



Bayesian optimization of massive material injection for disruption mitigation in tokamaks









Downloaded from: <https://research.chalmers.se>, 2025-12-04 23:24 UTC

Citation for the original published paper (version of record):

Pusztai, I., Ekmark, I., Bergström, H. et al (2023). Bayesian optimization of massive material injection for disruption mitigation in tokamaks. *Journal of Plasma Physics*, 89(2).
<http://dx.doi.org/10.1017/S0022377823000193>

N.B. When citing this work, cite the original published paper.

Bayesian optimization of massive material injection for disruption mitigation in tokamaks

I. Pusztai ^{1,†}, I. Ekmark ¹, H. Bergström ^{1,2}, P. Haldestam ^{1,2},
P. Jansson ³, M. Hoppe ⁴, O. Vallhagen ¹ and T. Fülöp ¹

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

²Max Planck Institute for Plasma Physics, Garching b. M 85748, Germany

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

⁴Swiss Plasma Center, Ecole Polytechnique Fédérale de Lausanne, Lausanne CH-1015, Switzerland

(Received 2 February 2023; revised 24 February 2023; accepted 24 February 2023)

A Bayesian optimization framework is used to investigate scenarios for disruptions mitigated with combined deuterium and neon injection in ITER. The optimization cost function takes into account limits on the maximum runaway current, the transported fraction of the heat loss and the current quench time. The aim is to explore the dependence of the cost function on injected densities, and provide insights into the behaviour of the disruption dynamics for representative scenarios. The simulations are conducted using the numerical framework DREAM (Disruption Runaway Electron Analysis Model). We show that, irrespective of the quantities of the material deposition, multi-megaampere runaway currents will be produced in the deuterium–tritium phase of operations, even in the optimal scenarios. However, the severity of the outcome can be influenced by tailoring the radial profile of the injected material; in particular, if the injected neon is deposited at the edge region it leads to a significant reduction of both the final runaway current and the transported heat losses. The Bayesian approach allows us to map the parameter space efficiently, with more accuracy in favourable parameter regions, thereby providing us with information about the robustness of the optima.

Key words: fusion plasma, runaway electrons, plasma simulation

1. Introduction

One of the threats to reliable tokamak operation are off-normal events known as disruptions, which are induced by a sudden loss of plasma confinement (Boozer 2012). When this occurs, the ensuing heat and particle transport results in a rapid temperature drop – a thermal quench (TQ) – that is accompanied by a decrease in the electrical conductivity of the plasma. The reduced conductivity leads to a decay in plasma current – a

[†] Email address for correspondence: pusztai@chalmers.se

current quench (CQ) – that is counteracted by the induction of an electric field, which may accelerate runaway electrons (REs) to relativistic energies (Breizman *et al.* 2019). The REs could potentially strike the wall and lead to subsurface melting of the wall components.

The plasma current in future devices will be around an order of magnitude higher than in present experiments. Correspondingly, the magnetic energy in the plasma will increase (~ 400 MJ in ITER vs ~ 10 MJ in JET) (Hender *et al.* 2007), along with the kinetic energy, thus the available energy that can be released in a disruption is significantly higher than in present devices. It is therefore essential to develop effective disruption mitigation systems.

An effective disruption mitigation system in a tokamak should limit the exposure of the wall to localized heat losses and to the impact of high current RE beams, and avoid excessive forces on the structure (Hollmann *et al.* 2015). To avoid damage to the first wall on ITER, at least 90 % of the thermal energy loss must be lost in the form of radiation. The RE current should be kept below 150 kA in order to avoid melting of plasma facing components, in the case of localized loss (Lehnen & the ITER DMS task force 2021). The CQ time, i.e. the time it takes for the ohmic component of the current to decay, should be kept between 50 and 150 ms. Current quench times below 50 ms will lead to excessive forces due to eddy currents in the structures surrounding the plasma. On the other hand, CQ times above 150 ms are expected to lead to intolerably large halo currents in plasma facing components.

In ITER, the envisaged disruption mitigation system is based on massive material injection (Lehnen *et al.* 2015). The injected material can radiate away a large fraction of the thermal energy and it can also inhibit RE generation by increasing the critical energy for electron runaway. Furthermore, it can also be used to control the temperature during the CQ, which directly influences the CQ duration. However, the question of what mixture of material should be injected, and how it should be deposited, to accommodate all requirements on the disruption mitigation system simultaneously, if it is at all possible, is still open.

In this paper, we describe a Bayesian optimization framework applied to simulations of ITER-like disruption scenarios mitigated with combined injection of deuterium and neon. The aim is to find the injected material quantities and deposition profiles for which the outcome of the disruption is tolerable with respect to the expected RE current, transported heat fraction and CQ time. Bayesian optimization has several attractive features: it does not rely on gradient information, it can handle non-deterministic (noisy) functions and it is suitable for relatively high-dimensional optimization problems and computationally expensive function evaluations. However, its main advantage concerning the current study is that it informs us about the properties of promising parameter regions – in particular the robustness of the optima to variations in the control parameters.

The rest of the paper is structured as follows. The methods are explained in § 2, detailing the set-up of the disruption simulations in 2.1 and the Bayesian optimization in 2.2. The results are presented in § 3, first mapping out the optimization landscape with constant injected densities in § 3.1 followed by a detailed analysis of some representative scenarios in § 3.2. Then, we present optimization results allowing radially varying injection in § 3.3. Finally, we study the parametric sensitivities of the optima and reflect on the beneficial effects of radial profile variations in § 3.4, before we conclude and discuss our findings in § 4.

2. Bayesian optimization of simulated disruptions

We employ an open source Bayesian optimization routine that treats the disruption simulations as a black-box function that produces a single scalar output, the cost function, and accepts inputs for injected material densities and deposition profiles in specified

ranges – these are the input parameters that we want to optimize. In the following we will discuss the disruption simulations, and provide details of the optimization algorithm.

2.1. Simulation set-up

The disruption simulations assume an initially ($t < 0$) pure fully ionized deuterium–tritium (D–T) plasma with 50–50 % isotope concentrations. Specifically, the initial electron density is spatially constant 10^{20} m^{-3} , the temperature is parabolic with 20 keV on axis and the total plasma current is 15 MA. The simulations use an ITER-like magnetic geometry with major radius $R_0 = 6 \text{ m}$, minor radius $a = 2 \text{ m}$, wall radius $b = 2.833 \text{ m}$, on-axis toroidal magnetic field $B(r = 0) = 5.3 \text{ T}$ and a resistive wall time of $\tau_w = 0.5 \text{ s}$, as well as a Miller model equilibrium (Miller *et al.* 1998) with realistic, radially varying shaping parameters; further information is given in [Appendix A](#).

The simulations are performed by the DREAM (Disruption Runaway Electron Analysis Model) code that captures the particle acceleration and energy dissipation processes following a disruption (Hoppe, Embreus & Fülöp 2021). It solves a set of coupled transport equations describing the evolution of temperature, ion charge state densities, current density and electric field in arbitrary axisymmetric geometry. The temperature evolution includes ohmic heating, radiated power using atomic rate coefficients, collisional energy transfer from hot electrons and ions, as well as dilution cooling.

DREAM allows modelling of the REs at different degrees of approximation ranging from fluid to fully kinetic. As we do not require kinetic outputs, we limit our modelling to the least computationally expensive, fluid treatment of the plasma. This means that the thermal bulk of cold electrons and the small runaway population are modelled as two separate fluid species. The former is characterized by a density n_e , a temperature T_e as well as an ohmic current density j_{ohm} , and the REs are described by their density n_{RE} . It is assumed that the REs move with the speed of light parallel to the magnetic field, hence their associated current density is $j_{\text{RE}} = ec n_{\text{RE}}$. The simulations include Dreicer, hot-tail and avalanche sources, as well as REs generated by Compton scattering of γ photons and tritium decay. These are modelled as quasi-stationary sources feeding electrons into the runaway population (Fülöp *et al.* 2020). The runaway generation rates used in the simulations have been benchmarked with the corresponding kinetic results (Hoppe *et al.* 2021). Further details on the simulations are given in [Appendix A](#).

Neutral neon and deuterium are introduced with zero temperature at the start of the simulation ($t = 0$). At the same time an elevated transport of electron heat and energetic electrons is activated, using a Rechester–Rosenbluth-type model (Rechester & Rosenbluth 1978) with a radially constant normalized magnetic perturbation amplitude $\delta B/B$. This is done to emulate the break-up of flux surfaces during the TQ, and leads to heat losses, with a heat diffusivity proportional to $R_0 v_{te} (\delta B/B)^2$, where $v_{te} = \sqrt{2T_e/m_e}$ is the local electron thermal speed. The full expression is given by (B.5) of Hoppe *et al.* (2021). In four optimization runs $\delta B/B$ is scanned over the range 0.2 %–0.5 % that falls within the range of values observed in magnetohydrodynamic simulations of the TQ (Hu *et al.* 2021).

During the TQ we also account for a diffusive transport of REs using a diffusion coefficient of similar form, but assuming a parallel streaming along the perturbed field lines at the speed of light, $D_{\text{RE}} = \pi R_0 c (\delta B/B)^2$. This approach neglects the momentum space variation of the transport coefficients (Särkimäki *et al.* 2020), as well as the form of the RE distribution function, which would reduce the effect of runaway transport. Thus, using this expression provides an upper bound on the effect of runaway transport for a given magnetic perturbation amplitude (Svensson *et al.* 2021). We employ here the same $\delta B/B$ as for electron heat transport for consistency.

The injected material is ionized by its interaction with the plasma, and cools it by radiation and dilution. When the average electron temperature falls below 10^{-3} times the maximum initial temperature (here 20 eV), we assume that the TQ is completed and the flux surfaces reform. After the TQ the transport of energetic electrons is switched off, and a significantly reduced, but finite electron heat diffusivity is used ($\delta B/B = 0.04\%$). This is to avoid the development of non-physical narrow hot ohmic channels during the CQ. Such ohmic channels are soliton-like solutions of the problem (Putvinski *et al.* 1997) without sufficient heat diffusivity. In a physical system the corresponding excessive temperature and current gradients would be expected to destabilize these formations well before they could fully form. Note, that the diffusive heat transport is subdominant compared with radiative heat losses at the low post-TQ temperatures, thus this heat transport has no effect besides not allowing hot channels to form.

2.2. Optimization

The optimization problem involves multiple objectives, i.e. multiple quantities need to be within certain limits simultaneously. The maximum value of the total RE current and the fraction of transported heat losses must be small, while the CQ time should be within certain limits. These quantities are normalized and combined into a single scalar cost function $\mathcal{L} \geq 0$ which is to be minimized. Denoting the control vector containing the parameters by \mathbf{x} , we wish to find the \mathbf{x}^* that minimizes \mathcal{L} , where \mathbf{x} resides in a specified volume $\mathcal{V} \subset \mathbb{R}^d$ of the control space (where d is the dimensionality of the optimization).

We employ Bayesian optimization (Brochu, Cora & de Freitas 2010) using Gaussian process regression (Rasmussen & Williams 2005), using the Bayesian Optimization (Nogueira 2014) Python package. A Gaussian process is fitted to the already sampled points $\{\mathbf{x}_i\}_{i=1}^n$, and the expected improvement acquisition strategy (described in Appendix B) is used to choose the next point to be sampled, \mathbf{x}_{n+1} . The Gaussian process contains information on both the expected value $\mu(\mathbf{x})$ and the uncertainty of the estimate of \mathcal{L} , quantified in terms of the covariance $k(\mathbf{x}, \mathbf{x}')$ between any two points \mathbf{x} and \mathbf{x}' . In this process there is a balance between exploration and exploitation, i.e. search within regions with high uncertainty, as well as in regions that are most likely to host the global optimum.

The cost function we use is of the form

$$\mathcal{L} = \frac{I_{\text{RE}}^{\text{max}}}{I_{\text{RE}}^{\text{tol}}} + \frac{I_{\text{ohm}}^{\text{fin}}}{I_{\text{ohm}}^{\text{tol}}} + 10 \frac{\eta_{\text{cond}}}{\eta_{\text{cond}}^{\text{tol}}} + 100 \theta(t_{\text{CQ}}), \quad (2.1)$$

where $I_{\text{RE}}^{\text{max}}$ is the maximum RE current in the simulation, $I_{\text{RE}}^{\text{tol}} = 150$ kA represents the tolerable RE current in ITER and $I_{\text{ohm}}^{\text{fin}}$ is the ohmic current at the end of the (150 ms long) simulation. A significant remnant ohmic current may be the sign of an incomplete TQ, and it can still potentially be converted to a RE current. Thus it is treated on equal footing with the RE current, so we also set $I_{\text{ohm}}^{\text{tol}} = 150$ kA. Here, $\eta_{\text{cond}}^{\text{tol}} = 0.1$ is the tolerable transported heat loss fraction. The prefactor 10 in the η_{cond} term is used to get a penalty for non-tolerable transported heat losses comparable to typical penalties obtained for mega-ampere (MA) size currents. Finally, to penalize CQ times t_{CQ} below $t_{\text{L}} = 50$ ms and above $t_{\text{U}} = 150$ ms we use the penalty function

$$\theta(t_{\text{CQ}}) = \tilde{\Theta}(t_{\text{L}} - t_{\text{CQ}}) + \tilde{\Theta}(t_{\text{CQ}} - t_{\text{U}}), \quad (2.2)$$

where $\tilde{\Theta}(t) = \frac{1}{2}[1 + \tanh(t/\Delta t)]$ is a function similar to a step function but smooth with a transition width set by $\Delta t = 3.3$ ms. Values of t_{CQ} outside the tolerable range yield a penalty as high as the maximum achievable penalty for any of the other terms in (2.1), due

to the prefactor 100 in front of θ . We calculate the CQ time as $t_{\text{CQ}} = [t(I_{\text{ohm}} = 0.2I_p^0) - t(I_{\text{ohm}} = 0.8I_p^0)]/0.6$ (Hender *et al.* 2007), where $I_{\text{ohm}}(t)$ is the total ohmic current and I_p^0 is the initial plasma current.

In addition we set $\mathcal{L} = 500$ for simulations where the TQ is not complete within 20 ms, our condition for which is that the average temperature falls to below $10^{-3}T_e(r=0, t=0)$. Finally, as it is difficult to completely avoid simulations that fail due to numerical issues, we use $\mathcal{L} = 500$ for these as well.

Independently of their dimensionality, the optimizations use 400 samples, chosen according to the acquisition function through sequential function evaluations, following 20 randomly selected initial samples. As the parameter space of injected quantities ranges across multiple orders of magnitude, the logarithm of the injected quantities is used as optimization parameters.

3. Bayesian optimization of disruption mitigation with material injection

The goal is to identify what densities of injected neon and deuterium produce the most favourable outcomes in a disruption mitigation, corresponding to the minimum of the cost function. Modelling the details of the material injection is outside the scope of the present work, instead, we assume the material to be instantaneously deposited in the form of neutrals, either uniformly distributed over the magnetic flux surfaces, described in §§ 3.1–3.2 or with radially varying distribution, described in § 3.3.

3.1. Optimization landscape with constant concentrations

First, we perform optimization in the two-dimensional (2-D) parameter space of radially constant injected deuterium and neon densities, $n_{\text{D},\text{inj}}$ and $n_{\text{Ne},\text{inj}}$. The ranges of injected densities we consider are $n_{\text{D},\text{inj}} \in [10^{18}, 3.16 \times 10^{22}] \text{ m}^{-3}$, and $n_{\text{Ne},\text{inj}} \in [10^{16}, 10^{20}] \text{ m}^{-3}$.

Figure 1 shows the estimated mean of the cost function μ on a logarithmic contour plot for four different values of $\delta B/B$, with blue shades representing favourable and red shades unfavourable values. Each panel used 420 samples, indicated by grey dots, while the optima are indicated with black stars. The area of favourable values (with blue shades) decreases with increasing $\delta B/B$, and this is mostly due to the increasing transported heat fraction, and to a lesser degree to an increasing RE current, to be discussed further in relation to figure 5. In general, the lower left corner of the plots is occupied by cases with an incomplete TQ. In this case the plasma tends to get reheated after the prescribed transport event, leading to long CQ times (i.e. $t_{\text{CQ}} > 150$ ms). With an increasing $\delta B/B$ the incomplete TQ region shrinks somewhat.

Another general feature is a relatively narrow corridor of favourable parameters in the vicinity of $n_{\text{D},\text{inj}} = 10^{22} \text{ m}^{-3}$, extending from the lowest $n_{\text{Ne},\text{inj}}$ values plotted to a bit above $n_{\text{Ne},\text{inj}} = 10^{19} \text{ m}^{-3}$. The optima also reside in these corridors at $n_{\text{Ne},\text{inj}}$ values of a few times 10^{18} m^{-3} . Additionally, a wider corridor of moderate values of \mathcal{L} extends to the left of the optima, in the range $n_{\text{Ne},\text{inj}} \approx 3 \times 10^{18} - 2 \times 10^{19} \text{ m}^{-3}$, which is most pronounced in the $\delta B/B = 0.3\%$ and 0.4% cases. At $n_{\text{Ne},\text{inj}}$ values above that, \mathcal{L} increases, and decreases again around the highest $n_{\text{Ne},\text{inj}}$ values included in this optimization. Before analysing the characteristic behaviours in these different regions of the current optimization landscape in § 3.2, we shall discuss the detailed dynamics at the optima.

We consider the behaviour of the representative optimum obtained in the $\delta B/B = 0.3\%$ case (indicated by the black star in figure 1b), located at $n_{\text{D},\text{inj}} = 9.4 \times 10^{21} \text{ m}^{-3}$ and $n_{\text{Ne},\text{inj}} = 2.9 \times 10^{18} \text{ m}^{-3}$, shown in figure 2. Following the instantaneous material injection at $t = 0$, the temperature profile drops by a factor of ≈ 100 within a μs , due to dilution. This is followed by an approximately exponential cooling with a characteristic time of




FIGURE 1. The estimated mean of the cost function of Bayesian optimizations in the $n_{D,inj}$ - $n_{Ne,inj}$ space for various normalized magnetic perturbation amplitudes. The colour code varies from blue to red tones, representing favourable and unfavourable values of μ ; (a) $\delta B/B = 0.2\%$, (b) 0.3% , (c) 0.4% , (d) 0.5% . Black stars indicate the locations of the optima. Grey dots show the samples taken; note that these are more numerous in the vicinity of the optima. Circles with case identifiers in panel (b) indicate the cases discussed in § 3.2.

$\tau \approx 1.5$ ms. After the initial exponential cooling, a cold front starts to propagate radially inward from the edge. This inward propagating cooling is seen in the $t = 2$ ms curve in figure 2(b). This cooling proceeds until almost the entire plasma settles at around 5 eV (see the $t = 10$ ms curve), representing the equilibrium between ohmic heating and radiation corresponding to the ion composition and current density of the plasma. Then there is another inward propagating cooling happening over the next 50 ms. As the ohmic current density drops during the CQ, the equilibrium temperature falls from ≈ 5 to ≈ 1.2 eV. The temperature is radially uniform at this level at 60 ms (black curve), and remains there until the end of the simulation.

The ohmic current density gradually decreases in the core and it drops rapidly across the cold front at the edge; compare the 10 ms curves in figure 2(b,c). This front propagates inward in the first 40 ms, after which the ohmic current gets rapidly replaced by RE current; see the process in terms of total current components in figure 2(a), and RE current density




FIGURE 2. The best performing case for the optimization in the $n_{D,inj}$ – $n_{Ne,inj}$ space, for $\delta B/B = 0.3\%$. (a) The time evolution of the total plasma current (dashed), and its ohmic (solid) and RE (dash-dotted) components. (b–d) Radial profiles of quantities in a few time points, indicated by their respective figure legends; with increasing time corresponding to darker colours. (b) Electron temperature. (c) Ohmic (solid) and RE (dashed) current density. (d) Parallel electric field (solid). The effective critical electric field is also indicated for $t = 60$ ms (dotted); note that it does not vary appreciably over time.

at 60 ms in [figure 2\(c\)](#) (dashed line). The electric field exceeds the effective critical field E_c^{eff} – calculated as in appendix C2 of [Hoppe *et al.* \(2021\)](#) – first in the edge, then it grows to an approximately radially constant value around 30 V m^{-1} , approximately 4 times E_c^{eff} , where it stays until the macroscopic RE conversion starts. Then it drops into the vicinity of E_c^{eff} , such that in the core $E_{||}$ is pinned to E_c^{eff} , and it takes radially decreasing values at the edge; compare $E_{||}$ (solid black curve) with E_c^{eff} (dotted) in [figure 2\(d\)](#). Then the electric field remains like that until most of the RE current dissipates away. Physically, the dissipation of the RE current, in the absence of transport losses, is caused by a collisional slowing down and thermalization of the REs when $E_{||} < E_c^{eff}$. In the employed fluid RE model it is technically accounted for by allowing the avalanche growth rate to become negative for $E_{||} < E_c^{eff}$ values. The corresponding decay of the RE current is quite pronounced in this case.

3.2. Characteristic cases with constant concentrations

In order to understand the typical dynamics in various regions in the $n_{D,inj}$ – $n_{Ne,inj}$ space, we consider six representative cases in the $\delta B/B = 0.3\%$ optimization, with case C1 being the optimum discussed above. The cases are indicated in [figure 1\(b\)](#) and corresponding injected quantities and figures of merit are listed in [table 1](#). Cases C2 and C3 are taken in the high $n_{D,inj}$ region of the space; C2 is located in the favourable channel at low $n_{Ne,inj}$, and C3 at even higher $n_{D,inj}$ than the optimum. Cases C4 to C6 are taken at a fixed $n_{D,inj} = 10^{20} \text{ m}^{-3}$, at respectively increasing value of $n_{Ne,inj}$. We discuss C1–C3 and C4–C6 in the following subsections.

Case ID	$n_{D,\text{inj}}$ (10^{20} m^{-3})	$n_{\text{Ne},\text{inj}}$ (10^{18} m^{-3})	$I_{\text{RE}}^{\text{max}}$ (MA)	$I_{\text{ohm}}^{\text{fin}}$ (MA)	t_{CQ} (ms)	η_{cond} (%)	\mathcal{L}
C1	93.9	2.88	4.2	0.33	59	8.9	39
C2	160	0.032	4.8	0.33	54	43	88
C3	316	2.88	8.2	0.0007	5	1.4	156
C4	1	5.01	6.3	0.059	88	80	122
C5	1	31.6	8.1	0.092	26	72	226
C6	1	100	8.9	0.159	15	23	163

TABLE 1. Characteristic cases from the $n_{D,\text{inj}}-n_{\text{Ne},\text{inj}}$ optimization landscape for $\delta B/B = 0.3\%$, their four figures of merit and corresponding cost function values. The cases are marked in [figure 1\(b\)](#); C1 is the optimum.




FIGURE 3. Time evolution of quantities of interest for the high $n_{D,\text{inj}}$ representative cases: C1–C3. Line colour darkens and dashes shorten with increasing case number, and case numbers are indicated with callouts. (a) Runaway electron current. (b) Electron temperature at mid-radius. (c) Electric field normalized to critical electric field at mid-radius. (Note the longer time range plotted in panel a.)

3.2.1. Representative cases at high $n_{D,\text{inj}}$

In the RE plateau the electric field tends to stay close to the effective field E_c^{eff} , as expected (Breizman 2014). In particular, following the RE conversion, all $E_{||}/E_c^{\text{eff}}$ values (taken at mid-radius) settle around unity, as shown in [figure 3\(c\)](#). The high $n_{D,\text{inj}}$ cases are all characterized by a significant decay rate of the RE current after it reaches its maximum value; see [figure 3\(a\)](#). This is consistent with their $E_{||}/E_c^{\text{eff}}$ being lower than unity towards the edge, as we have seen for C1 in [figure 2\(d\)](#). This is due to the relatively high value of E_c^{eff} typical at these high $n_{D,\text{inj}}$ values.

The dynamics of the RE current in C1 and C2 is fairly similar, as seen in [figure 3\(a\)](#). It may be surprising that $E_{||}/E_c^{\text{eff}}$ is almost all the time higher in C1 than C2 – shown in [figure 3\(c\)](#) – but the maximum RE current in C2 is still higher, and is reached a bit earlier. The reason for this is that the temperature drops to 1.09 eV in C2 already at $t \approx 10$ ms ([figure 3\(b\)](#)), a temperature where 44 % of the hydrogenic species is recombined, thereby increasing the total-to-free electron density ratio and the avalanche growth rate in proportion. Meanwhile in C1 the temperature does not drop to this low temperature until the RE conversion is over.

The effect of the hydrogen recombination is even more pronounced in C3, where the temperature drops to 1.02 eV within a millisecond. The reason for this fast cooling is the very high dilution that brings the temperature down to a range where radiative losses are strong and can effectively (and rapidly) cool the plasma further. That the temperature drops immediately to its final value, without stopping at some higher, intermediate value,

can be explained by how the temperature dependence of the total radiative losses (P) is affected by the very high hydrogen content. Depending on the hydrogen (including D and T) and neon densities, the curve $P(T_e)$ can exhibit a local minimum in the few eV range between a low T_e peak caused by hydrogen and a higher T_e peak from neon. The large hydrogen density in C3 leads to an elevated value of P at this minimum, thereby effectively eliminating the bottleneck this minimum represents concerning the cooling. While 1.02 eV is just slightly cooler than the final temperature in C2, now 70 % of the hydrogenic species are recombined, which, in combination with the early high value of $E_{\parallel}/E_c^{\text{eff}}$, leads to an extremely fast RE conversion and the highest RE current among these three cases.

In terms of figures of merit, C3 is not only problematic due to a high $I_{\text{RE}}^{\text{max}}$ value, but also because of the extremely short $t_{\text{CQ}} \approx 5$ ms. While $I_{\text{RE}}^{\text{max}}$ is not too much higher in C2 than in C1, it has a $\eta_{\text{cond}} \approx 44$ %, exceeding the tolerable 10 %, unlike C1 and C3. This is due to the small neon content in C2.

The remarkably short cooling times, of the order of 2 ms, observed at large deuterium injections, such as C3, may be partly due to our simplifying assumption of instantaneous deposition. However, in realistic material injection scenarios, the cooling at a given flux surface can be as rapid as observed here, even if the time scale needed for pellet shards flying at 500 m s^{-1} to travel between the edge and the centre of an ITER plasma is longer (≈ 4 ms). As the local cooling time is the crucial factor to get a large hot-tail seed, and furthermore, the rapid avalanche rate depends on the final temperature, similar behaviour is also observed in shattered pellet injection simulations (Vallhagen *et al.* 2022). Ion convection time scales across the radius in a TQ can also be in the ms range. The excessive runaway generation is thus not an artefact of the instantaneous deposition, however, the detailed temperature evolution is expected to be different once the injection dynamics is resolved.

3.2.2. Representative cases at low $n_{\text{D},\text{inj}}$

The cases at $n_{\text{D},\text{inj}} = 10^{20} \text{ m}^{-3}$ – C4 to C6 – are not affected by hydrogen recombination as their temperature never drops below 2 eV. They reach much higher values of $E_{\parallel}/E_c^{\text{eff}}$ than the high $n_{\text{D},\text{inj}}$ cases, as they have low E_c^{eff} ; compare figures 3(c) and 4(c). Their RE conversion timing and magnitude well correlates with when the peak of $E_{\parallel}/E_c^{\text{eff}}$ is reached, and its magnitude. This, in turn, depends on the first equilibrium temperature reached, varying between approximately 5 and 11 eV, see figure 4(b). This temperature decreases monotonically with increasing injected neon quantity, while the magnitude of the final RE current increases, and the time of RE conversion shifts earlier. Once the conversion is complete, the temperature falls further into the 2–4 eV range. Note that at these low $n_{\text{D},\text{inj}}$ cases the dissipation rate of the RE current in the RE plateau is negligible during the simulation, due to the lower E_c^{eff} values.

Only in C4 does t_{CQ} fall into the acceptable range, in the other two cases it is too short due to the early RE conversion. The reason for the non-monotonic dependence of \mathcal{L} with increasing $n_{\text{Ne},\text{inj}}$, i.e. it is higher for C5 than for C6, is that it is caused by the reduction in the transported heat loss fraction from the 70 %–80 % range to 23 % (that is still not acceptably low though).

3.3. Radially varying material injection

Next, we relax the assumption of spatially homogeneous injection, and allow profile variations with a simple model for the injected densities, where the inward or outward




FIGURE 4. Time evolution of quantities of interest for the low $n_{D,inj}$ representative cases: C4–C6. Line colour darkens and dashes shorten with increasing case number, and case numbers are indicated with callouts. (a) Runaway electron current. (b) Electron temperature at mid-radius. (c) Electric field normalized to critical electric field at mid-radius. (Note the longer time range plotted in panel a.)

peaking of the profile is set by a single parameter c_i per species i

$$\tilde{n}_{i,inj} \propto 1 + \tanh \left[c_i \left(\frac{r}{a} - \frac{1}{2} \right) \right], \quad (3.1)$$

where the tilde indicates that $\tilde{n}_{i,inj}$ is a radially varying quantity. The notation $n_{i,inj}$ is reserved to the scalar parameter that appears in the optimization. The factor multiplying the expression in (3.1) is determined such that the total number of injected particles in the plasma is the same as in an injection of a constant density $n_{i,inj}$. Negative/positive values of c_i correspond to densities peaked in the plasma centre/edge, and in the optimization we allow values in the $[-10, 10]$ range.

Figure 5 shows the I_{RE}^{max} and η_{cond} figures of merit, along with the cost function \mathcal{L} at the optima found for different $\delta B/B$ values, when radially constant injection is employed (dotted line, referred to as 2D) and when profile variation is allowed (dashed, 4D). In the latter case, the additional degrees of freedom allow us to find optima with better properties. Since in all cases the remaining ohmic current is much smaller than I_{RE}^{max} (in the 300–400 kA range), and t_{CQ} is also in the tolerable range, \mathcal{L} is dominated by the two figures of merit plotted. In none of the cases considered is I_{RE}^{max} tolerably small; it is around 4 MA independently of $\delta B/B$ in the 2-D optimization, and it reduces almost by a factor of 2 in four dimensions (without any clear trend with $\delta B/B$), as seen in figure 5(a).

There are two main reasons for obtaining such high values even in the optimal cases. We consider D–T plasmas, and we include RE seed sources relevant for activated operation, tritium decay and Compton scattering of γ photons, in addition to Dreicer and hot-tail RE generation. The tritium decay and Compton sources can provide a significant RE seed even after the TQ, during which the transport due to magnetic perturbations decimates the initial hot-tail and Dreicer seed population. This circumstance also explains the weak sensitivity of I_{RE}^{max} to $\delta B/B$ in the 2-D simulations. A simulation identical to the 2-D optimum at $\delta B/B = 0.3\%$, but without activated seed sources (i.e. only Dreicer, hot-tail and avalanche sources active) yields a negligibly small $I_{RE}^{max} = 4.1$ kA instead of 4.2 MA.

A similarly important factor is the realistic radius of the conducting wall, which is chosen to match the energy in the poloidal magnetic field due to the plasma current within the conducting wall to that observed in JOREK simulations. If in the 2-D optimum at $\delta B/B = 0.3\%$ we reduce the wall radius from 2.833 to 2.15 m, which was used in previous work, e.g. by Vallhagen *et al.* (2020), the RE current reduces to the – non-negligible, but still significantly lower – value of 1 MA.




FIGURE 5. Variation of (a) the maximum RE current, (b) the transported heat loss fraction and (c) the corresponding cost function in optimizations, for a range of $\delta B/B$ values, when optimizing only for injected densities (circle markers, blue short dashed curve) and when including profile variation as well in the optimization (squares, red long dashed). In panels (a,b), below the thin solid line the values are considered tolerable. In panel (a) simulations with the parameters corresponding to the 2-D optimum at $\delta B/B = 0.3\%$ but without activated sources is indicated with a black rectangle marker, and a simulation with a reduced wall radius of 2.15 m is shown with a black asterisk.

The fraction of transported heat losses, shown in figure 5(b), increases strongly with $\delta B/B$ in the 2-D cases, which is not surprising, since the heat transport during the TQ is then increasing, while the radiated losses are not directly impacted by $\delta B/B$. However, when profile variation is allowed η_{cond} is almost independent of $\delta B/B$; the reason for this will be explained in relation to figure 7.

3.4. Sensitivity of the optima

To gauge the sensitivity of the optima to the input parameters, we investigate the regions occupied by samples within some range of \mathcal{L} above the optimal values. The location of the optima in the optimization space is marked in figure 6 (\otimes markers). In the 2-D optimization study we also scatter plot all samples in the 10% vicinity of the optimum, figure 6(a); this is such a narrow range in \mathcal{L} , that any point in this point cloud can be considered equally well performing as the optimum itself. In the 4-D optimization study we show points in the 25% vicinity of the optima, figure 6(b–d). As the total number of samples is the same in both the 2-D and 4-D optimization studies, the higher dimensionality in four dimensions implies a sparser exploration in the vicinity of the optimum compared with two dimensions; hence the lower number of points in spite of the wider relative range included.

First, considering the 2-D optimization, figure 6(a), we find that the relative extent of the point clouds is significantly larger in the $n_{Ne,inj}$ direction, than in the $n_{D,inj}$ direction; for instance in the $\delta B/B = 0.2\%$ case $n_{Ne,inj}$ spans more than an order of magnitude, while $n_{D,inj}$ spans only a bit more than a factor of two. In practice it translates to the need of a higher precision concerning the injected amount of deuterium than that of neon. The negative correlation between $n_{D,inj}$ and $n_{Ne,inj}$ seen from the arrangement of the point cloud indicates that there are similarities in the effects of these two injected species. These features are also reflected in the favourable valleys (blue tone regions) seen in figure 1. The favourable parameter range indicated by the point clouds shrinks with $\delta B/B$. Note that the region covered by the optima at different $\delta B/B$ values is even smaller than the smallest (black) point cloud; thus we should not read much into how the actual location of the optima varies with $\delta B/B$.

In the 4-D optimization, the resulting point clouds are more scattered, when projected into the $n_{D,inj}$ – $n_{Ne,inj}$ subspace, see figure 6(b). If anything, there is still a weak




FIGURE 6. Scatter plot of input parameters for samples with the lowest \mathcal{L} values in each optimization case. When (a) optimizing only for injected densities (two dimensions) they represent an additional 10 % range above the optimum, and when (b–d) including profile variation as well in the optimization (four dimensions), they represent a 25 % range. Darkening colour indicates increasing value of $\delta B/B$, as given in panel (b), and the optima are indicated by \otimes markers. (a,b) Concentration space, (c,d) correlating concentration with profile parameter of an injected species. Note that in the 4-D $\delta B/B = 0.3$ % case there is no sample within the 25 % range above the optimum.

anti-correlation between the injected quantities, but the poor statistics makes it less clear. Similarly to the 2-D optimization, the range covered in $n_{\text{D},\text{inj}}$ is smaller than that in $n_{\text{Ne},\text{inj}}$. We can also see that there are no cases within a relative range of 25 % of the optimum for $\delta B/B = 0.3$ %. In addition, the optimum itself appears far in the parameter space from the other three overlapping clouds. Namely, it appears at the highest $n_{\text{D},\text{inj}}$ and lowest $n_{\text{Ne},\text{inj}}$ values. We omit this outlier case in the following discussion, but will return to it at the end of this section.

The point clouds occupy the relatively narrow $c_D \in [-1.5, 1.2]$ range, as seen in figure 6(c), corresponding to modest profile variation. We find a positive correlation between $n_{\text{D},\text{inj}}$ and c_D . It means that higher injected content corresponds to more edge-localized peaking. In particular, the injected densities at the plasma centre occupy a narrower range than at the edge (see table 2); apparently, the deuterium density value at the edge is less important. We also observe that lower $\delta B/B$ corresponds to higher c_D and $n_{\text{D},\text{inj}}$ values.

For the injected neon profiles, a strong outward peaking is preferred, with values of $c_{\text{Ne}} \in [5, 10]$, as seen in figure 6(d). The total injected quantities are typically higher than those in the 2-D optimization, covering mostly the $n_{\text{Ne},\text{inj}} \in [10^{19}, 10^{20}] \text{ m}^{-3}$ range – an order of magnitude higher than in two dimensions. It is interesting to note that, similarly to deuterium, there is a positive correlation between c_{Ne} and $n_{\text{Ne},\text{inj}}$.

$\delta B/B$	$n_D(0)$ (10^{20} m^{-3})	$n_D(a)$ (10^{20} m^{-3})	$n_{Ne}(0)$ (10^{18} m^{-3})	$n_{Ne}(a)$ (10^{18} m^{-3})
0.2 %	80	160	0.14	62
0.3 %	55	320	0.19	3×10^{-4}
0.4 %	87	110	0.10	20
0.5 %	140	60	0.02	54

TABLE 2. Total hydrogenic (including the background) and neon densities at the plasma centre ($r = 0$) and at the edge ($r = a$) in the 4-D optimization in the various $\delta B/B$ cases.




FIGURE 7. Comparison of the optimal cases in the 2-D (dashed curves) and the 4-D (solid curves) optimization for $\delta B/B = 0.5\%$. (a) Radial total hydrogenic density, $n_{D+T+D,inj}$ (blue), and neon density, n_{Ne} (red). (b) The RE current density profiles taken at the time point when the total RE current takes its maximum, $t = 42$ ms (50 ms) in the 2-D (4-D) case. (c) Time evolution of the heat loss power in the first millisecond, when most of the thermal energy is lost from the plasmas (note the log scale). Blue curves represent the transported heat losses, red curves are the radiated losses.

To understand why the optima in the 4-D optimization perform better than those of two dimensions, we compare the respective $\delta B/B = 0.5\%$ cases, where the figures of merit are most disparate. The hydrogenic (blue curves) and neon (red) density profiles of the 2-D (dashed curves) and 4-D (solid) optima, are shown in figure 7(a). For deuterium, the 4-D optimization finds a moderate inward peaking ($c_D = -0.92$), while the neon profile is strongly peaked at the edge ($c_{Ne} = 8.67$), covering a density range over three orders of magnitude.

The neon content has two major effects on our figures of merit. An increasing neon concentration corresponds to a lower quasi-equilibrium temperature during the RE conversion, typically leading to higher final RE currents. This is the same trend that we have witnessed moving from C4 to C6. At the same time, a higher neon concentration can help increase the radiated fraction of heat losses (this was also clear when comparing C5 with C6). However, the neon concentration affects the final RE current most strongly where the RE growth is strongest. This happens to be the plasma core in the parameter region of interest, without a radial variation of the neon density. In addition, to achieve a low η_{cond} value it is sufficient to have enough radiating impurities in the edge. Both requirements can be satisfied by an outward peaking neon concentration, which is indeed what the 4-D optima tend to develop.

We find that the 2-D optimum produces a centrally peaked RE current, as seen in figure 7(b) (dashed curve), while the 4-D optimum has a RE profile peaked off axis (solid curve), as expected for the low core concentration of neon. We note that, in this case, the runaway and the centrally peaked ohmic currents decay together after the RE current

reaches its maximum, and only towards the end of the simulation (≈ 140 ms) does the total current profile become truly hollow.¹

The time evolution of the volume-integrated heat losses is shown in [figure 7\(c\)](#) in the first millisecond. This is when the vast majority of the thermal energy content of the plasma is lost, while the fraction of magnetic-to-thermal energy conversion is still negligible. Again, the dashed curves correspond to the 2-D optimum; in this case, the transported loss (dashed blue) reaches comparable values to the radiated losses (dashed red). The entire energy loss process varies relatively smoothly over the plotted time range. In contrast, in the 4-D optimum case the transported heat losses (solid blue) are approximately two orders of magnitude lower than the radiated losses (solid red), and both of these channels have a strong peak at $t = 0$, related to the ionization and equilibration of the injected material.

After having discussed the representative behaviour at the optima we return to the analysis of the outlier case, the 4-D optimum at $\delta B/B = 0.3\%$. In this case the injected neon density is roughly three orders of magnitude lower than in the other three cases, and as such, it exhibits reheating following the TQ in the plasma centre. This reheated region supports a relatively slowly decaying ohmic current, hence the CQ time is on the long side $t_{\text{CQ}} = 123$ ms (while still tolerable). The slowly decaying ohmic current and the high value of the effective critical electric field E_c^{eff} , owing to the high $n_{\text{D, inj}}$, lead to the RE growth stopping just before the RE current grows to macroscopic values. The strong dilution is able to rapidly reduce the temperature to sufficiently low values at the edge, so that even in the presence of a low neon content the cooling can continue to ≈ 1 eV. As then most of the heat transported to the edge is radiated away by the recombined deuterium, the resulting transported heat loss also remains small in this case. This is a fragile case nevertheless; indeed there is no sample within 25 % of the \mathcal{L} value reached by this optimum. Some parameter combinations in the vicinity of this optimum yield a behaviour reminiscent of C3, with an extremely rapid RE conversion and then a strongly decaying RE current. Thus, even though this optimum performs better than the other three cases in four dimensions, it should not be targeted in a experiment, due to the lack of robustness.

Finally, we comment on the numerical efficiency of the Bayesian approach. We estimate that to achieve a similar level of resolution in the regions that contain samples within 25 % of the optima would require more than 12 000 points in two dimensions and 800 000 points in four dimensions, should we decide to use equidistant scans over the entire search domains. These estimates are based on the average minimum distance between samples (in the search space mapped to the unit hyper-cube). As a reference, we use only 420 samples in both the 2-D and the 4-D optimizations. In uninteresting regions with high cost function values the resolution is much lower.

The Bayesian results can be confirmed with calculations on a uniform grid. In a detailed study of a similar problem presented by Bergström & Haldestam (2022), it was shown that the mean function obtained by the Gaussian process regression accurately recovered the cost function calculated on a uniform grid in the vicinity of the optimum, and showed a good agreement even in regions with high cost values. In terms of finding the global optimum the Bayesian method outperformed Powell's method (Powell 1964).

¹The magnetohydrodynamic stability of the current density is not monitored in the DREAM simulations; a hollow current profile might well be unstable to macroscopic plasma instabilities; this aspect of the simulated current evolution is outside the scope of this study.

4. Discussion and conclusions

We have used Bayesian optimization to find optimal parameters characterizing massive material injection. This is a multi-objective problem where the cost function we aim to minimize accounts for the maximum RE current, the transported heat loss fraction, the CQ time and the final ohmic current. Bayesian optimization is well suited to this problem, as it is a computationally efficient method for finding global optima, providing also uncertainty quantification. In the disruption context, it has also been used recently for validation of simulations of a CQ in a JET plasma discharge with an argon induced disruption (Järvinen *et al.* 2022).

We find that, even in the optimal case, RE currents of several megaampere are predicted. Magnetic perturbations strongly affect the RE dynamics through inducing transport losses of heat and seed REs. Then the optimization is, to a large degree, searching for a balance between sufficiently low transported heat loss – typically favouring large injected impurity quantities and low magnetic perturbation amplitudes – and tolerable final RE current – favouring the opposite conditions. The importance of such a balance has previously been pointed out by Svenningsson *et al.* (2021). In each optimization we kept the normalized magnetic perturbation level constant, in the range 0.2 %–0.5 %. This range of magnetic perturbation levels is motivated by magnetohydrodynamic simulations. We note that higher values are also reached in some recent studies (Särkimäki *et al.* 2020; Nardon *et al.* 2021), which, based on the trends we observe in figure 5, is not expected to have a significant effect on the final RE current, while it would impact the transported heat loss fraction negatively.

The optimum is generally found at a rather high injected deuterium density $n_{D, inj} \approx 10^{22} \text{ m}^{-3}$, while at a lower neon density $n_{Ne, inj} \approx 3 \times 10^{18} \text{ m}^{-3}$. The sensitivity of the optimum to an inaccuracy of the injected deuterium quantity is much stronger than that of the injected neon. The strong sensitivity to the deuterium quantity is due to the possibility of extremely rapid cooling through dilution and subsequent radiation at sufficiently high deuterium densities, which leads to an effective seed generation. In addition, deuterium recombination steeply increases above a certain deuterium density, allowing the already large seed to avalanche more effectively. We also find that neon deposited at the edge is advantageous, where it can produce sufficient radiative heat losses, without making the avalanche RE generation problem more severe, for which the conditions are typically more favourable in the core. Whether an outward peaking impurity density can be sustained long enough to see these benefits can only be answered using higher fidelity simulations. In this sense, our 4-D optimization results can be considered as optimistic bounds.

We point out the importance of choosing the wall radius carefully, as it determines the magnetic energy reservoir for RE generation; a tightly fitted conducting wall may lead to too optimistic results concerning the maximum RE current (yielding 1 MA instead of 4 MA in our example). As we allow for activated RE seed generation mechanisms we cannot find parameter regions where all objectives fall within their respective tolerable ranges; we see, however, that this may not need to be the case with non-activated seed sources only.

The megaampere-scale RE currents predicted even in the optimal scenarios are concerning, thus these results should prompt further studies accounting for additional effects that can impact RE current generation. The most important effects to consider are: (i) magnetohydrodynamic and kinetic instabilities, (ii) vertical displacement and the associated interaction of the current-carrying plasma column with the wall, (iii) the possibility of magnetic surface re-healing taking place significantly later than the end of the TQ and (iv) the possible disappearance of closed flux surfaces below a finite – still

megaampere-level – plasma current. In addition, the dynamics of the injection – which is not resolved here – has a direct impact on the transported heat fraction, and more generally it may affect the temperature evolution and in turn the RE dynamics (mostly the Dreicer and hot-tail seed generation, and as such, it is expected to be more consequential in non-activated operation). Employing this Bayesian framework for the optimization of the more directly accessible parameters describing the injection (for instance the composition and timing of the injected pellets in shattered pellet injection) is thus a natural next step to pursue.

The results are quite robust with respect to the choice of the cost function. The most important trade-off between the various figures of merit appears between achieving a low runaway current and a low transported heat fraction. For instance, in the $\delta B/B = 0.3\%$ 2-D case, changing the weight of η_{cond} in the cost function by $\pm 10\%$ moves the optimum by $\pm 1.5\%$ in $n_{\text{Ne, inj}}$, and by 0.4% in $n_{\text{D, inj}}$. These figures are calculated relative to the extent of the 10% neighbourhood of the optimum on a logarithmic scale (i.e. the size of the corresponding point cloud in [figure 6a](#)). The lower bound of the 10% neighbourhood of the optimum changes by $\pm 5\%$, while the other bounds change by 1% or less. The functional form and weight of the various components in the cost function are ultimately chosen by the user. Currently this arbitrariness of the weights cannot be fully eliminated, partly because of a detailed knowledge about the (monetary) cost of a given value of a figure of merit is lacking, and such estimated figures may never be available. In addition, the current modelling provides too coarse information on the outcome of a given scenario. Indeed, RE beams with the same RE current may cause serious damage, or no detectable effect at all, depending on how the beam is lost to the wall. Recent results indicate that a combination of a low impurity concentration bulk plasma and large-scale magnetohydrodynamic instabilities may enable termination of megaampere-level RE currents without damage to the wall (Paz-Soldan *et al.* 2021; Reux, Cédric *et al.* 2021).

Acknowledgements

The authors are grateful to N. Botta, N. Smallbone, E. Berger, S. Newton, E. Nardon, J. Artola and M. Lehnen 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. 2018-03911 and Dnr. 2022-02862) and in part by the Swiss National Science Foundation. 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 author(s) 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. Simulation details

The magnetic geometry and the initial plasma temperature and current density profiles are shown in [figures 8\(a\)](#) and [8\(b,c\)](#), respectively. The parallel current density component




FIGURE 8. (a) Magnetic geometry with flux surfaces (grey curves), the outermost modelled flux surface $r = a$ is indicated by the thick blue line, and the effective wall is shown in red. The rest of the panels show initial plasma parameter profiles. (b) Electron temperature. (c) Current density.

j is taken at the outboard mid-plane. The magnetic geometry uses a model equilibrium parametrization similar to the Miller equilibrium (Miller *et al.* 1998), with the profiles of elongation, triangularity, Shafranov shift and toroidal magnetic field variation being identical to those shown in appendix A of Pusztai, Hoppe & Vallhagen (2022). The on-axis value is $B_0 = 5.3$ T. The magnetic equilibrium is not evolved self-consistently in the simulation, instead these shaping parameters, as well as the plasma position, are held fixed throughout the simulation.

The DREAM simulations are performed in fluid mode. The Dreicer RE generation rate is calculated using a neural network (Hesslow *et al.* 2019b), which takes effects of partial screening into account. Compton scattering and tritium decay seed sources are accounted for as in Vallhagen *et al.* (2020). The hot-tail seed is calculated using the model described in appendix C.4 in Hoppe *et al.* (2021). The avalanche growth rate accounts for partial screening (Hesslow *et al.* 2019a). Trapping effects are accounted for in the conductivity through the model by Redl *et al.* (2021), and in the avalanche and hot-tail RE generation rates.

The bulk electron temperature evolution is calculated from the time-dependent energy balance throughout the simulation, according to (43) in Hoppe *et al.* (2021), accounting for ohmic heating, line and recombination radiation and bremsstrahlung, as well as a radial heat transport. Since the RE population is not resolved in momentum space, the kinetic term – in (44) of Hoppe *et al.* (2021) – describing heating by REs is zero. However, the latter process is approximately accounted for by a term $j_{\text{RE}} E_c$, with $E_c = e^3 n_e \ln \Lambda_c / (4\pi\epsilon_0 m_e c^2)$ the critical electric field, ϵ_0 the vacuum permittivity and m_e the electron mass. We evolve the temperatures of the ion charge states separately according to (45) in Hoppe *et al.* (2021) which accounts for collisional heat exchange among various charge states as well as with electrons. We neglect current density profile flattening (Pusztai *et al.* 2022) associated with the flux surface breakup.

Opacity effects have been shown to have significant effect on the post-TQ plasma temperature and indirectly on the avalanche gain (Vallhagen *et al.* 2022). These effects are

taken into account by using ionization, recombination and radiation rates for the hydrogen isotopes that are based on the assumption of the plasma being opaque to Lyman radiation.

The simulations use 20 radial grid cells. During the TQ that takes a few milliseconds the solver uses adaptive time stepping with time steps estimated from the relative change of the free electron density within a time step (referred to as ionization-based adaptive time stepping), with allowed minimum and maximum time steps 10^{-11} s and 2×10^{-6} s. The rest of the 150 ms long simulation uses 2×10^4 – 2×10^5 equidistant time steps as needed for convergence.

Appendix B. Details of Bayesian optimization

After n steps our sample data $\mathcal{D}_n := (X_n, Y_n)$ are a collection of control vectors $X_n = \{\mathbf{x}_i\}$ and the corresponding function outputs $Y_n = \{f(\mathbf{x}_i)\}$ where the function f runs DREAM to obtain the four objectives and combines them using the cost function \mathcal{L} . The basic idea of Bayesian optimization is that $f(\mathbf{x})$ is a random variable for each \mathbf{x} and that, given the observations \mathcal{D}_n , the joint distribution of all these random variables is a Gaussian process. The corresponding mean and covariance functions are defined as the expected values

$$\mu(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})], \quad (\text{B1})$$

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - \mu(\mathbf{x}))(f(\mathbf{x}') - \mu(\mathbf{x}'))]. \quad (\text{B2})$$

In our case the DREAM simulation runs are deterministic, which means that the function μ will exactly coincide with f on the samples observed so far. In other points the Gaussian process model provides a smooth interpolation of the cost (something we used to visualize the cost function in [figure 1](#)).

The covariance between two points is modelled by the Matérn kernel (Matérn [1986](#); Stein [1999](#))

$$k_M(\mathbf{x}, \mathbf{x}') = \frac{1}{2^{\zeta-1} \Gamma(\zeta)} \left(2\sqrt{\zeta} |\mathbf{x} - \mathbf{x}'|\right)^{\zeta} K_{\zeta} \left(2\sqrt{\zeta} |\mathbf{x} - \mathbf{x}'|\right), \quad (\text{B3})$$

where Γ denotes the gamma function and K_{ζ} is the modified Bessel function of the second kind. We use a fixed smoothness parameter of $\zeta = 5/2$. The distance between two points in the D -dimensional parameter space is calculated as $|\mathbf{x} - \mathbf{x}'| = \sum_{i=1}^D (x_i - x'_i)^2 / \theta_i^2$, with the correlation length parameters θ_i (which are updated after each new sampling to maximize the marginal likelihood of \mathcal{D}_n).

We use the expected improvement $\text{EI}_n(\mathbf{x})$ acquisition function to find the most promising next point to sample. The following thought experiment (Frazier [2018](#)) illustrates this acquisition strategy. Let f_n^* be the minimal value of f based on the current sample, and let \mathbf{x}_n^* be the corresponding input. If the optimization procedure is terminated at this sample size, \mathbf{x}_n^* would be returned as the best estimate of the actual optimum location \mathbf{x}^* . Suppose that an additional evaluation is to be performed at any point \mathbf{x} yielding $f(\mathbf{x})$. After this, the minimal observed value of f is either $f(\mathbf{x})$ if $f(\mathbf{x}) < f_n^*$ or remain to be f_n^* otherwise. We might define the improvement we gain by performing this additional evaluation to be $f_n^* - f(\mathbf{x})$ in the former case – the amount we could decrease the best value found so far – and 0 in the latter. We aim to maximize this improvement, while $f(\mathbf{x})$ is, as of yet, still unknown. Instead, the next sample location is chosen to maximize the expectation value of the improvement, given the information at hand, that is

$$\mathbf{x}_{n+1} = \operatorname{argmax} \left\{ \mathbb{E}_n [\max(0, f_n^* - f(\mathbf{x}))] \right\} = \operatorname{argmax} \{ \text{EI}_n(\mathbf{x}) \}, \quad (\text{B4})$$

where $\mathbb{E}_n[\cdot]$ should be understood as the expectation under the posterior distribution, given the previously evaluated \mathcal{D}_n .

REFERENCES

- BERGSTRÖM, H. & HALLDESTAM, P. 2022 Optimization of tokamak disruption scenarios: Avoidance of runaway electrons and excessive wall loads. Master's thesis, Chalmers University of Technology.
- BOOZER, A.H. 2012 Theory of tokamak disruptions. *Phys. Plasmas* **19** (5), 058101.
- BREIZMAN, B.N. 2014 Marginal stability model for the decay of runaway electron current. *Nucl. Fusion* **54** (7), 072002.
- BREIZMAN, B.N., ALEYNIKOV, P., HOLLMANN, E.M. & LEHNEN, M. 2019 Physics of runaway electrons in tokamaks. *Nucl. Fusion* **59** (8), 083001.
- BROCHU, E., CORA, V.M. & DE FREITAS, N. 2010 A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. [arXiv:1012.2599](https://arxiv.org/abs/1012.2599).
- FRAZIER, P.I. 2018 A tutorial on Bayesian optimization. [arXiv:1807.02811](https://arxiv.org/abs/1807.02811).
- FÜLÖP, T., HELANDER, P., VALLHAGEN, O., EMBRÉUS, O., HESSLOW, L., SVENSSON, P., CREELY, A.J., HOWARD, N.T. & RODRIGUEZ-FERNANDEZ, P. 2020 Effect of plasma elongation on current dynamics during tokamak disruptions. *J. Plasma Phys.* **86** (1), 474860101.
- 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 growth rate in the presence of high-Z impurities using a neural network. *J. Plasma Phys.* **85**, 475850601.
- HOLLMANN, E.M., ALEYNIKOV, 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., EMBREUS, O. & FÜLÖP, T. 2021 DREAM: a fluid-kinetic framework for tokamak disruption runaway electron simulations. *Comput. Phys. Commun.* **268**, 108098.
- HU, D., NARDON, E., HOELZL, M., WIESCHOLLEK, F., LEHNEN, M., HUIJSMANS, G.T.A., VAN VUGT, D.C., KIM, S.-H. & JET CONTRIBUTORS AND JOREK TEAM 2021 Radiation asymmetry and MHD destabilization during the thermal quench after impurity shattered pellet injection. *Nucl. Fusion* **61** (2), 026015.
- JÄRVINEN, A.E., FÜLÖP, T., HIRVIJOKI, E., HOPPE, M., KIT, A. & ÅSTRÖM, J. 2022 Bayesian approach for validation of runaway electron simulations. *J. Plasma Phys.* **88** (6), 905880612.
- LEHNEN, M., ALEYNIKOVA, K., ALEYNIKOV, P.B., CAMPBELL, D.J., DREWELow, P., EIDIETIS, N.W., GASPARYAN, Y., GRANETZ, R.S., GRIBOV, Y., HARTMANN, N., *et al.* 2015 Disruptions in ITER and strategies for their control and mitigation. *J. Nucl. Mater.* **463**, 39–48.
- LEHNEN, M. & THE ITER DMS TASK FORCE 2021 The ITER disruption mitigation system – design progress and design validation. In *Presented at Theory and Simulation of Disruptions Workshop, PPPL*.
- MATÉRN, B. 1986 *Spatial Variation*. Springer.
- 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.
- NARDON, E., HU, D., ARTOLA, F.J., BONFIGLIO, D., HOELZL, M., BOBOC, A., CARVALHO, P., GERASIMOV, S., HUIJSMANS, G., MITTERAUER, V., *et al.* 2021 Thermal quench and current profile relaxation dynamics in massive-material-injection-triggered tokamak disruptions. *Plasma Phys. Control. Fusion* **63** (11), 115006.
- NOGUEIRA, F. 2014 Bayesian Optimization: open source constrained global optimization tool for Python. <https://github.com/fmfn/BayesianOptimization>.
- PAZ-SOLDAN, C., REUX, C., ALEYNIKOVA, K., ALEYNIKOV, P., BANDARU, V., BEIDLER, M., EIDIETIS, N., LIU, Y.Q., LIU, C., LVOVSKIY, A., *et al.* 2021 A novel path to runaway electron

- mitigation via deuterium injection and current-driven MHD instability. *Nucl. Fusion* **61** (11), 116058.
- POWELL, M.J.D. 1964 An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Comput. J.* **7** (2), 155–162.
- PUSZTAI, I., HOPPE, M. & VALLHAGEN, O. 2022 Runaway dynamics in tokamak disruptions with current relaxation. *J. Plasma Phys.* **88** (4), 905880409.
- 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.
- RASMUSSEN, C.E. & WILLIAMS, C.K.I. 2005 *Gaussian Processes for Machine Learning*. The MIT Press.
- 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., ANGIONI, C., BELL, 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.
- REUX, CÉDRIC, P.-S., CARLOS, A., PAVEL, B., VINODH, F., ONDREJ, S., SCOTT, H., MATTHIAS, J., STEFAN, E., NICHOLAS, L., MICHAEL, S., *et al.* 2021 Demonstration of safe termination of megaampere relativistic electron beams in tokamaks. *Phys. Rev. Lett.* **126**, 175001.
- SÄRKIMÄKI, K., EMBREUS, O., NARDON, E., FÜLÖP, T. & JET CONTRIBUTORS 2020 Assessing energy dependence of the transport of relativistic electrons in perturbed magnetic fields with orbit-following simulations. *Nucl. Fusion* **60** (12), 126050.
- STEIN, M.L. 1999 *Interpolation of Spatial Data*. Springer.
- SVENNINGSSON, I., EMBREUS, O., HOPPE, M., NEWTON, S.L. & FÜLÖP, T. 2021 Hot-tail runaway seed landscape during the thermal quench in tokamaks. *Phys. Rev. Lett.* **127**, 035001.
- SVENSSON, P., EMBREUS, O., NEWTON, S.L., SÄRKIMÄKI, K., VALLHAGEN, O. & FÜLÖP, T. 2021 Effects of magnetic perturbations and radiation on the runaway avalanche. *J. Plasma Phys.* **87**, 905870207.
- VALLHAGEN, O., EMBREUS, 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**, 475860401.
- 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.