

Bayesian approach for validation of runaway electron simulations

A.E. Järvinen ^{1,2,†}, T. Fülöp ³, E. Hirvijoki ⁴, M. Hoppe ⁵, A. Kit ²,
J. Åström ^{2,6} and JET Contributors[‡]

¹VTT Technical Research Centre of Finland, FI-02044 VTT, Finland

²University of Helsinki, FI-00014 Helsinki, Finland

³Chalmers University of Technology, SE-412 96 Göteborg, Sweden

⁴Aalto University, FI-00076 AALTO, Finland

⁵Ecole Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center, CH-1015 Lausanne, Switzerland

⁶CSC-IT Center for Science, FI-02101 Espoo, Finland

(Received 1 August 2022; revised 11 November 2022; accepted 14 November 2022)

Plasma-terminating disruptions in future fusion reactors may result in conversion of the initial current to a relativistic runaway electron beam. Validated predictive tools are required to optimise the scenarios and mitigation actuators to avoid the excessive damage that can be caused by such events. Many of the simulation tools applied in fusion energy research require the user to specify input parameters that are not constrained by the available experimental information. The conventional approach, where an expert modeller calibrates these input parameters based on domain knowledge, is prone to lead to an intractable validation challenge without systematic uncertainty quantification. Bayesian inference algorithms offer a promising alternative approach that naturally includes uncertainty quantification and is less subject to user bias in choosing the input parameters. The main challenge in using these methods is the computational cost of simulating enough samples to construct the posterior distributions for the uncertain input parameters. This challenge can be overcome by combining probabilistic surrogate modelling, such as Gaussian process regression, with Bayesian optimisation, which can reduce the number of required simulations by several orders of magnitude. Here, we implement this type of Bayesian optimisation framework for a model for analysis of disruption runaway electrons, and explore for simulations of current quench in a JET plasma discharge with an argon induced disruption. We use this proof-of-principle framework to explore the optimum input parameters with uncertainties in optimisation tasks ranging from one to seven

[†] Email address for correspondence: aaro.jarvinen@vtt.fi

[‡] See the author list of ‘Overview of JET results for optimising ITER operation’ by Mailloux *et al.* (*Nucl. Fusion*, vol. 62, 2022, 042026).

dimensions. The relevant Python codes that are used in the analysis are available via https://github.com/aejarvin/BO_FOR_RE_SIMULATIONS/.

Key words: fusion plasma, runaway electrons, plasma simulation

1. Introduction

Runaway electrons (REs) pose one of the leading concerns regarding the integrity and duty cycle of future fusion reactors (Boozer 2018). As RE generation in disruptions is exponentially sensitive to the plasma current, I_p (Rosenbluth & Putvinski 1997), and I_p is projected to increase from the level of a few MA in large present-day tokamaks to the range of 10–20 MA in power plant relevant tokamaks, an unmitigated RE beam at reactor scale would be expected to cause severe damage and extended downtime and repair costs (Boozer 2018; Breizman *et al.* 2019). Therefore, there is a strong need for validated tools to predict and avoid disruptions and RE beam generations when entering the era of reactor-scale devices.

To address this need, several numerical tools have been developed for disruption and RE analysis, such as the nonlinear magnetohydrodynamic codes JOEAK (Huysmans & Czarny 2007) and NIMROD (Sovinec *et al.* 2004), the kinetic code CQL3D (Harvey *et al.* 2000) and the fluid-kinetic framework DREAM (Hoppe, Embreus & Fülöp 2021). However, validating these simulation tools with experimental data is complicated, as typically some of the input parameters of the simulations are not well constrained by the available experimental information. In such a situation, the user must first specify values for these uncertain input parameters to calibrate the model to appropriately represent the investigated system. This challenge is common to other numerical tools applied in magnetic confinement fusion research as well, such as scrape-off layer plasma simulations conducted with SOLPS-ITER (Wiesen *et al.* 2015). A conventional approach to calibrating the model is to use previous experience and domain knowledge to conduct the necessary parameter fitting manually, aiming to find a set of input parameters that minimises the discrepancy between the synthetic and measured diagnostic data. The remaining discrepancy is then evaluated and documented as a degree of validity of the model.

However, this type of expert modeller approach becomes intractable as the number of uncertain parameters increases. With multiple uncertain input parameters, manual optimisation is prone to lead to subjective reasoning for the trajectory selection through the space of optimised parameters. As a result, the degree of confidence in the obtained solution and its uncertainty is likely to remain ambiguous. Furthermore, the *ad hoc* nature of subjective reasoning is prone to lack of scientific rigour, which is problematic from the point of view of aiming to establish objective and reproducible scientific results. These challenges can be alleviated by applying a regimented approach, such as grid search, instead of subjective reasoning, but such an approach is intractable to conduct manually with multiple uncertain input parameters and is computationally inefficient when operating with costly simulations. Due to these inefficiencies of the expert modeller approach, an optimisation algorithm that would take the human out-of-the-loop, provide a rigorous systematic approach with natural uncertainty quantification and select the samples from the search space in a computationally efficient manner would be a very attractive alternative approach (Brochu, Cora & de Freitas 2010; Shahriari *et al.* 2016; Frazier 2018).

In this study, approximate Bayesian computation (ABC) and Bayesian optimisation (BO) are applied to find the optimal values and provide inverse uncertainty quantification

of uncertain input parameters (Brochu *et al.* 2010; Marin *et al.* 2012; Shahriari *et al.* 2016; Frazier 2018) in DREAM RE simulations. Inverse uncertainty quantification here refers to quantifying the uncertainty of the uncertain input parameters of the model given observed experimental data (Oberkampf & Trucano 2002; Wu *et al.* 2018). As the algorithm provides the uncertainty distributions for the free input parameters of the model, the human scientist can focus on analysing the resulting optimal parameters and the degree of validity of the model instead of manual tuning of the fitting parameters. The aim of this study is to provide a proof-of-principle approach to using these methods in calibrating the uncertain input parameters in RE simulations, while the methodology could be used broadly in validating other predictive tools within magnetic confinement fusion research. The implementation is based on the BO for likelihood-free inference (BOLFI) method of the engine for likelihood-free inference Python software package (Gutmann & Corander 2016; Lintusaari *et al.* 2018).

2. Bayesian approach

This section describes the methodology used in this study. Section 2.1 provides an overview of ABC, § 2.2 an introduction to BO, § 2.3 describes the usage of Gaussian process regression as a probabilistic surrogate model in BO and § 2.4 describes the functionality of the acquisition functions in BO.

2.1. Approximate Bayesian computation

Bayesian inference aims to establish the conditional probability distribution, $P(\mathbf{x}|D)$, called the posterior, of the uncertain input parameters, \mathbf{x} , given observed experimental measurements, D . Here, $P(\mathbf{x}|D)$ represents the best estimate and uncertainty of the input parameters for the investigated system. Bayesian inference applies the Bayes' theorem

$$P(\mathbf{x}|D) = \frac{P(D|\mathbf{x})P(\mathbf{x})}{P(D)}, \tag{2.1}$$

which states that the posterior is proportional to the likelihood of D given \mathbf{x} , $P(D|\mathbf{x})$, multiplied by the prior probability distribution for \mathbf{x} , $P(\mathbf{x})$. The marginal probability of the experimental measurements, $P(D)$, represents an integral over all possible data generating input values, $\int P(D|\mathbf{x}')P(\mathbf{x}') d\mathbf{x}'$, which would be computationally challenging to evaluate. However, it is typically sufficient to establish the relative posterior probabilities of various values of \mathbf{x} . Therefore, $P(D)$ does not need to be directly evaluated, and it is sufficient to apply the proportionality

$$P(\mathbf{x}|D) \propto P(D|\mathbf{x})P(\mathbf{x}). \tag{2.2}$$

When the likelihood function, $P(D|\mathbf{x})$, is either not available analytically or cannot be evaluated within the available computational or time resources, the standard alternative is to use ABC (Marin *et al.* 2012; Järvenpää *et al.* 2018); ABC aims to establish the approximate Bayesian posterior

$$P_{\text{ABC}}(\mathbf{x}|D) \propto P(\mathbf{x}) \int H(\epsilon - \Delta(D, y))P(y|\mathbf{x}) dy, \tag{2.3}$$

where $y \in \mathbb{R}^d$ are data generated with the simulation model with input parameters \mathbf{x} , and $\Delta : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_{>0}$ is a discrepancy function between the simulated and measured data. Here, ϵ represents the threshold parameter controlling the trade-off between posterior estimation accuracy and efficiency, and $H(x)$ is the Heaviside step function which takes a

value 1 whenever ϵ is greater than the discrepancy. Small values of ϵ lead to more accurate approximate posteriors, but also increase the computational challenge.

One of the simplest ABC algorithms that could be applied to numerically estimate the integral (2.3) is rejection sampling (Marin *et al.* 2012; Lintusaari *et al.* 2017; Järvenpää *et al.* 2018):

- (i) Draw a random sample from $\mathbf{x}' \sim P(\mathbf{x})$.
- (ii) Generate $y' \sim P(y|\mathbf{x}')$.
- (iii) Accept y' if $\Delta(D, y') \leq \epsilon$.
- (iv) Go back to step (i) until sufficient number of accepted samples are collected.

The accepted values represent the approximate posterior distribution. The drawback of the standard rejection sampling is the number of required function evaluations. For a typical simulation tool in magnetic confinement fusion, a function evaluation takes at least several minutes and more typically hours or days. Therefore, it is not computationally feasible to collect sufficiently many samples to get an accurate ABC posterior distribution using a rejection sampler. Furthermore, a rejection sampler with a small threshold parameter, ϵ , is computationally very inefficient as a large fraction of the sampled function evaluations are rejected. While it is possible to improve the efficiency by applying approaches such as Markov chain Monte Carlo (MCMC) (Marjoram *et al.* 2003), the direct sampling based ABC algorithms are still expected to be computationally too costly for the type of applications targeted in this work.

The inefficiency of the direct sampling based ABC approach can be circumvented by applying Bayesian optimisation to traverse the space of optimised input parameters (Gutmann & Corander 2016; Järvenpää *et al.* 2019). This leads to a probabilistic surrogate model based ABC approach that is observed to be several orders of magnitude more efficient in terms of full function evaluations than the direct sampling based ABC algorithms. At each step of the algorithm, the ABC posterior is estimated using the surrogate model as $P_{\text{ABC}}(\mathbf{x}|D) \propto P(\mathbf{x})\mathbb{P}(\Delta_{\mathbf{x}} < \epsilon)$, where the probability is computed using the probabilistic surrogate model.

2.2. Bayesian optimisation

Bayesian optimisation offers a powerful approach for global optimisation of costly-to-evaluate, non-convex functions, without access to first- or second-order derivatives (Brochu *et al.* 2010; Shahriari *et al.* 2016; Frazier 2018). The problem of finding optimal values, \mathbf{x}_* , for the uncertain input parameters, \mathbf{x} , can be represented as a task of finding the optimum of a nonlinear function $f(\mathbf{x})$ of a compact set \mathcal{A} , called the search space herein. If $f(\mathbf{x})$ represents the discrepancy between the synthetic and measured diagnostic data, then the problem can be formulated as

$$\mathbf{x}_* = \arg \min_{\mathbf{x} \in \mathcal{A}} f(\mathbf{x}). \quad (2.4)$$

The target for a BO algorithm is to be able to traverse the search space efficiently in terms of function evaluations and to find the globally optimum solution by applying prior belief about the optimised function and by balancing exploration and exploitation of the search space (Brochu *et al.* 2010; Shahriari *et al.* 2016; Frazier 2018). In exploitation, samples are collected in regions of the search space that are known to lead to near optimal function values based on prior belief, and in exploration, samples are collected in regions that encompass a large uncertainty.

A standard BO algorithm consists of two main components (Brochu *et al.* 2010; Shahriari *et al.* 2016; Frazier 2018):

- (i) A probabilistic model for the objective function.
- (ii) An acquisition function for recommending the next sampling point.

The probabilistic model represents essentially a low evaluation cost surrogate model for the objective function, and the uncertainties retained in the probabilistic model represent the degree of confidence on the surrogate model predictions. The acquisition function applies the mean and variance of the probabilistic model to balance exploitation and exploration. The collected sample values are then used to update the probabilistic surrogate model.

2.3. Gaussian process regression

The usual choice for the probabilistic model is to use Gaussian process regression (GPR), also known as Kriging (Stein 1999; Rasmussen & Williams 2006, and references therein). Kriging surrogate-based optimisation was previously used for parameter optimisation of plasma transport codes by Rodriguez-Fernandez *et al.* (2018). Other examples of GPR applications in plasma physics can be found in von Nessi *et al.* (2012); Li *et al.* (2013); von Nessi & Hole (2013); Romero & Svensson (2013); Chilenski *et al.* (2015, 2017); Ho *et al.* (2019) and references therein.

Gaussian process regression is a Bayesian regression technique and is very powerful for interpolating small sets of data as well as retaining information about the uncertainty of the regression (Stein 1999; Rasmussen & Williams 2006, and references therein). Gaussian process (GP) is a stochastic process, for which any finite collection of random values has a multivariate normal distribution; GP is specified by the mean function, $m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})]$, and the covariance function, $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))]$ (Rasmussen & Williams 2006)

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \tag{2.5}$$

The GPR represents a family of functions, for which the point to point variance is described by the covariance function. Usually, the mean function is assumed to be zero, although other assumptions are possible. This means that the mean of the prior assumption on variation of the objective function value when propagating from a collected data point is zero. The covariance function or kernel describes the smoothness assumption on the possible functions f . The GP essentially describes a normal distribution over possible functions, f , conditioned with observations $\{(\mathbf{x}_i, f_i), i = 1, \dots, n\}$. Assuming a collection of possibly noisy observations with a Gaussian noise variance σ_n^2 , the posterior probability distribution function of f at point \mathbf{x} is Gaussian with posterior mean $\mu_n(\mathbf{x})$ and posterior variance $v_n(\mathbf{x}) + \sigma_n^2$

$$f(\mathbf{x}) \sim \mathcal{N}(\mu_n(\mathbf{x}), v_n(\mathbf{x}) + \sigma_n^2). \tag{2.6}$$

Assuming a zero mean function, $m(\mathbf{x})$, the mean and variance can be obtained as

$$\mu_n(\mathbf{x}_*) = \mathbf{k}_*^\top \mathbf{K}_n^{-1} \mathbf{f}_n, \tag{2.7}$$

$$v_n(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top \mathbf{K}_n^{-1} \mathbf{k}_*, \tag{2.8}$$

where \mathbf{k}_* represents the vector of covariances between the test point, \mathbf{x}_* , and the n observations, \mathbf{f}_n is a vector of the n observations and \mathbf{K} is the covariance matrix (Rasmussen & Williams 2006; Gutmann & Corander 2016). Since the function evaluations in this work are deterministic, the σ_n term is constrained to a low value that does not impact

the predictions. An estimate of the likelihood at w is given by

$$\mathbb{P}(\Delta_w < \epsilon) = F\left(\frac{\epsilon - \mu_n(w)}{\sqrt{v_n(w) + \sigma_n^2}}\right), \quad (2.9)$$

where $F(x)$ is the cumulative distribution function of $\mathcal{N}(0, 1)$. In this work, we will set ϵ equal to the current optimal value provided by the probabilistic surrogate model, which is also the default in BOLFI (Gutmann & Corander 2016).

A key step in building a GP regression is to select the covariance function or kernel (Rasmussen & Williams 2006). Usually the default choice is the radial basis function (RBF), also known as squared exponential or Gaussian kernel

$$k_{\text{RBF}}(x_i, x_j) = \sigma_f^2 \exp\left(-\sum_{k=1}^d \frac{(x_{i,k} - x_{j,k})^2}{2l_k^2}\right), \quad (2.10)$$

where $\mathbf{l} = [l_1, \dots, l_d]$ is a vector of covariance length scales for each dimension, d , and σ_f^2 is the variance. In the applications in this work, the single constant \mathbf{l} was observed to be often too restrictive and a rational quadratic (RQ) kernel was used instead

$$k_{\text{RQ}}(x_i, x_j) = \sigma_f^2 \left(1 + \sum_{k=1}^d \frac{(x_{i,k} - x_{j,k})^2}{2\alpha l_k^2}\right)^{-\alpha}, \quad (2.11)$$

which is equivalent to summing many RBF kernels with varying \mathbf{l} . The hyperparameter α represents the relative weighting between large and small \mathbf{l} values. Before applying the model, the hyperparameters (\mathbf{l} , α , σ_f^2 , σ_n) must be optimised. These can be estimated by maximising the marginal log-likelihood (Rasmussen & Williams 2006). The GPR library used in this work as well as in BOLFI is the Python GP framework, GPy (GPy since 2012), which encompasses the applied optimisation routines.

2.4. Acquisition function

The acquisition function applies the mean and variance of the probabilistic surrogate model to recommend the next sampling point for the objective function (Brochu *et al.* 2010; Shahriari *et al.* 2016; Frazier 2018). The acquisition functions are typically constructed to recommend sampling points that either encompass the optimal predicted mean for the objective function, exploitation, or large uncertainty, exploration. The sampling point is selected by optimising the acquisition function. Several acquisition functions have been developed and can be found in the reviews in Brochu *et al.* (2010), Shahriari *et al.* (2016) and Frazier (2018). The acquisition function used in sequential sampling in this work is the lower confidence bound selection criterion (LCBSC), which is also the default acquisition function in BOLFI (Brochu *et al.* 2010; Srinivas *et al.* 2010; Gutmann & Corander 2016, and references therein). This function can be written as

$$\mathcal{A}_n(\mathbf{x}) = \mu_n(\mathbf{x}) - \sqrt{\eta_n^2 v_n(\mathbf{x})}, \quad (2.12)$$

$$\eta_n^2 = 2 \ln\left(\frac{n^{2d+2} \pi^2}{3\epsilon_\eta}\right), \quad (2.13)$$

where the coefficient η_n is the trade-off parameter between exploration and exploitation, $\epsilon_\eta \in (0, 1)$ is a constant chosen by the user and d is the dimensionality of the search space.

Optimising the acquisition function provides a deterministic answer for the next sampling point. An example of the application of this acquisition function is shown in § 3.2.

Since the next sampling point is obtained deterministically for a given state of the surrogate model, the approach is naturally sequential: (i) the objective function is evaluated for the sampling location provided by the optimum of the acquisition function, (ii) the GPR surrogate model is updated and (iii) the acquisition function is optimised again, using the updated GPR, to recommend the next sampling point. However, with complicated, multi-dimensional optimisation tasks with computationally time consuming function evaluations, it would be more attractive to conduct several objective function evaluations in parallel with each other to reduce the overall time consumption of the optimisation task, especially when suitable high-performance computing resources are available.

To conduct parallel BO, stochastic acquisition rules can be used. Various approaches for parallel BO and batch acquisition have been developed (Thompson 1933; Chapelle & Li 2011; Järvenpää *et al.* 2018; Kandasamy *et al.* 2018; Järvenpää *et al.* 2019; Palma *et al.* 2019). In this work, the randmaxvar approach developed by Järvenpää *et al.* (2019) is used as the stochastic acquisition method. The approach is based on the maxvar acquisition rule also presented in Järvenpää *et al.* (2019). The maxvar acquisition method recommends a sample in a location that encompasses the maximum variance of the unnormalised ABC posterior. Basically, due to limited information, represented by the collected samples, there is uncertainty in the GPR representation of the objective function and this uncertainty is propagated as uncertainty of the unnormalised ABC posterior. The maxvar method aims to collect samples that lead to maximum reduction of this uncertainty. In the stochastic version of this method, samples are collected from the distribution that represents the variance of the unnormalised ABC posterior. Since samples are collected stochastically, several samples can be collected without updating the GPR surrogate in between the samples, enabling parallelisation. Furthermore, sampling can be done asynchronously by simply updating the GPR surrogate whenever new results are added to the dictionary of collected samples and by sampling new values whenever an idle processor becomes available.

3. Application to a JET RE experiment

This section describes proof-of-principle applications of the methodology discussed in § 2 for a RE experiment at JET. Section 3.1 describes the investigated JET plasma discharge and the DREAM set-up, § 3.2 documents a one-dimensional (1-D) proof-of-principle optimisation, §§ 3.3, 3.4 and 3.5 extend the search to four- to seven-dimensional search spaces and § 3.6 evaluates the number of samples required for convergence as a function of dimensionality of the search.

3.1. Simulated experiment and DREAM set-up

The BO and ABC approach discussed in the previous section is applied for DREAM RE simulations of current quench (CQ) in the disruption of JET discharge #95135. This was a deuterium limiter plasma with an argon massive gas injection induced disruption, described in detail by Reux *et al.* (2021) and Brandström (2021). DREAM is a numerical tool for self-consistently simulating the evolution of temperature, poloidal flux and impurity densities, along with the generation and transport of REs in tokamak disruptions (Hoppe *et al.* 2021). The DREAM simulations in this manuscript are similar to those presented by Hoppe *et al.* (2021) and Brandström (2021), with the exception that only a fluid model for RE electrons is used here, as kinetic simulations were not necessary for the proof-of-principle of the Bayesian approach. For a full description of the physics model in DREAM, we refer the reader to Hoppe *et al.* (2021).

The simulations are started at the peak of the total plasma current, I_p , obtained during the disruption. An instantaneous thermal quench (TQ) is assumed, after which all background plasma quantities, except the electron temperature, T_e , are evolved self-consistently. The post-disruption T_e is instead given as an uncertain input parameter. The background plasma density is obtained from pre-disruption measurement with the high resolution Thomson scattering (Pasqualotto *et al.* 2004; Frassinetti *et al.* 2012). Even though the uncertainty of the electron density measurement could be taken into account in the Bayesian approach, for simplicity, we neglect it here. The argon density in the plasma is obtained from the estimated amount of injected argon, volume of the vessel and fraction of argon that is assimilated. Argon is assumed to be uniformly distributed in the plasma. While the injected amount can be obtained from the experiment, $N_{\text{Ar}} \sim 8 \times 10^{20}$ atoms (Brandström 2021), the assimilated fraction is given as an uncertain input parameter, f_{Ar} .

In our simulations, the initial total plasma current I_p (combination of ohmic and prescribed runaway seed current) is constrained to match the experimentally measured peak value, $I_p \sim 1.42$ MA, by adjusting the initial electric field in order for the avalanche multiplication factor to be constrained. The electric field is evolved self-consistently, during the CQ and RE plateau simulations. The conductivity of the wall is controlled by a characteristic wall time, $\tau_{\text{wall}} = L_{\text{ext}}/R_{\text{wall}}$, that is provided as an uncertain input parameter. Here, L_{ext} is the external inductance and R_{wall} the resistance of the wall. The full list of uncertain input parameters are T_e , f_{Ar} , RE seed distribution and τ_{wall} . In this study, these parameters are allowed to vary independently from each other, while in reality T_e and f_{Ar} would be closely linked due to the strong radiation losses from argon. It is acknowledged that this risks the system having the capability to overfit the model, by applying mutually inconsistent input parameters. However, this is considered acceptable in this proof-of-principle study and further constraints to the degrees of freedom will be investigated in follow-up studies. In addition, the RE seed distribution is scaled with a multiplier such that the RE plateau current matches the experimentally measured value. This is done iteratively with a binary search algorithm. Runaway electrons are generated through a combination of the Dreicer, hot-tail and avalanche mechanisms. Due to the presence of magnetic perturbations during the TQ a part of the seed runaways are transported out of the plasma. After the thermal quench, when the conductivity is very low, the plasma current will be carried essentially entirely by the runaways. The final plateau runaway current profile is determined by the surviving post-TQ runaway seed and the avalanche multiplication factor. The latter is a robust quantity and depends on the initial and final plasma currents (or more precisely the change in the poloidal flux), see e.g. Boozer (2015) and Hesslow *et al.* (2019). Since the magnitude of the runaway seed current is uniquely determined by the runaway plateau current and the avalanche multiplication factor it does not need to be added to the search space of the BO algorithm, while the seed profile must since it cannot be similarly directly inferred from experiment. To confirm this, a test was conducted to include the seed scale into the search space of the BO algorithm in the 7-D search (§ 3.4), leading essentially to the same result as obtained with the binary search. The algorithm is initialised by multiplying or dividing the multiplier by 10^3 at each iteration until the predicted plateau current passes the experimentally measured value, after which binary search is applied to converge to the optimum. If the multiplier is reduced below a small value, $\sim 10^{-3}$, the search algorithm stops without finding a solution that matches the RE plateau current and the objective function value at the point of reaching that threshold is propagated through the algorithm. This is simply due to the fact that it cannot be assumed that for all input values within the search space it is possible to find a multiplier that would enable matching the RE plateau current.

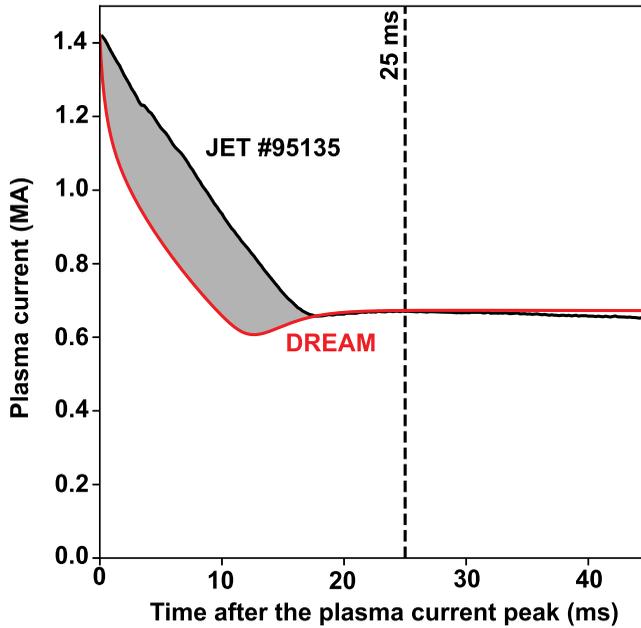


FIGURE 1. Illustration of the discrepancy between the measured (black) and simulated (red) plasma current for an example simulation. The grey area represents the L1-norm used as the discrepancy metric in this study. The vertical dashed line indicates the temporal extent of the application of the discrepancy function.

In the proof-of-principle examples that follow, the objective function is chosen as the L1-norm of the discrepancy between the measured and predicted I_p during the CQ (figure 1)

$$\Delta = \|I_p^{\text{measured}} - I_p^{\text{DREAM}}\|_1. \tag{3.1}$$

Effectively this discrepancy function calculates the area between the two curves in figure 1. To avoid accumulating excessive integration within the RE plateau, the discrepancy is only calculated between 0 and 25 ms from the current peak.

The chosen proof-of-principle case is relatively simple to facilitate exploration of the methodology. With the released Python code, this can also provide a low threshold environment for the scientific community to become familiar with this methodology. A follow-up study will apply these methods to a more complicated validation exercise with a significantly higher computational overhead. The BO approach can also naturally extend to multiple objectives and this will be explored in follow-up studies.

3.2. One-dimensional search space

The first proof-of-principle test is to apply the Bayesian approach to search for the post-TQ T_e that minimises the discrepancy function (3.1). The other uncertain input parameters are fixed as $f_{Ar} = 15\%$, uniform RE seed profile and $\tau_{\text{wall}} = 5$ ms. A RQ kernel is used and the length scale is constrained to be below 1.0. This is done to avoid the model becoming overconfident in regions that have not been sampled yet. If the length scale converges to a large value, it can suppress exploration prematurely and prevent the algorithm from finding the optimal solution. More generally, it seems that, in BO algorithms, it is probably better to have a model that has the capability to overfit rather than underfit the data. The BO algorithm is interpolating solutions within the search space and an overfitting model will

just encourage exploration while an underfitting model might not have the generalisation capability to fit the solution near the optimum. The acquisition function is the LCBSC (2.12) with exploration constant $\epsilon_\eta = 0.2$. The search space for T_e is bounded between 1.0 and 20 eV. A uniform prior uncertainty distribution is assumed.

Within less than 10 iterations, the algorithm starts to converge to the optimum value (figure 2). The first 3 samples are collected randomly from the uniform prior distribution. After these are collected, the GPR surrogate model is fitted to the data (figure 2). The mean and variance of the GPR are applied by the acquisition function to recommend the next sampling location. The algorithm proceeds by choosing the location of the minimum value of the acquisition function. By proceeding like this, the algorithm converges near the optimum value with a narrow confidence interval around the sample number 7 to 8 (figure 3a). Since the prior distribution is uniform, the posterior probability distribution can be obtained from the GPR by applying (2.9) (figure 3b). Finally, the simulated current value with the temperature value providing the highest posterior probability, $T_e \sim 5.84$ eV, can be compared with the experimental measurements (figure 3c). The results indicate that by only adjusting the constant post-TQ T_e , the model is not able to reproduce the experimentally measured CQ rate. Alternatively, these results indicate that the background plasma resistivity is changing during the CQ.

3.3. Five-dimensional search space

After the 1-D proof-of-principle, the next step is to extend the search space to include the other uncertain parameters as well. Since the RE seed is a 1-D profile, it is parameterised as the probability distribution function of the gamma distribution

$$f_{\text{RE}}(r; \alpha, \beta) = \frac{\beta^\alpha r^{\alpha-1} e^{-\beta r}}{\Gamma(\alpha)}, \quad (3.2)$$

where α and β are free parameters and $\Gamma(\alpha)$ is the gamma function (figure 4). The intention is to provide a general parameterisation that constrains as little as possible the possible RE seed profile shapes. The search spaces for both α and β are set as uniform distributions between 0.001 and 10. Ten randomly sampled RE seed profile shapes are shown in figure 4. It can be clearly seen that this parameterisation allows flexible representations of profiles peaking at the centre, middle or edge of the plasma. The search space for the argon assimilation fraction is set as uniform between 0.001 % and 100 %. The search space for the characteristic wall time is set such that $\ln(\tau_{\text{wall}}/1 \text{ ms})$ is sampled uniformly between 0 and 7. As a result, the τ_{wall} values range between approximately 1 and 1100 ms. This allows sampling for very large τ_{wall} parameters exceeding 1.0 s, while encouraging collection of samples at low values.

Due to the expanded volume of the search space, more samples are expected to be needed than in the 1-D example. Therefore, the randmaxvar acquisition function was used with a batch size of 10 samples conducted in parallel. The first 50 samples were collected by sampling the search space randomly, after which the acquisition function was used to recommend sampling locations.

Similar to the 1-D search, a RQ kernel is used. The kernel parameters are restricted such that the power parameter in the RQ kernel was constrained to be between 10^{-10} and 0.03. The length scale constraints are altered for every batch of 10 samples. For even round batches, the length scales are constrained to be positive and manually initialised as the distance of the search domain for the dimension divided by the number of collected samples. After this preconditioning step, GP optimisation is conducted. There is no maximum length scale set up during the even rounds. During the odd round batches, the

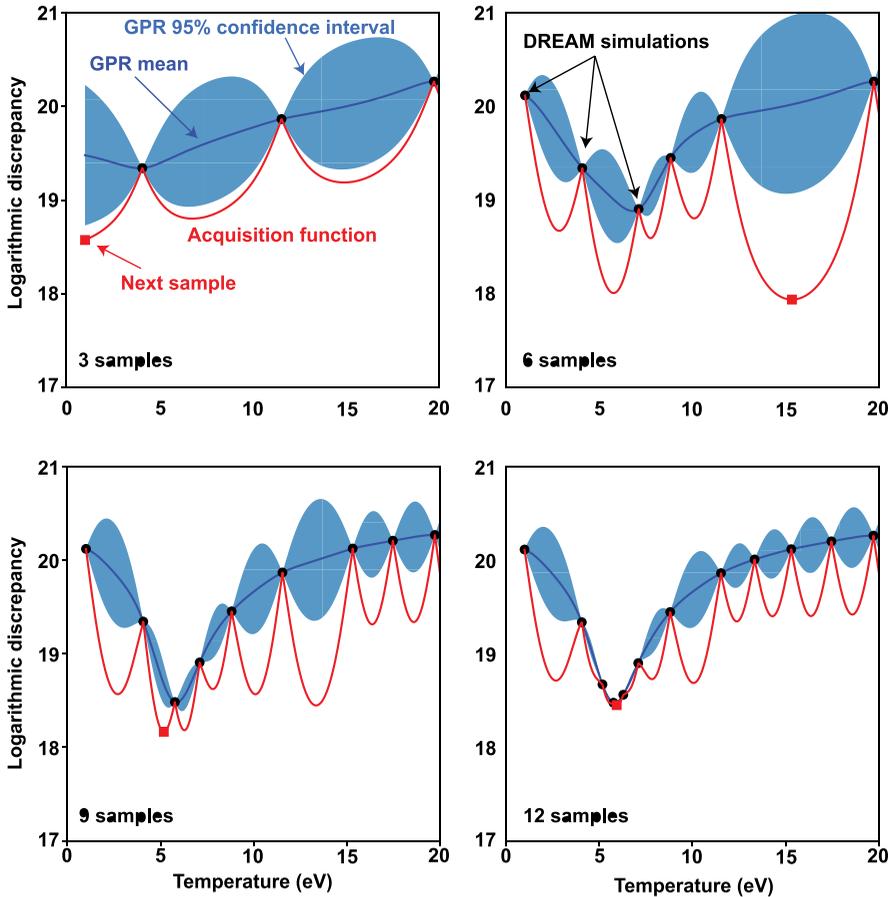


FIGURE 2. Illustration of the progress of the BO algorithm in the 1-D example after 3, 6, 9 and 12 samples. The dark blue lines represent the mean and the light blue regions represent the 95 % confidence interval of the GPR. The red lines represent the acquisition function and the red squares illustrate the optimum of the acquisition function that provides the next sampling location. The black circles represent the collected objective function values obtained through DREAM simulations.

length scales for variations are constrained to be below 1 for the T_e and f_{Ar} dimensions, below 0.5 for the α and β dimensions and below 0.1 for the $\ln(\tau_{wall})$ dimension. The lower limits for the length scales were set to 10^{-3} . The even rounds perform essentially automatic relevance determination obtaining very long length scales for dimensions that do not show a significant impact on the objective function value, guiding the search to prioritise optimisation of input parameters that have a more significant impact on the objective function value. However, this approach alone would risk the surrogate model becoming overly confident early in the search and stop exploration for input parameters deemed unimportant. To counteract this risk, the odd rounds apply restrictions of length scales, such that the algorithm understands to explore input parameter values, which would simply be extrapolated and interpolated over by the long length scale surrogate model.

The sampling algorithm finds clear optima in the search space for the background plasma temperature and characteristic wall time (figure 5). It can be observed that there is a global optimum at T_e around 5–7 eV and τ_{wall} less than approximately 2.7 ms

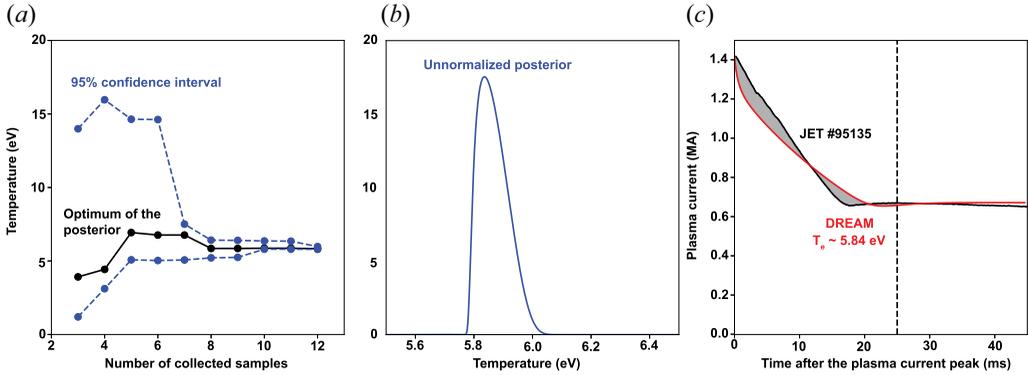


FIGURE 3. One-dimensional search space results. (a) Optimum of the posterior (black) and 95 % confidence interval (blue dashed) as a function of number of collected samples. (b) Unnormalised posterior distribution after collecting 12 samples. (c) Measured (black) and predicted (red) plasma current as a function of time. The predicted value is conducted at the optimum of the posterior after collecting 12 samples.

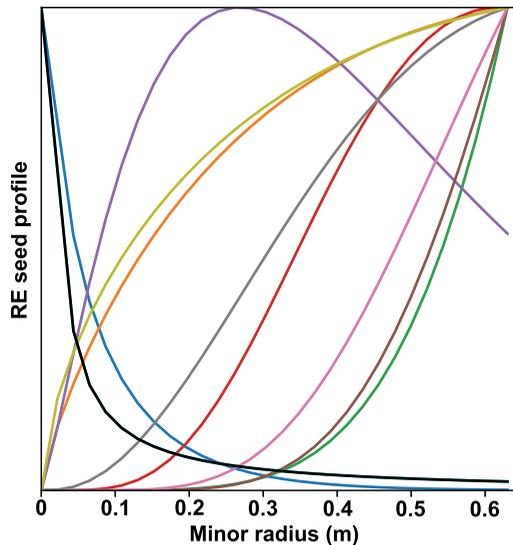


FIGURE 4. Ten randomly sampled RE seed profiles.

($\ln(\tau_{\text{wall}}/1 \text{ ms}) < 1.0$), and a local optimum at T_e approximately 15–20 eV and τ_{wall} larger than approximately 50 ms ($\ln(\tau_{\text{wall}}/1 \text{ ms}) > 4.0$) (figure 5 red and black ellipses). The two solution branches can also be observed in the plot as a function of the argon assimilation fraction, such that the low temperature solution branch reaches the optimum at values above 20 %, while the higher temperature solution branch at lower values between 5 % and 30 % (figure 5c). The shape of the RE seed distribution does not seem to impact the discrepancy significantly (figure 5d,e).

After 290 samples, the algorithm estimates the global optimum to be $T_e \approx 6.1$ eV, $f_{\text{Ar}} \approx 54.2$ %, $\tau_{\text{wall}} \approx 1.1$ ms, $\alpha \approx 7.9$ and $\beta \approx 4.3$, where the 95 % confidence intervals for f_{Ar} , α and β span most of the search space (figure 6). It should be noted that the problem set-up, where the assimilated argon fraction is kept independent from the temperature, is

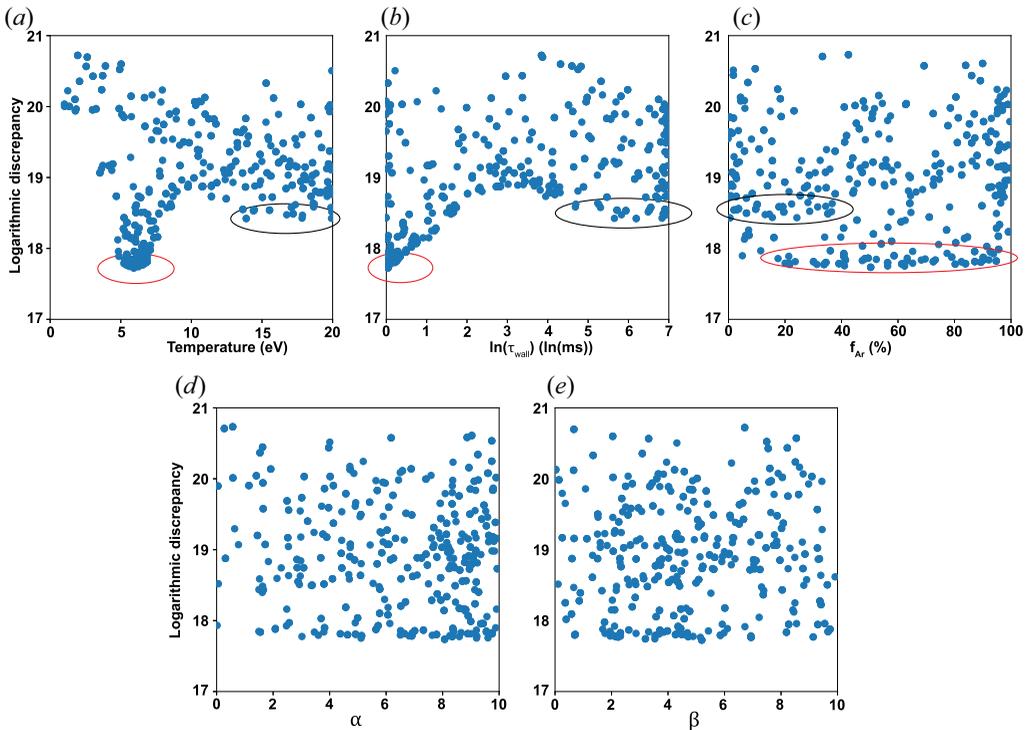


FIGURE 5. Discrepancies of the collected samples as a function of the dimensions of the search space: (a) temperature, (b) logarithmic characteristic wall time, (c) argon assimilation fraction and parameters of the runaway seed distribution (d,e). The total number of collected samples is 290. The two optima are highlighted with red (global) and black (local) ellipses.

likely to impact this result, and a general conclusion about insensitivity of the plasma current on f_{Ar} should not be drawn here. To further investigate this, an f_{Ar} scan was conducted, while keeping all the other parameters at their inferred optimal values. It is observed that, at these low temperatures, scanning f_{Ar} from 40% to 100% only increases the average Z_{eff} in the plasma from approximately 1.5 to 1.9. While this does increase the CQ rate, the impact is sufficiently small that the optimisation algorithm is able to counteract this by small changes of T_e and τ_{wall} . The local uncertainties for the global optimum can be obtained by evaluating the local properties of the approximate posterior. The approximate posterior can be extracted from the probabilistic surrogate model by applying (2.9). While a global ABC posterior can be obtained through, for example, MCMC sampling, the local inverse uncertainty near the global optimum is likely of more practical interest in fusion energy research and can be obtained directly from the GPR representation of the posterior. This type of analysis was done by evaluating the 1-D posterior distribution for each search dimension from the global optimum, which can be integrated to obtain confidence intervals (figure 6). It should be noted that this analysis does not take into account the secondary optimum at higher temperatures as that would require nonlinear analysis of the approximate posterior, while MCMC sampling of the posterior would collect some distribution in that area also. However, since multimodality of the optimised function can be observed from the discrepancy plot already (figure 5), it is considered more important to obtain local uncertainty estimations near the global optimum than sample the full approximate posterior. For completeness, MCMC sampling

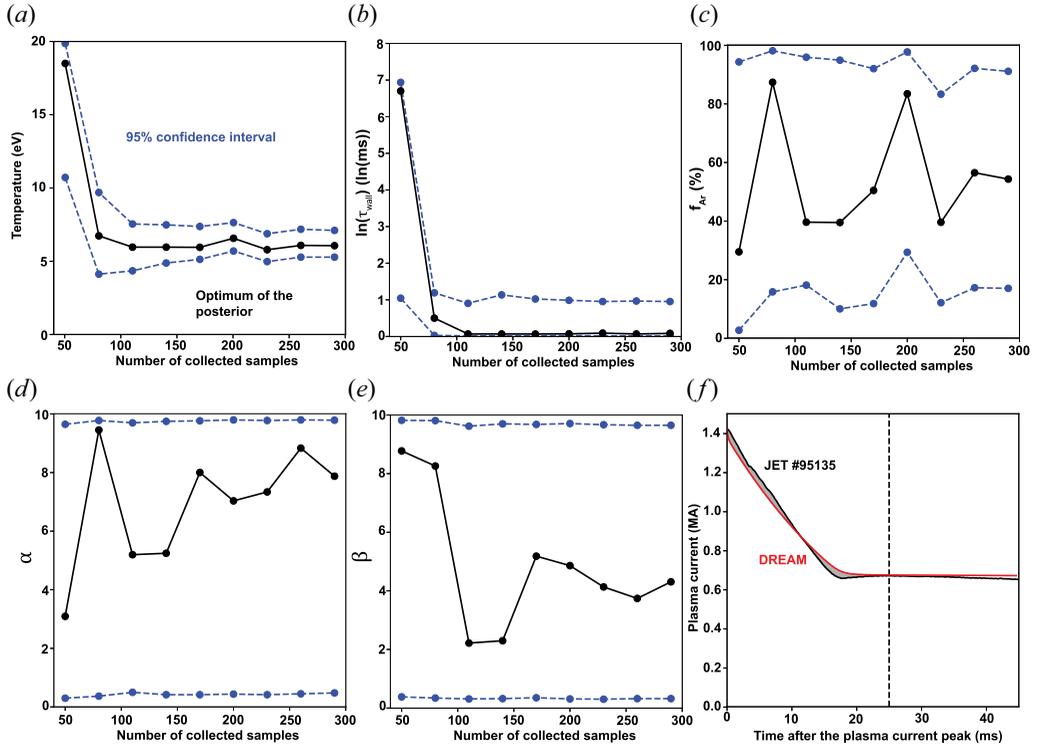


FIGURE 6. Convergence of the posterior distribution near the global optimum for temperature (a), logarithmic characteristic wall time (b), argon assimilation fraction (c) and RE seed profile parameters α and β (d,e) as a function of number of collected samples. The black dots illustrate the current optimum of the posterior distribution and the blue dots with dashed lines the 95 % confidence interval. (f) The predicted total plasma current (red) with the optimal input parameters after 290 samples compared with the experimental plasma current (black). The GPR in this convergence figure applies the length scale restrictions as used for the odd round samples.

was conducted for the approximate posterior and the resulting distributions for T_e and $\ln(\tau_{\text{wall}})$ show the global optimum, also visible in (figure 6), as well as the secondary local optimum at higher temperatures (figure 7).

Analysing the convergence of the search it can be observed that, after approximately 80–110 samples, the posterior distributions for T_e and $\ln(\tau_{\text{wall}})$ have converged (figure 6a,b). For the other parameters, the 95 % confidence intervals remain large, indicating that the discrepancy value is not very sensitive to these input parameters, as was observed in the discrepancy plot already (figure 5). The step observed at 200 samples highlights the stochastic nature of the GPR surrogate model fitting. Depending on the initial conditions, the optimisation algorithm might find somewhat different hyperparameters for the GPR, leading to a different shape of the approximate posterior and shift of the optimum and boundaries of the confidence interval as well. However, the large steps of the optimum for f_{Ar} , as well as the large steps at other sample numbers for α and β , are a result of the relatively flat posterior distribution shapes, while posterior shape for T_e and $\ln(\tau_{\text{wall}})$ are not changed significantly.

Finally, the predicted and measured total plasma current are compared for the global optimum extracted by the algorithm (figure 6f). Comparing the result with the optimum

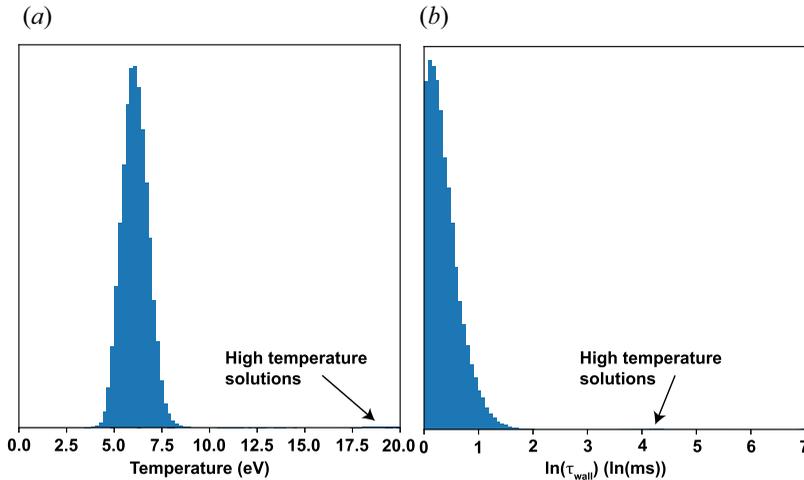


FIGURE 7. The MCMC samples of the posterior distribution for T_e (a) and $\ln(\tau_{\text{wall}})$ (b).

case in the 1-D search space example, it can be clearly seen that, with a 5-D search space, the algorithm is able to obtain a significantly better fit to the experimentally measured current (figures 3c and 6f). In the 1-D example, the current was reducing significantly faster than experimentally measured during the early part of the CQ and the end of the CQ happened several ms later than experimentally measured, such that the average L1-norm discrepancy was minimised, while the rate of change of plasma current was poorly matched (figure 3c). On the other hand, in the 5-D example, the initial drop is also faster than measured experimentally, but soon the rate of change of the plasma current is matched to the experimentally measured rate of change, such that the end of CQ happens near the experimentally measured end of the CQ when the L1-norm between the two currents is minimised (figure 6f). The fact that the current is reducing faster than experimentally measured during the early parts of the CQ indicates that the plasma resistance that on average matches the plasma current evolution probably overestimates the resistance during early parts of the CQ. To address this, the final proof-of-principle test conducted in this manuscript is to allow linear variation of T_e during the simulation by extending the search space to seven dimensions.

3.4. Seven-dimensional search space

As a next extension of the search, a parameterised variation of the background plasma T_e during the simulation is allowed. Linear variation of T_e is assumed, such that the search space is extended to seven dimensions, by adding final temperature, $T_{e,\text{final}}$, and the time at which $T_{e,\text{final}}$ is reached, t_{final} . After t_{final} is reached, T_e is assumed to stay constant. Same search spaces are used for the initial and final T_e , and the search space for t_{final} is set as uniform between 1 and 44 ms.

The RQ kernel is used in the GPR. The power of the kernel was restricted similarly to the set-up in the 5-D search. Similar to the 5-D search, the length scale restrictions were altered between even and odd round batches. Batch size was set to 50. For the odd round batches, the length scale constraints are similar to those in the 5-D search with the same length scale constraint applied for the initial and final T_e . For the t_{final} the minimum length scale is set as 10^{-3} ms and maximum as 1 ms.

After 950 samples, the local 95% confidence intervals around the optimum point, recommended by the algorithm, are $T_{e,\text{initial}} \in [8.5, 11.7]$ eV, $T_{e,\text{final}} \in [3.3, 5.8]$ eV, $t_{\text{final}} \in [11, 19]$ ms, $\tau_{\text{wall}} \in [1.1, 2.1]$ ms, $f_{\text{Ar}} \in [48, 97]$ %, $\alpha \in [0.2, 9.4]$ and $\beta \in [0.5, 9.7]$ (figure 8). The optimum point is $T_{e,\text{initial}} \approx 9.5$ eV, $T_{e,\text{final}} \approx 4.3$ eV, $t_{\text{final}} \approx 14$ ms, $\tau_{\text{wall}} \approx 1.4$ ms, $f_{\text{Ar}} \approx 76$ %, $\alpha \approx 4.3$ and $\beta \approx 2.2$. With these input parameters, the predicted CQ rate is very close to the measured values (figure 8f). Initially, the rate is faster than measured and also the transition to runaway plateau in the simulation occurs approximately 1 ms earlier than measured (figure 8f). However, between 3 and 12 ms, the predicted current is nearly exactly on top of the measured current.

3.5. Constraining τ_{wall} to 5 ms

In both the 5-D and 7-D searches, the algorithm found optimum τ_{wall} around 1.1–1.4 ms. This result was somewhat surprising when the conventional prior expectation would have suggested higher values in the range of 5–10 ms. Since the optimisation algorithm finds the set of parameters that minimise the discrepancy, it is possible that the algorithm compensates for missing physics in the model by reducing τ_{wall} below values that would actually be realistic. As a final test, further 4-D and 6-D optimisations were conducted, where τ_{wall} was fixed to 5 ms.

In the 4-D search, the optimum point recommended by the algorithm is $T_e \approx 7.0$ eV, $f_{\text{Ar}} \approx 52$ %, $\alpha \approx 4.6$ and $\beta \approx 5.5$ (figure 9a). The local 95% confidence intervals are $T_e \in [6.4, 8.0]$ eV and $f_{\text{Ar}} \in [32, 91]$ %, $\alpha \in [0.3, 9.6]$ and $\beta \in [0.3, 9.7]$. In the 6-D search, the optimum point recommended by the algorithm is $T_{e,\text{initial}} \approx 16.2$ eV, $T_{e,\text{final}} \approx 5.9$ eV, $t_{\text{final}} \approx 10$ ms, $f_{\text{Ar}} \approx 82$ %, $\alpha \approx 1.3$ and $\beta \approx 6.1$ (figure 9b). The local 95% confidence intervals are $T_{e,\text{initial}} \in [13.9, 19]$ eV, $T_{e,\text{final}} \in [5.1, 6.7]$ eV, $t_{\text{final}} \in [8, 12]$ ms, $f_{\text{Ar}} \in [48, 98]$ % and $\alpha \in [0.1, 8.8]$, $\beta \in [0.5, 9.7]$.

As f_{Ar} , α and β do not impact the discrepancy significantly, the best match obtained by the 4-D search seems very similar to the best match obtained by the 1-D search (figures 3c and 9a). When allowing linearly varying background plasma T_e , the algorithm is able to find a solution that matches the experimentally measured plasma current nearly as well as in the 7-D search (figures 8f and 9b). However, when fixing $\tau_{\text{wall}} = 5$ ms, the recommended initial T_e is increased from 9.5 to 16.2 eV, and t_{final} reduced from 14 ms to 10 ms, highlighting the nonlinear dependencies between the optimal input parameters.

3.6. Convergence as a function of dimensions

The computational challenge of the optimisation task increases with the number of dimensions of the search space. Figure 10 illustrates the discrepancy as a function of sample numbers for the 4-D, 5-D, 6-D and 7-D search tasks in this work. Evaluating the convergence based on the sample number after which the minimum discrepancy saturates, the 4-D search converges after approximately 40 samples, the 5-D search after approximately 80 samples, the 6-D search after approximately 250 samples and the 7-D search after approximately 300 samples. Beyond this point, increasing sample numbers will reduce the uncertainty of the posterior distribution while the minimum discrepancy is not reduced anymore.

Comparing with a grid search of eight samples for each dimension, the BO algorithm is very efficient (figure 11). Beyond four dimensions, the grid search algorithm would be calling for over 10 000 samples and soon become intractable. The Bayesian approach, on the other hand, obtains samples near the minimum discrepancy after a few hundred samples even in the case of the seven dimensional search space.

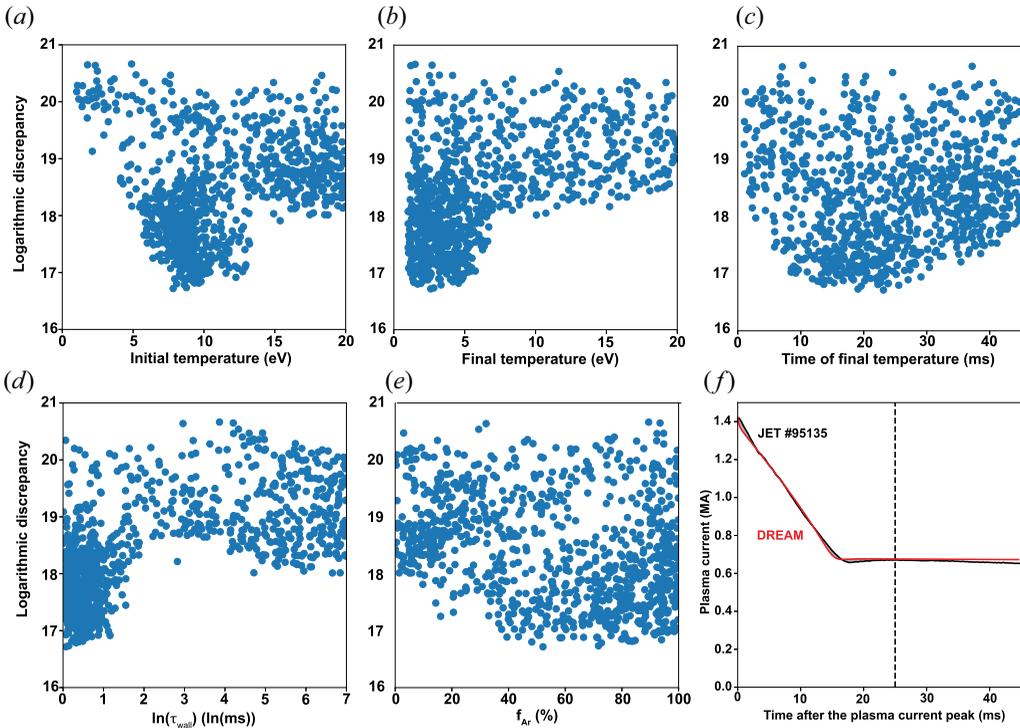


FIGURE 8. Discrepancies of the collected samples as a function of the dimensions of the search space: (a) initial temperature, (b) final temperature, (c) time at which the final temperature is reached, (d) logarithmic characteristic wall time, (e) argon assimilation fraction. The α and β parameters are not shown as those do not show any significant impact on the discrepancy. (f) The predicted total plasma current (red) with the recommended optimal input parameters compared with the experimental plasma current (black). The total number of collected samples is 950.

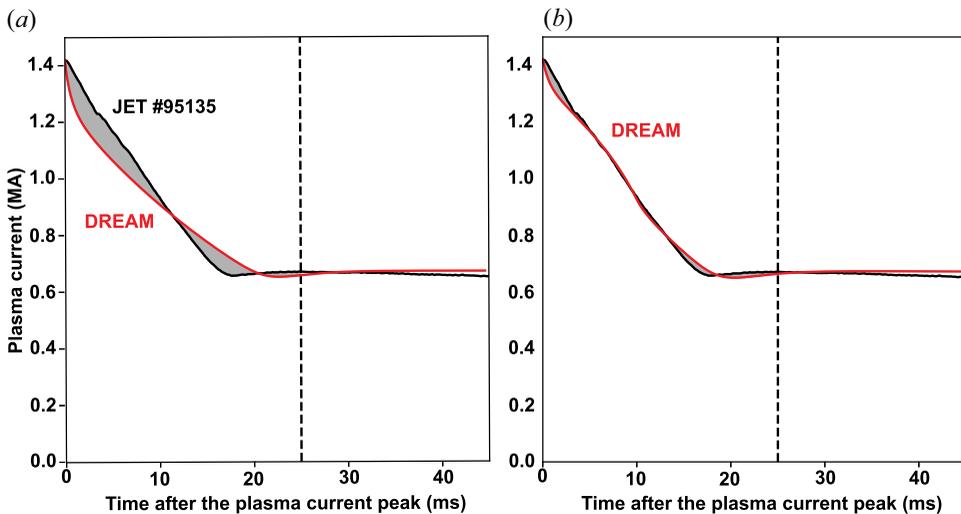


FIGURE 9. The predicted total plasma current (red) with the recommended optimal input parameters in the 4-D (a) and 6-D (b) search tasks compared with the experimental plasma current (black).

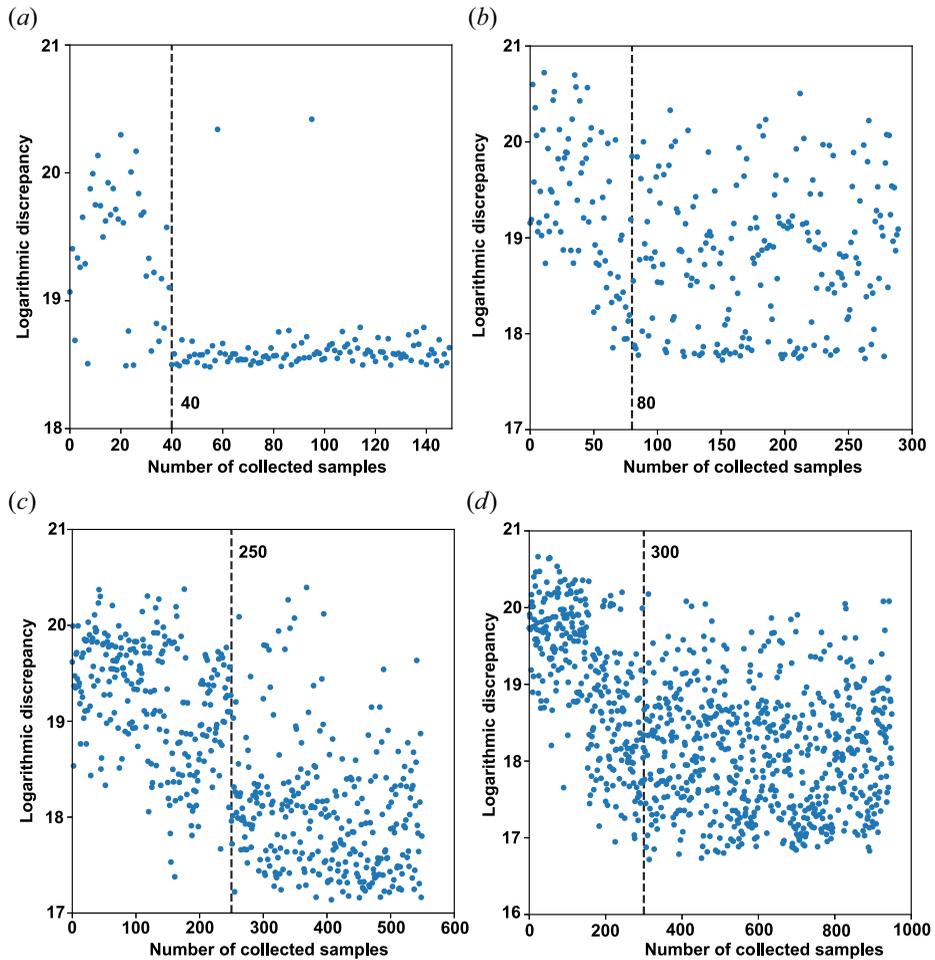


FIGURE 10. Discrepancy as a function of sample number for the 4-D (a), 5-D (b), 6-D (c) and 7-D (d) search tasks. The vertical dashed lines illustrate the approximate point when the minimum discrepancy saturates.

4. Summary

A Bayesian approach has been explored for validation of RE simulations. Many of the simulation tools applied in fusion energy research require the user to specify several input parameters that are not constrained by the available experimental information. Bayesian inference algorithms offer a promising approach to determine these free parameters with uncertainty quantification and are less subject to user bias than approaches based on manual parameter calibration. The main challenge in using an algorithmic approach to parameter calibration is the computational cost of simulating enough samples to construct the posterior distributions for the uncertain input parameters. By using probabilistic surrogate modelling, through GPR, with BO, it is possible to reduce the number of required simulations by several orders of magnitude. This type of BO framework was implemented in this work for a disruption RE analysis model, and explored for CQ simulations for a JET plasma discharge with an argon induced disruption. The algorithm is able to find optimal input parameters with uncertainties in 1-D to 7-D proof-of-principle

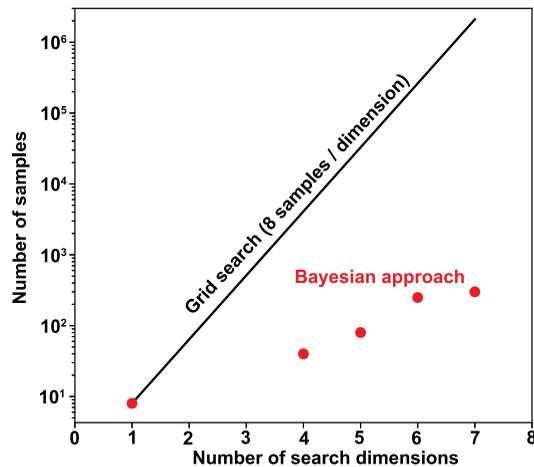


FIGURE 11. Number of samples as a function of number of search dimensions: grid search with 8 samples for each dimension (black line), Bayesian approach in this work (red circles).

cases, and is several orders of magnitude more sample efficient than a regimented grid search algorithm would have been.

The proof-of-principle application in this manuscript was based on fluid RE simulations with DREAM. These simulations are computationally relatively light, requiring only a few minutes of computing per sample. However, many simulations in fusion research require several hours, days or weeks of computing time. Obviously, even these relatively sample efficient BO algorithms will eventually begin to struggle with increasing search space dimensions, number of objective functions and computational requirements of the optimised model. Nevertheless, the discussed methodology does provide a principled, algorithmic approach to parameter fitting for these complicated models that is more sample efficient than, for example, grid or random search. Furthermore, running a few hundred simulations of a computational model that takes hours to converge is still a task that is manageable with a relatively standard high-performance computing environment.

Surrogate model specification is central to the performance of the search algorithm. Using the GP approach, the kernel parameters need to be appropriately constrained for the surrogate model to provide meaningful guidance for the search through the acquisition function. An overly smooth kernel with long maximum correlation length scales can make the surrogate model overly confident and not find the actual global optimum. On the other hand, limiting the length scales to small values will encourage exploration but also require more iterations for convergence. Finding the appropriate surrogate model specifications is an area that requires attention from the user of these algorithms. The most appropriate constraints are likely to be specific to each search task. Furthermore, both specifying appropriate kernel constraints and diagnosing potential issues with kernel constraints become more challenging with increasing number of search dimensions. Therefore, it would be desired to find default kernel constraints that are likely to work acceptably well in most circumstances. The approach chosen here was to alternate the kernel constraints at specific sample intervals between unconstrained but positive length scales and length scales constrained to be below a certain threshold that is a fraction of the width of the search dimension. By alternating the kernel constraints, the risk of the algorithm either oversmoothing a region or getting into a mode of infinite exploration is reduced. However,

more generally applicable methods for constraining the surrogate model are likely to exist, and could improve the performance of the algorithm further.

Acknowledgements

The authors wish to acknowledge CSC – IT Center for Science, Finland, for computational resources.

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

Funding

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 authors only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. This work was carried out under the EUROfusion E-TASC Advanced Computing Hub project. The work of Eero Hirvijoki was supported by the Academy of Finland grant no. 315278. The work was supported in part by the Swedish Research Council (Dnr. 2018-03911).

Declaration of interests

The authors report no conflict of interest.

REFERENCES

- BOOZER, A.H. 2015 Theory of runaway electrons in ITER: equations, important parameters, and implications for mitigation. *Phys. Plasmas* **22** (3), 032504.
- BOOZER, A.H. 2018 Pivotal issues on relativistic electrons in ITER. *Nucl. Fusion* **58**, 036006.
- BRANDSTRÖM, B. 2021 Spatio-temporal analysis of runaway electrons in a JET disruption with material injection. Master's thesis, Chalmers University of Technology.
- BREIZMAN, B.N., ALENYIKOV, P., HOLLMANN, E.M. & LEHNEN, M. 2019 Physics of runaway electrons in tokamaks. *Nucl. Fusion* **59**, 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).
- CHAPELLE, O. & LI, L. 2011 An empirical evaluation of Thompson sampling. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS' 11* (eds J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira & K.Q. Weinberger), pp. 2249–2257. Curran Associates Inc.
- CHILENSKI, M.A., GREENWALD, M.J., HUBBARD, A.E., HUGHES, J.W., LEE, J.P., MARZOUK, Y.M., RICE, J.E. & WHITE, A.E. 2017 Experimentally testing the dependence of momentum transport on second derivatives using Gaussian process regression. *Nucl. Fusion* **57**, 126013.
- CHILENSKI, M.A., GREENWALD, M., MARZOUK, Y.M., HOWARD, N.T., WHITE, A.E., RICE, J.E. & WALK, J.R. 2015 Improved profile fitting and quantification of uncertainty in experimental measurements of impurity transport coefficients using Gaussian process regression. *Nucl. Fusion* **55**, 023012.
- FRASSINETTI, L., BEURSKENS, M.N.A., SCANNELL, R., OSBORNE, T.H., FLANAGAN, J., KEMPENAARS, M., MASLOV, M., PASQUALOTTO, R., WALSH, M. & JET-EFDA CONTRIBUTORS 2012 Spatial resolution of the JET Thomson scattering system. *Rev. Sci. Instrum.* **83**, 013506.
- FRAZIER, P.I. 2018 A Tutorial on Bayesian Optimization. [arXiv:1807.02811](https://arxiv.org/abs/1807.02811).
- GPY SINCE 2012 GPY: a Gaussian process framework in python.

- GUTMANN, M.U. & CORANDER, J. 2016 Bayesian optimization for likelihood-free inference of simulator-based statistical models. *J. Machine Learning Res.* **17**, 1–47.
- HARVEY, R.W., CHAN, V.S., CHIU, S.C., EVANS, T.E., ROSENBLUTH, M.N. & WHYTE, D.G. 2000 Runaway electron production in DIII-D killer pellet experiments, calculated with the CQL3D/KPRAD model. *Phys. Plasmas* **7**, 4590.
- HESSLOW, L., EMBRÉUS, O., VALLHAGEN, O. & FÜLÖP, T. 2019 Influence of massive material injection on avalanche runaway generation during tokamak disruptions. *Nucl. Fusion* **59** (8), 084004.
- HO, A., CITRIN, J., AURIEMMA, F., BOURDELLE, C., CASSON, F.J., KIM, H.-T., MANAS, P., SZEPESI, G., WEISEN, H. & JET CONTRIBUTORS 2019 Application of Gaussian process regression to plasma turbulent transport model validation via integrated modelling. *Nucl. Fusion* **59**, 056007.
- HOPPE, M., EMBREUS, O. & FÜLÖP, T. 2021 DREAM: a fluid-kinetic framework for tokamak disruption runaway electron simulations. *Comput. Phys. Commun.* **268**, 108098.
- HOPPE, M., HESSLOW, L., EMBREUS, O., UNNERFELT, L., PAPP, G., PUSZTAI, I., FÜLÖP, T., LEXELL, O., LUNT, T., MACUSOVA, E., *et al.* 2021 Spatiotemporal analysis of the runaway distribution function from synchrotron images in an ASDEX upgrade disruption. *J. Plasma Phys.* **87** (1), 855870102.
- HUYSMANS, G. & CZARNY, O. 2007 MHD stability in X-point geometry: simulation of ELMs. *Nucl. Fusion* **47**, 659–666.
- JÄRVENPÄÄ, M., GUTMANN, M.U., PLESKA, A., VEHTARI, A. & MARTTINEN, P. 2019 Efficient acquisition rules for model-based approximate Bayesian computation. *Bayesian Anal.* **14** (2), 595–622.
- JÄRVENPÄÄ, M., GUTMANN, M.U., VEHTARI, A. & MARTTINEN, P. 2018 Gaussian process modelling in approximate Bayesian computation to estimate horizontal gene transfer in bacteria. *Ann. Appl. Stat.* **12** (4), 2228–2251.
- KANDASAMY, K., KRISHNAMURTHY, A., SCHNEIDER, J. & POCZOS, B. 2018 Parallelised Bayesian optimisation via Thompson sampling. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics* (ed. A. Storkey & F. Perez-Cruz), Proceedings of Machine Learning Research, vol. 84, pp. 133–142. PMLR.
- LI, D., SVENSSON, J., THOMSEN, H., MEDINA, F., WERNER, A. & WOLF, R. 2013 Bayesian soft X-ray tomography using non-stationary Gaussian Processes. *Rev. Sci. Instrum.* **84**, 083506.
- LINTUSAARI, J., GUTMANN, M.U., DUTTA, R., KASKI, S. & CORANDER, J. 2017 Fundamentals and recent developments in approximate Bayesian computation. *Syst. Biol.* **66** (1), e66–e82.
- LINTUSAARI, J., VUOLLEKOSKI, H., KANGASRÄÄSIÖ, A., SKYTÉN, K., JÄRVENPÄÄ, M., MARTTINEN, P., GUTMANN, M.U., VEHTARI, A., CORANDER, J. & KASKI, S. 2018 ELFI: engine for likelihood-free inference. *J. Machine Learning Res.* **19** (16), 1–7.
- MARIN, J.M., PUDLO, P., ROBERT, C.P. & RYDER, R.J. 2012 Approximate Bayesian computational methods. *Stat. Comput.* **22**, 1167–1180.
- MARJORAM, P., MOLITOR, J., PLAGNOL, V. & TAVARÉ, S. 2003 Markov chain Monte Carlo without likelihoods. *Proc. Natl Acad. Sci.* **100** (26), 15324–15328.
- OBERKAMPF, W.L. & TRUCANO, T.G. 2002 Verification and validation in computational fluid dynamics. *Prog. Aerosp. Sci.* **38** (3), 209–272.
- PALMA, A.D., MENDLER-DÜNNER, C., PARNELL, T.P., ANGHIEL, A. & POZIDIS, H. 2019 Sampling Acquisition Functions for Batch Bayesian Optimization. CoRR Vol. abs/1903.09434.
- PASQUALOTTO, R., NIELSEN, P., GOWERS, C., BEURSKENS, M., KEMPENAARS, M., CARLSTROM, T., JOHNSON, D. & JET-EFDA CONTRIBUTORS 2004 High resolution Thomson Scattering for Joint European Torus (JET). *Rev. Sci. Instrum.* **75**, 3891.
- RASMUSSEN, C.E. & WILLIAMS, C.K.I. 2006 *Gaussian Processes for Machine Learning*. MIT Press.
- REUX, C., PAZ-SOLDAN, C., ALEYNIKOV, P., BANDARU, V., FICKER, O., SILBURN, S., HOELZL, M., JACHMICH, S., EIDIETIS, N., LEHNEN, M., *et al.* 2021 Demonstration of safe termination of megaampere relativistic electron beams in tokamaks. *Phys. Rev. Lett.* **126**, 175001.
- RODRIGUEZ-FERNANDEZ, P., WHITE, A.E., CREELY, A.J., GREENWALD, M.J., HOWARD, N.T., SCIORTINO, F. & WRIGHT, J.C. 2018 VITALS: a surrogate-based optimization framework for the accelerated validation of plasma transport codes. *Fusion Sci. Technol.* **74** (1–2), 65–76.

- ROMERO, J.A. & SVENSSON, J. 2013 Optimization of out-vessel magnetic diagnostics for plasma boundary reconstruction in tokamaks. *Nucl. Fusion* **53**, 033009.
- ROSENBLUTH, M.N. & PUTVINSKI, S.V. 1997 Theory for avalanche of runaway electrons in tokamaks. *Nucl. Fusion* **37** (10), 1355–1362.
- SHAHRIARI, B., SWERSKY, K., WANG, Z., ADAMS, R.P. & FREITAS, N. 2016 Taking the human out of the loop: a review of Bayesian optimization. *Proc. IEEE* **104**, 148–175.
- SOVINEC, C.R., GLASSER, A.H., GIANAKON, T.A., BARNES, D.C., NEBEL, R.A., KRUGER, S.E., SCHNACK, D.D., PLIMPTON, S.J., TARDITI, A., CHU, M.S., *et al.* 2004 Nonlinear magnetohydrodynamics simulation using high-order finite elements. *J. Comput. Phys.* **195**, 355–386.
- SRINIVAS, N., KRAUSE, A., KAKADE, S. & SEEGER, M. 2010 Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, Haifa, Israel (eds. J. Fürnkranz & T. Joachims). Omnipress. <https://dl.acm.org/doi/proceedings/10.5555/3104322>
- STEIN, M.L. 1999 *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics.
- THOMPSON, W.R. 1933 On the likelihood that one unknown probability exceeds another in views of the evidence of two samples. *Biometrika* **25** (3–4), 285–294.
- VON NESSI, G.T. & HOLE, M.J. 2013 Using Bayesian analysis and Gaussian processes to infer electron temperature and density profiles on the Mega-Ampere Spherical Tokamak experiment. *Rev. Sci. Instrum.* **84**, 063505.
- VON NESSI, G.T., HOLE, M.J., SVENSSON, J. & APPEL, L. 2012 Evidence cross-validation and Bayesian inference of MAST plasma equilibria. *Phys. Plasmas* **19**, 012506.
- WIESEN, S., REITER, D., KOTOV, V., BAELMANS, M., DEKEYSER, W., KUKUSHKIN, A.S., LISGO, S.W., PITTS, R.A., ROZHANSKY, V., SAIBENE, G., *et al.* 2015 The new SOLPS-ITER code package. *J. Nucl. Mater.* **463**, 480–484.
- WU, X., KOZLOWSKI, T., MEIDANI, H. & SHIRVAN, K. 2018 Inverse uncertainty quantification using the modular Bayesian approach based on Gaussian process. Part 1: theory. *Nucl. Engng Des.* **335**, 339–355.