

CONFERENCE PRE-PRINT

**NON-LINEAR 3D HYBRID KINETIC-MHD STUDIES OF
RUNAWAY ELECTRON BEAM TERMINATION EVENTS**

H. Bergström
Max Planck Institute for Plasma Physics
Garching b. M., Germany
Email: hannes.bergstroem@ipp.mpg.de

N. Schoonheere
CEA, IRFM
Saint-Paul-lez-Durance, France

P. Haldestam
Max Planck Institute for Plasma Physics
Garching b. M., Germany

V. Bandaru
Indian Institute of Technology Guwahati
Assam, India

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

S-J. Liu
Max Planck Institute for Plasma Physics
Garching b. M., Germany

F. Wouters
Max Planck Institute for Plasma Physics
Garching b. M., Germany

Abstract

Predictive modeling is crucial for assessing the threat posed by runaway electrons (REs) during disruptions in future reactor scale machines. This manuscript details ongoing work and challenges using the 3D non-linear MHD code JOREK to study the benign RE beam termination in JET discharge #95135. JOREK includes a novel full- f particle-in-cell model for REs, which is able to capture accurate transport in stochastic fields, while accounting for the mutual interaction between REs and MHD. The model can be used to study the sensitivity of the termination dynamics with respect to the particle momentum.

While earlier studies adopted a simple mono-energetic assumption for the REs, in this work we use DREAM simulations to derive estimates for the RE momentum distribution at different points during the disruption. The argon and deuterium densities are treated as unknown quantities and Bayesian optimization is applied to find the best agreement with regards to the experimentally measured current evolution. For the best matching case, the mean energy of the RE population is found to decrease from 18 MeV to 10 MeV in the phase leading up to the second injection.

Steps towards applying the resulting RE distribution in JOREK are then shown. The large drift orbits of the energetic particles are found to substantially alter the current density and the force balance inside the plasma. As a result, the profiles differ notably compared to those in earlier work. There are ongoing efforts to initialize the particles in a way such that the resulting profiles are consistent for different distributions, allowing for a more direct comparison. Finally some preliminary results from 3D JOREK simulations are shown where a $(m, n) = (4, 1)$ double tearing mode instability leads to large scale field line stochastization, but only a partial flush out of the RE population is observed.

1. INTRODUCTION

Disruptions and the associated generation of runaway electrons (REs) remain one of the biggest challenges in the design of future large scale tokamak reactors. The REs undergo an almost unhindered acceleration by the strong electric field induced during a disruption, quickly reaching relativistic energies [1]. An initial seed of these particles can then generate new REs through collisions with the thermal bulk, resulting in an exponential growth of the RE population [2]. Unless mitigated, disruptions in high current devices such as ITER are expected to result in the formation of a multi-MA beam of REs [3, 4]. As the beam eventually becomes deconfined, the highly energetic particles could strike the plasma facing components (PFCs) in a localized manner, causing intolerable heat loads and posing a threat to the integrity of the cooling channels.

To properly assess the severity of disruptions in future machines and find pathways for mitigation, one cannot only extrapolate from present day experiments, which lie in a very different parameter range compared to the foreseen reactor scale devices. For this purpose, predictive modeling, including high-fidelity simulations, is an important tool which can help inform future machine designs.

Currently, one of the most promising strategies to mitigate the damage caused by REs is through so-called *benign termination*, where an injection of low-Z material into an already formed RE beam largely recombines the plasma, facilitating the growth of a large scale MHD instability [5]. This eventually leads to a global stochastization of the magnetic field lines which flushes out the REs. The idea of this approach is that by inducing a violent MHD event, the particles will be spread over a larger area, reducing the localization of the heat loads. In addition, a near complete loss of the RE population is required to avoid reformation of the RE beam. Benign termination has been observed on a number of different machines, including JET, DIII-D, TCV and AUG [5, 6, 7].

While the scenario can be consistently reproduced in experiments, there is an ongoing effort to better understand the sequence of events from a modeling point of view. This includes the dynamics of the bulk plasma and the recombination, but also how the plasma can reach such an unstable state without terminating earlier.

This work aims to use a recently developed hybrid kinetic-MHD model for REs in the JOREK code to study different dependencies of the beam termination, based on earlier simulations of JET discharge #95135 [8, 9]. The model allows for simulating the mutual coupling between REs and MHD, while resolving the full phase space distribution of the particles and capturing accurate transport in the 3D fields. While earlier applications of this model used simple mono-energetic assumptions for the RE momenta, in this work we move towards more realistic distributions based on DREAM simulations of the phases with low MHD activity.

The manuscript is structured as follows: section 2 describes the considered JET experiment and the sequence of events leading up to the termination. section 3 then explains how DREAM is used to obtain a rough estimate for the distribution function at different stages, using Bayesian optimization to estimate the impurity content in the plasma. section 4 details current progress on implementing more sophisticated momentum distributions in JOREK simulations, highlighting some of the challenges and showing initial 3D results. Finally, section 5 provides some concluding remarks and outlook.

2. EXPERIMENTAL SCENARIO

The benign termination strategy and JET discharge #95135 are described in Ref. [5]. In this section we provide a general overview of the experiment but for the full details the reader is referred to the original article.

The discharge is carried out using a limiter configuration for the plasma to ensure vertical stability. A disruption is triggered by the injection of $8.21 \cdot 10^{20}$ argon atoms, converting the initial Ohmic current of 1.3 MA to a RE current of roughly 650 kA. Fig. 1 shows the evolution of the total plasma current during the disruption. About 400 ms into the runaway plateau, a second injection of $1.43 \cdot 10^{23}$ deuterium atoms is made using SPI, after which the free electron density quickly drops to nonmeasurable values, i.e. $n_e < 10^{18} \text{ m}^{-3}$. In addition, the decrease in argon line radiation suggests that the argon has been, at least partially, flushed out of the plasma. With a recombined plasma, the loop voltage can be used to further increase the RE current (and decrease the edge safety factor q_{edge}) which ramps up to ~ 760 kA before the beam terminates, which can be seen as a spike in the hard x-ray and neutron diagnostics, stemming from REs becoming deconfined and colliding with the wall. During and after the termination no localized heat loads can be seen in the thermography measurements.

Experimental observations over several machines point to two actuators of the benign termination: the recombination of the plasma and expulsion of high-Z impurities as well as the reduction of q_{edge} [5, 7]. The decrease in free electron density reduces the Alfvén time which is believed to contribute to the large MHD crash observed later. A reduction of high-Z material also means that RE generation through avalanching might not be as efficient [10], which helps prevent the reformation of a RE beam. It has been found that the reduction of q_{edge} which eventually destabilizes the plasma can be achieved either by increasing the plasma current, or by compressing the plasma, with both approaches still resulting in a benign termination [7].

The simulations in this work will consider different phases of the runaway plateau, which have been indicated in Fig. 1. To study the evolution of the momentum space distribution, two different DREAM simulations are carried out corresponding to the more static periods, where we do not expect to see pronounced MHD activity or large variations in the particle balance. The first one represents the time shortly after the completion of the current quench, while the latter starts a few ms after the




FIG. 1. Evolution of the total plasma current throughout the disruption. The two phases which are to be simulated using DREAM are indicated by the blue and green regions. The time points of the second material injection, as well as the start of the JOREK simulation are indicated by the dashed and dotted line respectively. Following Ref. [8], the chosen reference time corresponds to $t = 48.0233$ s in the experiment

second injection and ends some time shortly before the termination. Finally, the JOREK simulations use as a starting point the equilibrium right before the final crash.

3. DREAM SIMULATIONS

DREAM is a 1D flux surface averaged code with both fluid and kinetic models for bulk and runaway electrons [11]. In this work the REs are resolved kinetically while the bulk plasma is modeled as a fluid. Furthermore, since the goal is the study of RE momentum space dynamics, the simulation is run in 0D, meaning only one radial point is considered. The RE distribution is evolved according to the 3D Fokker-Planck equation, accounting for synchrotron losses, bremsstrahlung and small angle collisions with electrons, ions and neutrals, including the effects of partial screening. It also includes models for the different RE sources, although, in this work, only the avalanche source is considered as it is assumed to be dominant. The electric field is evolved self-consistently in the simulations. As a boundary condition the time evolution of the loop voltage on the wall has been prescribed based on experimental data.

For the first phase, i.e. before the second injection, the temperature and particle charge states are evolved self-consistently, assuming an initial temperature of $T_e = T_i = 10$ eV. The initial RE distribution function is based on the analytical formula given in Ref. [12, 13]:

$$f_{\text{RE}}(p, \xi) = \frac{n_{\text{RE}} A(p)}{2\pi m_e c \gamma_0 p^2} \frac{\exp\left[-\frac{\gamma}{\gamma_0} - A(p)(1 + \xi)\right]}{1 - e^{-2A}}, \quad (1)$$

where

$$A(p) = \frac{E_{\parallel}/E_c + 1}{Z_{\text{tot}} + 1} \gamma, \quad \gamma_0 = c_Z \ln \Lambda, \quad \gamma = \sqrt{1 + (p/m_e c)^2}, \quad c_Z \approx \sqrt{5 + Z_{\text{eff}}}.$$

Here p is the RE momenta, $\xi = p_{\parallel}/p$ is the particle pitch, m_e denotes the rest mass of the electron and c the speed of light. Z_{tot} and Z_{eff} is the total and effective charge respectively, n_{RE} is the runaway density and E_c the critical electric field. The distribution is derived under the assumption that RE generation is dominated by avalanching under the influence of a constant parallel electric field $E_{\parallel}$, which mainly holds true during the current quench phase. During the runaway plateau, at near-threshold values of the electric field, the momentum distribution is instead expected to peak around some attractor region, with a lower cut-off determined by the collisional drag on the particles and an upper limit set by synchrotron losses [14]. In this work we assume that $E_{\parallel}/E_c = 1$ and set n_{RE} such that the runaway current obtained from taking the moment of f_{RE} agrees with the total current at this point in time in the experiment. The choice of $E_{\parallel}/E_c$ is somewhat arbitrary but it was found to not significantly alter the simulation results.

The total deuterium and argon densities, n_D and n_{Ar} are kept constant during the simulation, while their charge states are evolved consistently with the temperature. The amount of material after the current quench is difficult to infer from measurements but plays an important role in the RE dynamics, determining the drag on the particles (mainly resulting from

collisions with free electrons), as well as the rate of pitch angle scattering (dominated by collisions with ions), which enhance the synchrotron losses of the REs. They are also essential in determining the temperature, electric field strength, and overall current evolution. Because of this, the material densities are free parameters in the simulations and scans are carried out to study the variation in the resulting RE distributions. As a measure of the accuracy of the results, the current evolution is compared to the experiment.

An efficient way to scan the parameter space, while avoiding regions with large discrepancies compared to experimental data, is to use Bayesian optimization [15]. Similar to earlier work done with DREAM [4, 16, 17], we adopt Gaussian process regression, in which a set of sampled points is used to construct multivariate gaussian distributions, describing the expected value $\mu(n_D, n_{Ar})$ of some objective function $\mathcal{L}(n_D, n_{Ar})$, which we aim to maximize. One of the strengths of this approach is that in addition to sampling regions where the objective function is expected to be the highest, the algorithm can also prioritize regions with greater uncertainty, as quantified by the covariance of the Gaussian distribution. This helps in searching the whole domain, increasing the chance of finding the global maxima. For our simulations we define the objective function as the root mean square error (RMSE) between the total simulated current and the total current measured in the experiment:

$$\mathcal{L} = -\sqrt{\frac{\sum_{i=0}^{N_t} (I_i^{\text{sim}} - I_i^{\text{exp}})^2}{N_t}}, \quad (2)$$

where I_i^{sim} denotes the simulated current at time step i and I_i^{exp} the experimental current interpolated at the same time. N_t is the total number of time steps in the simulations. The minus sign is simply added to reformulate the problem such that the best fit will maximize the objective function.

Initially 32 points are sampled uniformly in the parameter space. The limits are set such that $n_D \in [0.5 \cdot 10^{19}, 4 \cdot 10^{19}] \text{ m}^{-3}$ and $n_{Ar} \in [0.1 \cdot 10^{19}, 5.3 \cdot 10^{19}] \text{ m}^{-3}$, where the upper argon limit corresponds to 100% assimilation of the injected material. The deuterium content in the vessel pre-disruption corresponds to about half the amount of injected argon atoms, but due to the convergence properties seen shortly, the limits were set a bit broader to get a better understanding of the impact of the deuterium amount on the optimization results. After the initial samples, new samples are generated according to the expected improvement acquisition function [4], which weights "exploitation", i.e. sampling the region where it predicts the highest value for the objective function, and "exploration", meaning sampling the region with the highest uncertainty. Consequently, this approach enables an efficient balance between global search and convergence speed.

The results from an optimization using a total of 96 samples are shown in Fig. 2. Fig. 2a shows the expectation value of the RMSE determined from the Gaussian process. Using the same coloring based on the RMSE, Fig. 2b, c and d show the current evolution, resulting pitch-averaged RE distribution and mean energy of the REs respectively. From Fig. 2a and b it is clear that the current evolution doesn't change notably when varying n_D but is highly dependent on the amount of argon in the plasma. The optimal value is found around $n_D = 2.85 \cdot 10^{19} \text{ m}^{-3}$ and $n_{Ar} = 2.82 \cdot 10^{19} \text{ m}^{-3}$, which would correspond to a roughly 53% assimilation of the injected argon. It should be emphasized that as the objective function is essentially constant with respect to the amount of background deuterium, it is not possible to infer this quantity from the analysis. In principle additional components could be added to the objective function, for instance considering the free electron density measured in the experiment. It is however important to note that from Fig. 2c and d we find that the final momentum distribution is also mainly determined by the amount of argon in the plasma. The initial mean energy is in most cases around 18 MeV and the mean energy at the end of the simulation converges near 10.8 MeV when plotted as a function of the RMSE. It is apparent that the highly energetic tail of the distribution has substantially decreased, which is mainly attributed to synchrotron losses, while a small bump has formed in the region $p/(m_e c) \in [9, 32]$, at low RMSE.

Since the objective of this exercise is to derive more representative distribution functions to be used in termination studies, and the distributions calculated by DREAM show clear convergence for decreasing RMSE, the results of the optimization were deemed satisfactory. Work on applying a similar approach to the second phase, after the recombination of the plasma, is currently underway.

4. TERMINATION SIMULATIONS

The previous section detailed how DREAM was used to give an estimate for the RE distribution, prior to the recombination of the plasma, which does not assume continuous avalanching under a constant electric field. For the JOREK simulations, the aim is to use characteristic distributions at the end of the current quench (approximated by Eq. 1), before the recombination of the plasma and at the onset of the final termination. Because of the computational and time requirements of a hybrid RE simulation in JOREK, at present only the distribution derived in section 3 is considered. But the avalanche distribution as well as distributions in the presence of a recombined plasma are planned to be included in future simulations.

The simulation set up is described in subsection 4.1. It closely follows earlier JOREK work with a fluid RE model [8], but additional complexities are introduced when the REs are resolved kinetically. Initial tests of the 3D termination dynamics are then given in subsection 4.2. Results are still preliminary as profiles will continue to be adjusted to simplify comparisons to earlier studies.




FIG. 2. Results from Bayesian optimization of DREAM simulations. (a): RMSE with respect to the total current, as predicted by the Gaussian process. Sample points have been indicated by the black dots and the red star denotes the optimum. The same color scale is used in figures (b)-(d). (b): Total current throughout the simulations along with the experimentally measured current. (c): Pitch angle averaged RE distribution at the end of the simulation. The initial distribution at the start of the simulation corresponding to the optimum has been included for comparison. (d): Average RE energy at the beginning and end of the simulations as a function of the RMSE.

4.1. Particle initialization and equilibration

The simulations use as a starting point the equilibrium at the onset of the final crash. A constant density and (low) temperature is prescribed, assuming that these quantities will not change considerably on the time scale of the termination. The resistivity is set to be uniform in space and corresponds to the Spitzer resistivity at $T_e = 10$ eV, i.e. $\eta = 3.3 \cdot 10^{-5} \Omega\text{m}$. As described in Ref. [8], the RE density in the original simulations was chosen to yield qualitative agreement with synchrotron data, and such that the q-profile was susceptible to the $(m, n) = (4, 1)$ double tearing mode (TM) observed in the experiment. This corresponded to a hollow profile for the RE current. As will be discussed shortly, using the same particle distribution with the kinetic model turns out to be non-trivial.

Since DREAM is a flux surface averaged code, the momentum space distribution obtained from the simulations describes the momentum and pitch of a particle when it is at the point of minimum magnetic field along its orbit, i.e. at the midplane on the low field side. To obtain a distribution in JOEREK, where full orbit particle tracing is used to model the REs kinetically, that is consistent with the DREAM results, the particles are therefore initialized along the midplane low field side with momenta sampled from the DREAM distribution, and then randomly pushed over a period of $t_{\text{init}} = u \cdot N_t \cdot \Delta t_{\text{part}}$. Here $\Delta t_{\text{part}} = 5 \cdot 10^{-12}$ s is the particle time step, $u \sim \mathcal{U}(0, 1)$ is a uniformly sampled random number and $\Delta t_{\text{part}} \cdot N_t = 10 \mu\text{s}$ is the maximal initialization time. This step helps evenly distribute the particles along their orbits in 3D space.

When determining where to place the marker along the midplane, an inconsistency between the models becomes apparent. Both the kinetic model in DREAM and the fluid RE model used in preceding JET simulations with JOEREK use the zero orbit width approximation, meaning that particle orbits follow the magnetic flux surfaces exactly. At energies of 10 MeV and above, the curvature drift of the REs will however become quite notable. The true constant of motion describing these orbits is the canonical angular momentum:

$$P_\varphi = -e\psi + Rb_\varphi p_\parallel, \quad (3)$$

where e is the elementary charge, ψ the poloidal magnetic flux, R the major radial coordinate and $b_\varphi = B_\varphi/B$ with B_φ denoting the toroidal magnetic field. From Eq. 3 it can be seen that the particles will follow the poloidal flux surfaces exactly in the limit where $|Rb_\varphi p_\parallel| \ll |e\psi|$. At higher energies however, there will be a relative shift between the flux surfaces and drift orbit surfaces which skews the spatial profiles. In addition, the equilibrium also changes as a result of the drift orbits, leading to a Shafranov like shift of the flux surfaces [18]. As such, it is important to note that since the drift orbits fundamentally change the plasma profiles and equilibrium, comparisons between different distributions, as well as between kinetic and fluid models for REs, are not straightforward.

Nonetheless, to achieve profiles for the current density and safety factor that resemble those in earlier work, the particle position is sampled proportionally to

$$p(R|p, \xi) = \frac{\bar{J}(R)r'(p, \xi)R}{\int dR \bar{J}(R)r'(p, \xi)R}. \quad (4)$$

Here $\bar{J}$ denotes the target current density along the midplane on the low field side, while the minor radius r' and major radius R are included to account for the difference in surface area as the particles are distributed along their orbits. Note that r' corresponds to the minor radius of the drift surface, calculated using Eq. 3, and not the minor radius of the flux surface. At




FIG. 3. Midplane current profiles (a) and safety factor (b) from the initial RE fluid simulations of Ref. [8] and for the new equilibrium reached using the RE distribution from section 3. In addition, the modified q-profiles used in the 3D simulations discussed in subsection 4.2 are shown.

high energies, using the minor radius of the flux surface can lead to large peaking of the particle density near the axis. To allow for some tweaking of the profile Eq. 4 was adjusted such that $r' \rightarrow r' + h(r')^2$, where h is a hyper parameter which increases the hollowness of the profile. Due to the drift orbits, this was found to be needed to reach a sufficiently high q_0 for the (4, 1) double TM instability to occur.

Applying this methodology to the distribution from section 3, 10^7 marker particles are initialized using $h = 0.8$. As described in Ref. [9], the plasma is then allowed to relax into a new equilibrium accounting for the changes in the force balance. This implies running the coupled simulation in an axisymmetric configuration for roughly $300 \mu\text{s}$, during which there is a notable shift of the flux surfaces. The resulting current profile together with the one from the preceding fluid RE simulation are shown in Fig. 3, along with the corresponding q-profiles. From the figure it is apparent that even when adjusting the hollowness of the profile, the two current densities are qualitatively very different, meaning that a direct comparison between the hybrid simulation, and the fluid RE simulation is difficult. In the future, we aim to minimize the discrepancy between the profiles by introducing additional shaping parameters to Eq. 4.

4.2. Initial termination results

To assess whether the same (4, 1) TM termination dynamics could be observed with the equilibrium obtained using the updated RE distribution, initial 3D tests were carried out including mode numbers $n = 1, 2, 3$. As a simple scan, the q-profiles were varied by artificially increasing B_φ , which shifts the q-profile upward and importantly decreases the distance between the two $q = 4$ surfaces, which we denote $\Delta r_{q=4}$. It was found that for the (4, 1) mode to be dominant, the toroidal field needed to be increased by roughly $\sim 23\%$, for which $\Delta r_{q=4} = 22$ cm. In this case only a partial loss of the REs was observed however. Further increasing B_φ by 26% of the original amplitude such that $\Delta r_{q=4} = 19$ cm, the stochastization becomes more pronounced but still a large fraction of REs are not lost, with roughly 400 kA of RE current remaining. The modified q-profiles are also shown in Fig. 3b.

Fig. 4 shows the evolution of the currents (a) and the magnetic energy (b and c) for the different modes throughout the 3D simulation. While particles are still being transported out at the end of the $\Delta r_{q=4} = 22$ cm simulation, the loss rate is much slower than that seen in preceding work [8, 9]. It should be mentioned that a similar non-complete termination was observed in earlier simulations using higher energy REs (8 MeV) [9], but at present it is too early to suggest any trend with respect to particle momenta, as the profiles differ notably between the cases.

5. CONCLUSION AND OUTLOOK

In this manuscript, ongoing studies of RE kinetics and their impact on the MHD during termination events was presented. Utilizing DREAM simulations in combination with Bayesian optimization, an estimate for the RE momentum distribution prior to the second injection could be obtained. The best agreement with respect to the experimentally measured current was achieved assuming a 53% assimilation of the injected argon, in which case the mean particle energy decreased from ~ 18 MeV down to ~ 10.8 MeV during the simulation. There is ongoing work to run similar DREAM simulations for the post second injection phase, leading up to the beam termination.




FIG. 4. Current evolution (a) and the evolution of the magnetic energy of the different modes for the case where $\Delta r_{q=4} = 22$ cm (b) and $\Delta r_{q=4} = 19$ cm (c).

The calculated distribution was then used in JOREK, where particles were initialized along the midplane low field side to be consistent with the DREAM results. A point of inconsistency is however the particle orbits, where accounting for the curvature drift qualitatively changes the current profile and force balance if the REs are sufficiently energetic. As a result, the simulations are difficult to compare to earlier work owing to discrepancies in current densities and safety factor. Some initial 3D tests showed that the expected (4, 1) double TM instability could be recovered in the simulations, although this required shifting the q-profiles. In both simulations only a partial loss of REs was observed. This is qualitatively similar what has been observed at higher energies in previous studies, but more work is needed to make the profiles comparable.

The work presented here contributes to a complete understanding of the tools and workflows needed for simulating benign termination fully consistently with a hybrid fluid-kinetic approach. We do, however, not address in this work the question of how the plasma evolves into the highly unstable state causing the benign termination. Very likely, kinetic effects with stabilizing influence on the MHD modes will play an important role in this respect such that this open question can be addressed in future work based on the models and methods we develop here. Furthermore, in contrast to the JET scenario we consider here, the most relevant benign termination scenario for ITER is based on external modes that get suddenly excited while the edge safety factor drops to values around two. All these aspects require simulating not just the termination itself, but also the longer time scales before. These questions are of highest relevance and interest for our future work and will become accessible with recent developments and code porting to accelerated high performance computing architectures, which allows to efficiently cross these longer time scales.

ACKNOWLEDGEMENTS

This 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.

REFERENCES

- [1] Boris N Breizman et al. “Physics of runaway electrons in tokamaks”. In: *Nuclear Fusion* 59.8 (2019), p. 083001. DOI: 10.1088/1741-4326/ab1822.
- [2] MN Rosenbluth and SV Putvinski. “Theory for avalanche of runaway electrons in tokamaks”. In: *Nuclear fusion* 37.10 (1997), p. 1355. DOI: 10.1088/0029-5515/37/10/I03.

- [3] O. Vallhagen et al. “Runaway electron dynamics in ITER disruptions with shattered pellet injections”. In: *Nuclear Fusion* 64.8 (June 2024), p. 086033. DOI: 10.1088/1741-4326/ad54d7.
- [4] I. Pusztai et al. “Bayesian optimization of massive material injection for disruption mitigation in tokamaks”. In: *Journal of Plasma Physics* 89.2 (2023), p. 905890204. DOI: 10.1017/S0022377823000193.
- [5] Cédric Reux et al. “Demonstration of Safe Termination of Megaampere Relativistic Electron Beams in Tokamaks”. In: *Phys. Rev. Lett.* 126 (17 Apr. 2021), p. 175001. DOI: 10.1103/PhysRevLett.126.175001.
- [6] EM Hollmann et al. “Trends in runaway electron plateau partial recombination by massive H2 or D2 injection in DIII-D and JET and first extrapolations to ITER and SPARC”. In: *Nuclear Fusion* 63.3 (2023), p. 036011. DOI: 10.1088/1741-4326/acb4aa.
- [7] Umar Sheikh et al. “Benign termination of runaway electron beams on ASDEX Upgrade and TCV”. In: *Plasma Physics and Controlled Fusion* 66.3 (2024), p. 035003. DOI: 10.1088/1361-6587/ad1e31.
- [8] V Bandaru et al. “Magnetohydrodynamic simulations of runaway electron beam termination in JET”. In: *Plasma Physics and Controlled Fusion* 63.3 (Jan. 2021), p. 035024. DOI: 10.1088/1361-6587/abdbcf.
- [9] Hannes Bergström et al. “Introduction of a 3D global non-linear full-f particle-in-cell model for runaway electrons in JOREK”. In: *Plasma Physics and Controlled Fusion* 67.3 (2025), p. 035004. DOI: 10.1088/1361-6587/adaee7.
- [10] Linnea Hesslow et al. “Effect of partially ionized impurities and radiation on the effective critical electric field for runaway generation”. In: *Plasma Physics and Controlled Fusion* 60.7 (2018), p. 074010. DOI: 10.1088/1361-6587/aac33e.
- [11] Mathias Hoppe, Ola Embréus, and Tünde Fülöp. “DREAM: a fluid-kinetic framework for tokamak disruption runaway electron simulations”. In: *Computer Physics Communications* 268 (2021), p. 108098. DOI: 10.1016/j.cpc.2021.108098.
- [12] Tünde Fülöp et al. “Destabilization of magnetosonic-whistler waves by a relativistic runaway beam”. In: *Physics of Plasmas* 13.6 (2006). DOI: 10.1063/1.2208327.
- [13] Ola Embréus et al. “Dynamics of positrons during relativistic electron runaway”. In: *Journal of Plasma Physics* 84.5 (2018), p. 905840506. DOI: 10.1017/S0022377818001010.
- [14] Pavel Aleynikov and Boris N. Breizman. “Theory of Two Threshold Fields for Relativistic Runaway Electrons”. In: *Phys. Rev. Lett.* 114 (15 Apr. 2015), p. 155001. DOI: 10.1103/PhysRevLett.114.155001.
- [15] Eric Brochu, Vlad M Cora, and Nando De Freitas. “A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning”. In: *arXiv preprint arXiv:1012.2599* (2010).
- [16] I. Ekmark et al. “Runaway electron generation in disruptions mitigated by deuterium and noble gas injection in SPARC”. In: *Journal of Plasma Physics* 91 (3 2025), E82. DOI: 10.1017/S0022377825000455.
- [17] Hannes Bergström and Peter Haldestam. “Optimization of tokamak disruption scenarios”. MA thesis. Chalmers University of Technology, 2022. DOI: 20.500.12380/305257.
- [18] V Bandaru and M Hoelzl. “Tokamak plasma equilibrium with relativistic runaway electrons”. In: *Physics of Plasmas* 30.9 (2023). DOI: 10.1063/5.0165240.