielding significantly different RE-current dynamics, as shown in figure 3, due to the difference in hot-tail generation implementation in the two models. For the same MMI densities in a fluid simulation, we obtain a several MA runaway current, where the seed REs are exclusively generated by the hot-tail mechanism, while the kinetic simulation gives practically zero runaway current. In figure 4, the distribution function derived by Smith & Verwichte (2008) and the distribution function evolved in the kinetic simulation are plotted just before and at end of the TQ for $r = 0.85$ m, which is the radial position with the highest hot-tail generation rate in the fluid simulation. The two distribution functions $f_{\text{hot}}^{\text{fl}}$ and $f_{\text{hot}}^{\text{kin}}$ are similar initially, as shown in figure 4(a). The discrepancy at $p \gtrsim 1$ can be fully explained by $f_{\text{hot}}^{\text{fl}}$ being a Maxwell–Jüttner distribution function, while $f_{\text{hot}}^{\text{kin}}$ is a non-relativistic Maxwell–Boltzmann distribution.

For this radius, the hot-tail generation peaks at the end of the TQ in the fluid simulation and then the two distribution functions are significantly different, as shown in figure 4(b). Here, $f_{\text{hot}}^{\text{kin}} \gg f_{\text{hot}}^{\text{fl}}$ around p_c^{fl} (by many orders of magnitude; note the log-scale on a wide range), which can be explained by the kinetic simulations accounting for a radial diffusive transport of $f_{\text{hot}}^{\text{kin}}$ due to magnetic perturbations, while the evaluation of $f_{\text{hot}}^{\text{fl}}$ does not. A generalization of the fluid hot-tail rate, accounting for transport of the hot electrons, would not be straightforward, due to the velocity dependence of the Rechester–Rosenbluth transport. The neglect of this transport will lead to a significant overestimation of the runaway seed due to the factor of $f_{\text{hot}}^{\text{fl}}(t, p_c^{\text{fl}})$ in (4.2). This was confirmed by comparing the fluid and kinetic simulations with the radial transport removed.

As shown in figure 1, the optima found with both the fluid and kinetic models are located at high deuterium densities. As noted in § 3.1, at high deuterium densities, the effects of accounting for opacity to Lyman radiation become important. When not accounting for opacity to Lyman radiation at these optimum, the temperature decays faster, resulting

⁵The correct form of (C.25) of Hoppe *et al.* (2021a) is (4.2).




FIGURE 4. The distribution function for the optimal non-activated case, found using the kinetic model (a) just before and (b) at the end of the TQ. The distribution function f_{hot} (dashed) has been evolved in the kinetic simulation and $f_{\text{hot}}^{\text{fl}}$ (solid) is the model distribution used in the fluid simulation, given by (4.1). The latter is based on values of the plasma parameters evolved in the fluid simulation.

in a significantly larger runaway current, which is in agreement with the conclusions of Vallhagen *et al.* (2022).

In conclusion, the model distribution function used by the fluid model lacks electric field acceleration and radial transport effects, and it is the transport that dominates the differences observed between the two models. We note that previous work showed that if the losses due to magnetic perturbations are neglected, i.e. taking into account all the hot-tail electrons, the final runaway current is overestimated by a factor of approximately four in ASDEX Upgrade (Hoppe *et al.* 2021b), which is in qualitative agreement with our conclusions here.

4.2. Optimization of activated discharges

In activated scenarios, the RE seed is expected to be dominated by the generation of REs from tritium beta decay and Compton scattering in cases when the hot-tail and Dreicer generation sources have been successfully mitigated. The same intervals and number of samples were used for the optimization of the activated scenarios as for the non-activated scenarios, except that the neon density interval was changed to $\log(n_{\text{Ne}}/[1 \text{ m}^{-3}]) \in [15, 17.5]$ for the optimization with the kinetic model to include regions of lower runaway current.

Figure 5 shows the result from the optimizations of the activated scenarios. There are no regions of safe operation, neither when using the fluid nor the kinetic model – in fact, $\mathcal{L} > 10$ (based on the mean prediction μ) everywhere for both models and both models yield almost identical results. From the optimization using the fluid model, there are two local minima, which both are around $\mathcal{L} = 15$. Since the cost function was designed to distinguish safe from unsafe disruptions, and this distinction happens around $\mathcal{L} \sim 1$, it is not as reliable at determining which of the two disruptions is better if their cost function values are similar and large while the individual figures of merit are different. Therefore, the global minimum (in figure 5a) is not necessarily safer than the other minimum.

The optimization using the kinetic model presented in figure 5(b) has not explored the region around the global minimum found with the fluid model, thus its minimum has the same location as the other local minimum found using the fluid model. However, the location of this global optimum did not change when performing a similar optimization




FIGURE 5. Logarithmic contour plots of the cost function estimate μ for the activated scenario, generated using (a) the fluid and (b) the kinetic model in DREAM. The black star (diamond) indicates the optimal samples found using the fluid (kinetic) model, while the black dots indicate all optimization samples.

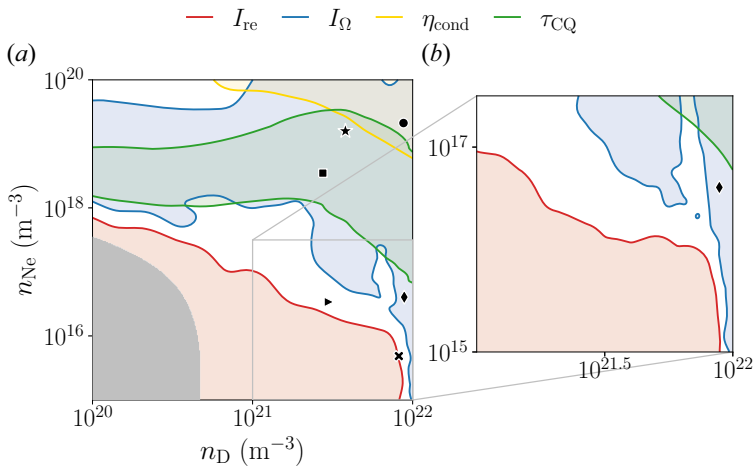


FIGURE 6. Regions of safe operation (shaded) for the activated case with regards to I_{re}^{repr} (red), I_{Ω}^{final} (blue), τ_{CQ} (green) and η_{cond} (yellow). Panel (a) uses fluid simulations and panel (b) uses kinetic simulations. The optimal samples are indicated by a star in panel (a) and a diamond in panel (b). The other markers correspond to the cases in table 1.

using the kinetic model which covered regions surrounding both optima found using the fluid model.

That the two plasma models yield very similar results is also clear from figure 6, where the regions of safe operation for each figure of merit are plotted. The two models yield qualitatively the same results, since the shapes of the safe regions are remarkably alike to the point where it is difficult to distinguish differences due to plasma models from stochastic differences of the sampled locations.

In table 1, fluid and kinetic simulations are compared for different combinations of the deuterium and neon densities, including the two optima (the locations in n_D – n_{Ne} space are indicated by symbols in figure 6). The difference between the kinetic and fluid models are

Marker	n_D (m ⁻³)	n_{Ne} (m ⁻³)	Model	\mathcal{L}	I_{re}^{repr} (kA)	I_{Ω}^{final} (kA)	τ_{CQ} (ms)	η_{cond} (%)
★	3.8×10^{21}	1.6×10^{19}	Fluid	14.1	4170	103	68.0	23.8
			Kinetic	14.7	4320	103	66.3	26.9
◆	8.9×10^{21}	4.0×10^{16}	Fluid	14.3	2790	92.0	201	75.2
			Kinetic	14.4	2820	88.8	201	75.8
✕	8.2×10^{21}	4.9×10^{15}	Fluid	19.4	218	526	408	79.3
			Kinetic	19.5	225	539	408	79.9
■	2.7×10^{21}	3.5×10^{18}	Fluid	16.5	4030	80.5	108	70.7
			Kinetic	17.1	4210	83.6	103	71.8
●	8.8×10^{21}	2.1×10^{19}	Fluid	22.7	6810	82.5	43.8	2.55
			Kinetic	22.8	6830	80.5	43.7	4.73
▶	3.0×10^{21}	3.4×10^{16}	Fluid	29.3	3960	2200	500	92.9
			Kinetic	29.4	3950	2200	501	94.0

TABLE 1. Disruption figures of merits for fluid and kinetic simulations for parameters indicated in figure 6 with the same markers.

in most cases just a few percent, and the relative difference is larger when the figures of merit are small. The largest relative error occurs for the conducted heat load, since it is consistently a few percent larger in the kinetic case; this can be explained by the kinetic simulation capturing the transport of the hot electrons.⁶

In absolute terms, only the runaway current varies significantly between the kinetic and fluid results – for large runaway current cases, the difference can be ~ 150 kA, while the relative difference is small. As shown in table 1, the kinetic simulations have a higher runaway current in all cases but one. The reason for the kinetic model yielding larger runaway current is indirectly due to the activated sources. The direct generation of electrons above the critical momentum from the activated sources is practically identical for both the fluid and kinetic simulations. In the kinetic simulations however, the hot electrons generated by the activated sources with momenta below the critical momentum will cause an increase in RE generation due to flux in velocity space.

The dynamics of the RE seed generation in the activated scenarios are substantially different from those in the non-activated scenarios. In the non-activated case, the seed generation was dominated by the hot-tail mechanism. For an activated scenario, f_{hot} (which includes not only the hot-tail, but also superthermal electrons generated by the Dreicer, Compton scattering and tritium decay mechanisms) will not be depleted due to magnetic perturbations, because there will be a constant generation of hot electrons due to the activated sources. This means that the runaway seed generation will not necessarily be overestimated with the fluid model, as it was in the non-activated case.

In figure 7, the distribution functions f_{hot} and f_{hot}^{Ω} are plotted at the beginning and end of the TQ, at $r = 1.05$ m, for the optimal activated scenario found using the fluid model. We see that, in contrast to the non-activated case, $f_{hot}^{\Omega} \ll f_{hot}$ around p_c^{Ω} at the end of the TQ for the activated case (note the scale of the y-axis). The fluid simulation should still overestimate the runaway seed generation rate due to the neglect of superthermal

⁶Note that for most of the scenarios presented in table 1, the conducted heat loads are high enough to cause severe wall damage if localized.




FIGURE 7. The distribution function for the optimal activated case found using the fluid model (a) just before and (b) at the end of the TQ. The distribution function f_{hot} (dashed) has been evolved in the kinetic simulation and $f_{\text{hot}}^{\text{fl}}$ (solid) is the model distribution used in the fluid simulation, given by (4.1). The latter is based on values of the plasma parameters evolved in the fluid simulation.

electron transport in the fluid hot-tail model, which should lead to $f_{\text{hot}}^{\text{fl}} \gg f_{\text{hot}}$ as in the non-activated cases. However, the fluid model does not capture the additional dynamics in the superthermal momentum range caused by the generation of hot electrons from the activated sources, as discussed above, and this explains why $f_{\text{hot}}^{\text{fl}} \ll f_{\text{hot}}$ in the activated scenarios. Furthermore, for the activated cases, the inaccurate runaway seed generation will have a smaller impact on the runaway current evolution since the activated sources, and particularly the Compton source, will constitute a significant fraction of the seed. For both optima found in the activated scenarios, the Compton scattering generation mechanism dominates the runaway seed generation by several orders of magnitude.

The physical model and methodology of this study are similar to that of Pusztai *et al.* (2023), but the cost function landscapes and the found optima are significantly different from those presented in that work. The most significant difference compared with the simulations presented by Pusztai *et al.* (2023) is the cost function. Pusztai *et al.* (2023) used the following cost function:

$$\mathcal{L}_{\text{IP}}(I_{\text{re}}^{\text{max}}, I_{\Omega}^{\text{final}}, \tau_{\text{CQ}}, \eta_{\text{cond}}) = \frac{I_{\text{re}}^{\text{max}}}{150 \text{ kA}} + \frac{I_{\Omega}^{\text{final}}}{150 \text{ kA}} + 100 \theta(\tau_{\text{CQ}}) + 10 \frac{\eta_{\text{cond}}}{10 \%}, \quad (4.3a)$$

$$\theta(\tau_{\text{CQ}}) = \tilde{\Theta}(50 \text{ ms} - \tau_{\text{CQ}}) + \tilde{\Theta}(\tau_{\text{CQ}} - 150 \text{ ms}), \quad (4.3b)$$

where $\tilde{\Theta}(\tau) = [1 + \tanh(\tau/3.3 \text{ ms})]/2$. In (4.3a), $I_{\text{re}}^{\text{max}}$ is used instead of $I_{\text{re}}^{\text{repr}}$, but this does not have a significant impact on the results. Rather, differences between how τ_{CQ} and η_{cond} are treated are responsible for most of the difference in the results.

First, the step function behaviour of $\theta(\tau_{\text{CQ}})$ limits the optimization of the other parameters to only the region of safe operation for τ_{CQ} . All samples with $\tau_{\text{CQ}} \notin [50, 150]$ ms will have $\mathcal{L}_{\text{IP}} > 100$, which is an order of magnitude larger than the optimum found by Pusztai *et al.* (2023), making it difficult to distinguish the impact of the other components on the cost function value.

Second, the transported heat fraction is weighted one order of magnitude higher than the runaway current and final Ohmic current. This means that $\eta_{\text{cond}} = 10 \%$ will have the same impact on \mathcal{L}_{IP} as $I_{\text{re}}^{\text{max}} = 1.5 \text{ MA}$, but while $\eta_{\text{cond}} = 10 \%$ is marginally tolerable, $I_{\text{re}}^{\text{max}} = 1.5 \text{ MA}$ is not.

The cost function differences explain the difference in landscape of \mathcal{L}_{IP} compared with \mathcal{L} here. Furthermore, it explains why no case with $I_{\text{re}}^{\text{max}} < 4 \text{ MA}$ is found by Pusztai *et al.* (2023) – the cases in table 1 with $I_{\text{re}}^{\text{max}} < 4 \text{ MA}$ have $\tau_{\text{CQ}} > 150 \text{ ms}$ and $\eta_{\text{cond}} \approx 80 \%$. This yields $\mathcal{L}_{\text{IP}} \sim 200$ which is a rather high value of the cost function according to the scale in figure 1 of Pusztai *et al.* (2023).

Apart from the difference in cost function landscape, the optimum found using the fluid model is shifted to higher values of n_{Ne} and lower values of n_{D} compared with the optimum found by Pusztai *et al.* (2023). The main reason for this shift is the decrease of the photon flux from $\Gamma_{\text{flux}} = 10^{18} \text{ (m}^2 \text{ s)}^{-1}$ to $\Gamma_{\text{flux}} = 10^{15} \text{ (m}^2 \text{ s)}^{-1}$ after the TQ, which was not accounted for by Pusztai *et al.* (2023). This change in the photon flux leads to a delay in the runaway current ramp up, resulting in a longer decay time for the Ohmic current. Therefore, τ_{CQ} is increased from 49 ms to 67 ms, and thus would have a large impact also on \mathcal{L}_{IP} .

An important similarity between our result and those of Pusztai *et al.* (2023) is that we do not find any $n_{\text{D}}-n_{\text{Ne}}$ combinations yielding tolerable or close to tolerable figures of merit in the activated case. For the full explored region, $\mathcal{L} \gg 1$, and as we mentioned earlier, the cost function was not designed to be particularly informative at these values. To study correlations and trade-offs between different figures of merit, another approach is needed.

We can use that each simulation which we have performed corresponds to a point in the four dimensional parameter space spanned by our four figures of merit, i.e. $I_{\text{re}}^{\text{repr}}$, $I_{\Omega}^{\text{final}}$, τ_{CQ} , η_{cond} . In figure 8, these points have been projected on two-dimensional subspaces spanned by pairs of the figures of merit, which illustrates potential trade-offs between each pair. Figure 8(a–c) clarify that for scenarios with $I_{\text{re}}^{\text{repr}} < 1 \text{ MA}$, we get $\tau_{\text{CQ}} > 400 \text{ ms}$ and $\eta_{\text{cond}} > 75 \%$. More specifically, the correlation between $I_{\text{re}}^{\text{repr}}$ and η_{cond} suggests that it is not possible to achieve moderate values for the runaway current and the transported heat fraction simultaneously – when $I_{\text{re}}^{\text{repr}} < 4 \text{ MA}$, then $\eta_{\text{cond}} > 75 \%$ and *vice versa*. This is related to the almost rectangular envelope bounding the point cloud from the side of low values in the $I_{\text{re}}^{\text{repr}}-\eta_{\text{cond}}$ panel.

Disregarding the runaway current, there are some points of safety overlap for the other three figures of merit; see figure 8(d–f). All simulations with $\tau_{\text{CQ}} \in [50, 150] \text{ ms}$ have $I_{\Omega}^{\text{final}} < 150 \text{ kA}$, and since there is a slight overlap between $\tau_{\text{CQ}} \in [50, 150] \text{ ms}$ and $\eta_{\text{cond}} < 10 \%$, here $I_{\Omega}^{\text{final}} < 150 \text{ kA}$ as well. This can be related back to figure 6, where we see a small region of overlap for safe values of τ_{CQ} and η_{cond} , as well as the safe region for τ_{CQ} being almost fully covered by the safe region for $I_{\Omega}^{\text{final}}$.

Being able to visualize correlations between the cost function components and which regions of the $n_{\text{D}}-n_{\text{Ne}}$ space correspond to safe regions of the figures of merits is among the benefits of using a stochastic and exploratory optimization algorithm such as Bayesian optimization. This feature has enabled the exploration of how kinetic and fluid simulations compare for a large variety of scenarios, while simultaneously yielding information on the disruption evolution not just for the most optimal case, but for the full region of interest in the $n_{\text{D}}-n_{\text{Ne}}$ space. Such comprehensive comparison or thorough exploration of MMI of deuterium and neon would not have been possible with a more traditional optimization algorithm, such as gradient descent or the comparison of just a few scenarios.

5. Conclusions

We have compared kinetic and fluid simulations of disruptions with massive material injection for an ITER-like tokamak set-up. With this objective in mind, we derived and implemented kinetic runaway electron generation sources for the Compton scattering and tritium beta decay mechanisms. Furthermore, we investigated how the hot-tail generation

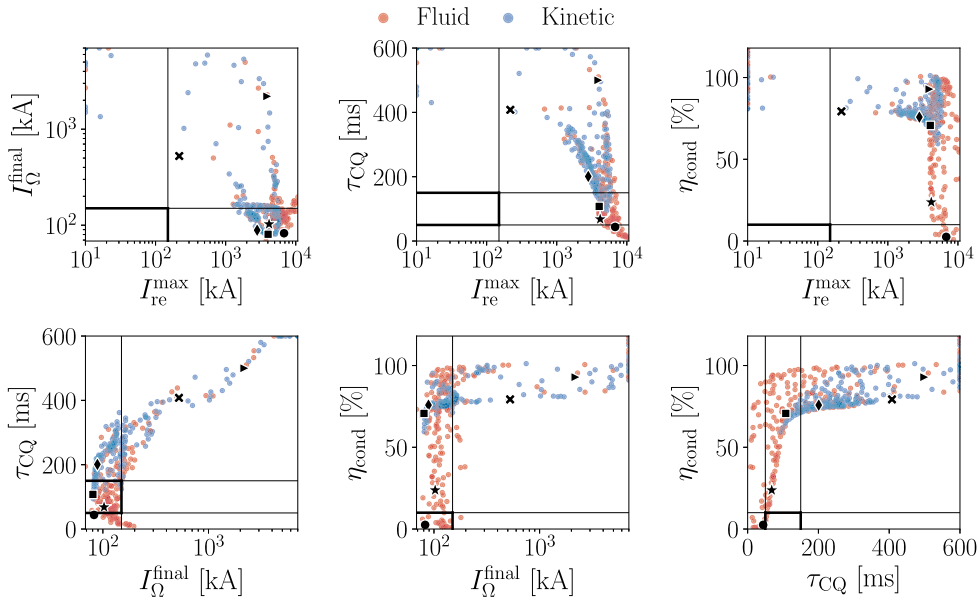


FIGURE 8. Projections of the simulation dataset to all the two-dimensional subspaces of the figure of merit space $(I_{re}^{repr}, I_{\Omega}^{final}, \tau_{CQ}, \eta_{cond})$, corresponding to the optimizations of the activated scenario. Both kinetic (blue) and fluid (red) simulations are shown. The intervals of safe operation for each cost function component is indicated by the black lines. This figure illustrates the trade-off between the different cost function components. The optimal samples found during the optimizations using the fluid (kinetic) model is indicated by a black star (diamond), while the other cases in table 1 are indicated as black markers of different shapes. Samples located outside of the plotted domains are placed at the edges.

compares between the two models, both for activated and non-activated scenarios. Comprehensive fluid-kinetic comparisons were facilitated using Bayesian optimization of the injected densities of deuterium and neon, which enabled exploration of a large variety of different disruption scenarios.

We found that while kinetic and fluid disruption simulations representing activated scenarios generated similar results, there were major differences for the non-activated scenarios, and this was caused by differences in the fluid and kinetic hot-tail generation dynamics. Since the kinetic model evolves the distribution function of the superthermal electrons in the simulations, it takes transport of these particles due to magnetic perturbations into account, while the model distribution used to derive the hot-tail generation rate used in the fluid simulations does not. This has a high impact on the non-activated simulations since this transport leads to a depletion of the distribution function, which limits the runaway seed generation in the kinetic simulations.

In the activated case, however, there is a constant generation of superthermal electrons in the kinetic simulations due to the activated sources, which acts against the depletion due to transport, and thus makes the difference between the fluid and kinetic models less severe. Furthermore, the additional generation of runaway electrons from the activated sources reduces the impact of the hot-tail and Dreicer generation on the magnitude of the runaway current during the CQ. This limits the significance of the overestimation of the runaway seed generation due to the neglect of superthermal electron transport in the fluid hot-tail model.

The cost function used for the Bayesian optimization was carefully designed, with the feature of safe disruption simulations corresponding to cost function values below unity. This had a significant impact on the cost function landscape in the explored MMI density region and the optima found during the optimization, when compared with a similar optimization done by Pusztai *et al.* (2023).

Outstanding issues that remain to be investigated include the effect of vertical displacements events and other instabilities on the runaway electron dynamics. Additionally, there might be remnant magnetic perturbations even after the thermal quench, which will affect the final runaway current. Finally, the injected material will in reality not be distributed uniformly. The material injection can also be divided into several stages, and recent studies indicate that multi-pellet injections are advantageous for disruption mitigation (Vallhagen *et al.* 2024). The impact of radial inhomogeneities in the assimilated material were, to some extent, explored by Pusztai *et al.* (2023), finding beneficial effects of outward peaking neon and inward peaking deuterium distributions. However, such inhomogeneities might be less consequential for the runaway current evolution in the presence of an effective particle transport during the TQ (Vallhagen *et al.* 2024).

Acknowledgements

The authors are grateful to H. Bergström, P. Halldestam and S. Newton for fruitful discussions.

Editor P. Helander thanks the referees for their advice in evaluating this article.

Funding

This work was supported by the Swedish Research Council (Dnr. 2022-02862 and 2021-03943), and by the Knut and Alice Wallenberg foundation (Dnr. 2022.0087). The work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 – EUROfusion). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Declaration of interests

The authors report no conflict of interest.

Appendix A. Cost function details

Here we detail two aspects of the cost function presented in § 3.2 – the representative runaway current and the impact of choice of p -norm.

A.1. Representative runaway current

As described in § 3.2, we considered two plausible ways to give a representative value for the runaway current; the maximum or the value at the time when $I_{\text{re}} = 0.95I_p$. 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 makes the resulting value arbitrary – see figure 9(a).

Using the 95% runaway current would solve this problem, but it is undetermined when $I_{\text{re}}(t) < 0.95I_p(t)$ throughout the simulation (figure 9b) or when $I_{\text{re}}^{95\%}$ for several t

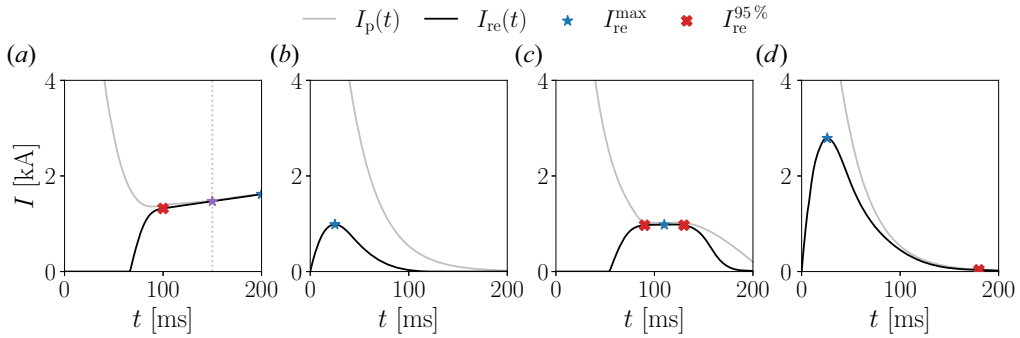


FIGURE 9. (a) A runaway plateau is reached during the disruption, such that the runaway current is slowly increasing and I_{re}^{max} depends on the simulation length. If the simulation is stopped at 200 ms (150 ms), the blue (purple) star marks I_{re}^{max} . (b) Criterion for $I_{re}^{95\%}$ is never fulfilled. (c) Criterion for $I_{re}^{95\%}$ is fulfilled more than once. (d) Runaway current peaks significantly early during the simulation, but reaches 95 % of the total plasma current when both currents are negligible.

(figure 9c). Furthermore, for a hypothetical case where the runaway current near the end of the simulation fulfils $I_{re} = 0.95I_p$ while being small, but peaks early during the simulation with an order ~ 1 MA without reaching 95 % of the plasma current, as in figure 9(d), $I_{re}^{95\%}$ would not be a representative value.

To address all of these shortcomings, we have chosen to use a combination of the maximum and 95 % current as our representative value for the runaway current. More specifically, we will define $I_{re}^{repr} = \min_t \{I_{re}^{max}, I_{re}^{95\%}\}$.

A.2. Choice of p -norm

In § 3.2, the cost function is described as being designed such that if all components are within their intervals of safe operation, the cost function should be smaller than one. This was partly achieved by combining the cost function components using the Euclidean norm. For safety to be equivalent with the cost function $\mathcal{L} \leq 1$, we would need to use the maximum norm instead, i.e.

$$\mathcal{L} = \max \left(\frac{I_{re}^{repr}}{150 \text{ kA}}, \frac{I_{\Omega}^{final}}{150 \text{ kA}}, \frac{\eta_{cond}}{10\%}, \frac{|\tau_{CQ} - 100 \text{ ms}|}{50 \text{ ms}} \right). \quad (\text{A1})$$

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

However, the components could be combined by simply averaging them, i.e.

$$\mathcal{L} = \frac{1}{4} \left(\frac{I_{re}^{repr}}{150 \text{ kA}} + \frac{I_{\Omega}^{final}}{150 \text{ kA}} + \frac{\eta_{cond}}{10\%} + \frac{|\tau_{CQ} - 100 \text{ ms}|}{50 \text{ ms}} \right). \quad (\text{A2})$$

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 four times its safe value with $\mathcal{L} \leq 1$ still satisfied.

Both (A1) and (A2) 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




FIGURE 10. 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 . 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 $1/C$.

normalized component dominating the cost function value and thus losing information about the other parameters, and having unity as the exact border between safe and unsafe parameter choices. For safety to be implied by $\mathcal{L} \leq 1$, the weight of each component must be set as $C = 1/(n)^{1/p}$ for n cost function components. In figure 10, 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. With four components, 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.

Appendix B. Radial current profiles

The current density profiles throughout the simulations of the non-activated and activated optimal scenarios are presented in figure 11. The Ohmic current density profiles in figure 11(b) are non-monotonic or highly peaked in the core, while the runaway current density profiles are hollow. In figure 11(c), especially the runaway current density profiles are strongly peaked in the core. These features suggest that some of the scenarios might become prone to MHD instabilities. Thus, to improve the predictive capability of these simulations, evolving the magnetic equilibrium and monitoring MHD stability would be desired.

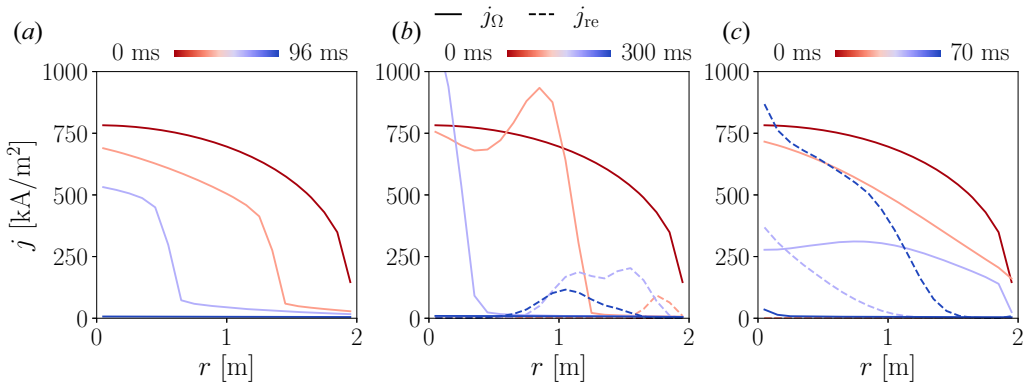


FIGURE 11. Radial profiles of the current density from kinetic simulations of the optimum (a) for non-activated scenarios, and for activated scenarios found using (b) the fluid model and (c) the kinetic model. Both Ohmic (solid) and runaway (dashed) current density profiles are shown.

REFERENCES

- BERGER, E., PUSZTAI, I., NEWTON, S.L., HOPPE, M., VALLHAGEN, O., FIL, A. & FÜLÖP, T. 2022 Runaway dynamics in reactor-scale spherical tokamak disruptions. *J. Plasma Phys.* **88** (6), 905880611.
- BREIZMAN, B.N., ALEJNIKOV, P., HOLLMANN, E.M. & LEHNEN, M. 2019 Physics of runaway electrons in tokamaks. *Nucl. Fusion* **59** (8), 083001.
- BUCHHOLZ, B. 2023 Calculation of the runaway electron current in tokamak disruptions. Master thesis, Helmut-Schmidt-University/ University of the German Federal Armed Forces Hamburg, [arXiv:2309.10827](https://arxiv.org/abs/2309.10827).
- CONNOR, J.W. & HASTIE, R.J. 1975 Relativistic limitations on runaway electrons. *Nucl. Fusion* **15** (3), 415–424.
- FEHÉR, T., SMITH, H.M., FÜLÖP, T. & GÁL, K. 2011 Simulation of runaway electron generation during plasma shutdown by impurity injection in ITER. *Plasma Phys. Control. Fusion* **53** (3), 035014.
- FERMI, E. 1934 Versuch einer Theorie der β -Strahlen. I. *Z. Phys.* **88** (3–4), 161–177.
- GARNETT, R. 2023 *Bayesian Optimization*. Cambridge University Press, stochastic processes: pp. 8, 16–17; Matérn kernel: pp. 51–52; Expected improvement acquisition function: pp. 127–129.
- HELANDER, P., ERIKSSON, L.-G. & ANDERSSON, F. 2002 Runaway acceleration during magnetic reconnection in tokamaks. *Plasma Phys. Control. Fusion* **44** (12B), B247.
- HENDER, T.C., WESLEY, J.C., BIALEK, J., BONDESON, A., BOOZER, A.H., BUTTERY, R.J., GAROFALO, A., GOODMAN, T.P., GRANETZ, R.S., GRIBOV, Y., *et al.* 2007 Chapter 3: MHD stability, operational limits and disruptions. *Nucl. Fusion* **47** (6), S128–S202.
- HESSLOW, L., EMBRÉUS, O., VALLHAGEN, O. & FÜLÖP, T. 2019a Influence of massive material injection on avalanche runaway generation during tokamak disruptions. *Nucl. Fusion* **59** (8), 084004.
- HESSLOW, L., UNNERFELT, L., VALLHAGEN, O., EMBRÉUS, O., HOPPE, M., PAPP, G. & FÜLÖP, T. 2019b Evaluation of the Dreicer runaway generation rate in the presence of high-Z impurities using a neural network. *J. Plasma Phys.* **85** (6), 475850601.
- HOLLMANN, E.M., ALEJNIKOV, P.B., FÜLÖP, T., HUMPHREYS, D.A., IZZO, V.A., LEHNEN, M., LUKASH, V.E., PAPP, G., PAUTASSO, G., SAINT-LAURENT, F., *et al.* 2015 Status of research toward the ITER disruption mitigation system. *Phys. Plasmas* **22** (2), 021802.
- HOPPE, M., EMBRÉUS, O. & FÜLÖP, T. 2021a DREAM: a fluid-kinetic framework for tokamak disruption runaway electron simulations. *Comput. Phys. Commun.* **268**, 108098.
- HOPPE, M., HESSLOW, L., EMBRÉUS, O., UNNERFELT, L., PAPP, G., PUSZTAI, I., FÜLÖP, T., LEXELL, O., LUNT, T., MACUSOVA, E., *et al.* 2021b Spatiotemporal analysis of the runaway distribution

- function from synchrotron images in an ASDEX Upgrade disruption. *J. Plasma Phys.* **87** (1), 855870102.
- HU, D., NARDON, E., HOELZL, M., WIESCHOLLEK, F., LEHNEN, M., HUIJSMANS, G.T.A., VAN VUGT, D.C., KIM, S.-H., JET CONTRIBUTORS & JOREK TEAM 2021 Radiation asymmetry and MHD destabilization during the thermal quench after impurity shattered pellet injection. *Nucl. Fusion* **61** (2), 026015.
- IZZO, V.A., PUSZTAI, I., SÄRKIMÄKI, K., SUNDSTRÖM, A., GARNIER, D.T., WEISBERG, D., TINGUELY, R.A., PAZ-SOLDAN, C., GRANETZ, R.S. & SWEENEY, R. 2022 Runaway electron deconfinement in SPARC and DIII-D by a passive 3D coil. *Nucl. Fusion* **62** (9), 096029.
- KRANE, K.S. 1988 *Introductory Nuclear Physics*. John Wiley & Sons.
- LEHNEN, M. 2021 The ITER disruption mitigation system – design progress and design validation. Presented at *Theory and Simulation of Disruptions Workshop, PPPL*. URL: <https://tsdw.pppl.gov/Talks/2021/Lehnen.pdf>
- MARTÍN-SOLÍS, J.R., LOARTE, A. & LEHNEN, M. 2017 Formation and termination of runaway beams in ITER disruptions. *Nucl. Fusion* **57** (6), 066025.
- MILLER, R.L., CHU, M.S., GREENE, J.M., LIN-LIU, Y.R. & WALTZ, R.E. 1998 Noncircular, finite aspect ratio, local equilibrium model. *Phys. Plasmas* **5** (4), 973–978.
- NOGUEIRA, F. 2014 Bayesian optimization: open source constrained global optimization tool for Python. <https://github.com/fmfn/BayesianOptimization>
- PUSZTAI, I., EKMARK, I., BERGSTRÖM, H., HALLDESTAM, P., JANSSON, P., HOPPE, M., VALLHAGEN, O. & FÜLÖP, T. 2023 Bayesian optimization of massive material injection for disruption mitigation in tokamaks. *J. Plasma Phys.* **89** (2), 905890204.
- PUTVINSKI, S., FUJISAWA, N., POST, D., PUTVINSKAYA, N., ROSENBLUTH, M.N. & WESLEY, J. 1997 Impurity fueling to terminate tokamak discharges. *J. Nucl. Mater.* **241–243**, 316–321.
- RECHESTER, A.B. & ROSENBLUTH, M.N. 1978 Electron heat transport in a tokamak with destroyed magnetic surfaces. *Phys. Rev. Lett.* **40**, 38–41.
- REDL, A., ANGIIONI, C., BELLI, E. & SAUTER, O. 2021 A new set of analytical formulae for the computation of the bootstrap current and the neoclassical conductivity in tokamaks. *Phys. Plasmas* **28** (2), 022502.
- ROSENBLUTH, M.N. & PUTVINSKI, S.V. 1997 Theory for avalanche of runaway electrons in tokamaks. *Nucl. Fusion* **37** (10), 1355.
- SMITH, H.M. & VERWICHTE, E. 2008 Hot tail runaway electron generation in tokamak disruptions. *Phys. Plasmas* **15** (7), 072502.
- SVENNINGSSON, I. 2020 Hot-tail runaway electron generation in cooling fusion plasmas. Master thesis, Chalmers University of Technology, <https://hdl.handle.net/20.500.12380/300899>.
- SVENNINGSSON, I., EMBRÉUS, O., HOPPE, M., NEWTON, S.L. & FÜLÖP, T. 2021 Hot-tail runaway seed landscape during the thermal quench in tokamaks. *Phys. Rev. Lett.* **127** (3), 035001.
- TINGUELY, R.A., IZZO, V.A., GARNIER, D.T., SUNDSTRÖM, A., SÄRKIMÄKI, K., EMBRÉUS, O., FÜLÖP, T., GRANETZ, R.S., HOPPE, M., PUSZTAI, I., *et al.* 2021 Modeling the complete prevention of disruption-generated runaway electron beam formation with a passive 3D coil in SPARC. *Nucl. Fusion* **61** (12), 124003.
- TINGUELY, R.A., PUSZTAI, I., IZZO, V.A., SÄRKIMÄKI, K., FÜLÖP, T., GARNIER, D.T., GRANETZ, R.S., HOPPE, M., PAZ-SOLDAN, C., SUNDSTRÖM, A., *et al.* 2023b On the minimum transport required to passively suppress runaway electrons in sparc disruptions. *Plasma Phys. Control. Fusion* **65** (3), 034002.
- TINGUELY, R.A., PUSZTAI, I., IZZO, V.A., SÄRKIMÄKI, K., FÜLÖP, T., GARNIER, D.T., GRANETZ, R.S., HOPPE, M., PAZ-SOLDAN, C., SUNDSTRÖM, A., *et al.* 2023a On the minimum transport required to passively suppress runaway electrons in SPARC disruptions. *Plasma Phys. Control. Fusion* **65** (3), 034002.
- VALLHAGEN, O., EMBRÉUS, O., PUSZTAI, I., HESSLOW, L. & FÜLÖP, T. 2020 Runaway dynamics in the DT phase of ITER operations in the presence of massive material injection. *J. Plasma Phys.* **86** (4), 475860401.

- VALLHAGEN, O., HANEBRING, L., ARTOLA, F.J., LEHNEN, M., NARDON, E., FÜLÖP, T., HOPPE, M., NEWTON, S.L. & PUSZTAI, I. 2024 Runaway electron dynamics in ITER disruptions with shattered pellet injections. *Nucl. Fusion* (submitted), [arXiv:2401.14167](https://arxiv.org/abs/2401.14167).
- VALLHAGEN, O., PUSZTAI, I., HOPPE, M., NEWTON, S.L. & FÜLÖP, T. 2022 Effect of two-stage shattered pellet injection on tokamak disruptions. *Nucl. Fusion* **62** (11), 112004.
- WALKOWIAK, J., HOPPE, M., EKMARK, I., JARDIN, A., BIELECKI, J., KRÓL, K., SAVOYE-PEYSSON, Y., MAZON, D., DWORAK, D. & SCHOLZ, M. 2024 First numerical analysis of runaway electron generation in tungsten-rich plasmas towards ITER. *Nucl. Fusion* **64** (3), 036024.