

PAPER • OPEN ACCESS

Machine learning techniques for sequential learning engineering design optimisation

To cite this article: L R Humphrey *et al* 2024 *Plasma Phys. Control. Fusion* **66** 025002

View the [article online](#) for updates and enhancements.

You may also like

- [MOJAVE. X. PARSEC-SCALE JET ORIENTATION VARIATIONS AND SUPERLUMINAL MOTION IN ACTIVE GALACTIC NUCLEI](#)
M. L. Lister, M. F. Aller, H. D. Aller et al.
- [A review on machine learning-guided design of energy materials](#)
Seongmin Kim, Jiaxin Xu, Wenjie Shang et al.
- [Roadmap on metamaterial theory, modelling and design](#)
Bryn Davies, Stefan Szyniszewski, Marcelo A Dias et al.

Machine learning techniques for sequential learning engineering design optimisation

L R Humphrey^{*} , A J Dubas, L C Fletcher and A Davis

United Kingdom Atomic Energy Authority, Culham Campus, Abingdon, Oxfordshire, OX14 3DB, United Kingdom

E-mail: luke.humphrey@ukaea.uk

Received 4 September 2023, revised 20 November 2023

Accepted for publication 4 December 2023

Published 21 December 2023



Abstract

When designing a fusion power plant, many first-of-a-kind components are required. This presents a large potential design space across as many dimensions as the component's parameters. In addition, multiphysics, multiscale, high-fidelity simulations are required to reliably capture a component's performance under given boundary conditions. Even with high performance computing (HPC) resources, it is not possible to fully explore a component's design space. Thus, effective interpolation between data points via machine learning (ML) techniques is essential. With sequential learning engineering optimisation, ML techniques inform the selection of simulation parameters which give the highest expected improvement for the model: balancing exploitation of the current best design with exploration of uncertain areas in the design space. In this paper, the application of an ML-driven design of experiment procedure for the sequential learning engineering design optimisation of a fusion component is shown. A parameterised divertor monoblock is taken as a typical example of a fusion component requiring HPC simulation to model. The component's geometry is then optimised using Bayesian optimisation, seeking the design which minimises the stress experienced by the component under operational conditions.

Keywords: machine learning, design optimisation, FEM

1. Introduction

To deliver commercially-viable fusion energy, a fusion power plant must not only be capable of generating more energy than is put into the fusion plasma, but the entire balance of plant must be sufficiently optimised to produce electrical energy at a low enough cost to compete with other sources. Such optimisation must include all of the systems between the plasma and

the electrical grid. Many of these systems are first-of-a-kind, presenting large regions of unexplored design space across any number of dimensions. Digital models of these systems are often multiphysics and multiscale, requiring high-fidelity simulations running on high performance computing (HPC) resources to model effectively [1, 2].

A straightforward approach to exploring a component's design space would be to perform a parameter sweep, evaluating how varying one parameter affects some metric of the component's performance. However, even if a full sweep is performed over each design parameter, the component would not be completely characterised unless the component's design parameters are independent in how they affect the performance metric. The number of evaluations required to fully explore a design space of n parameters with x number of

^{*} Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

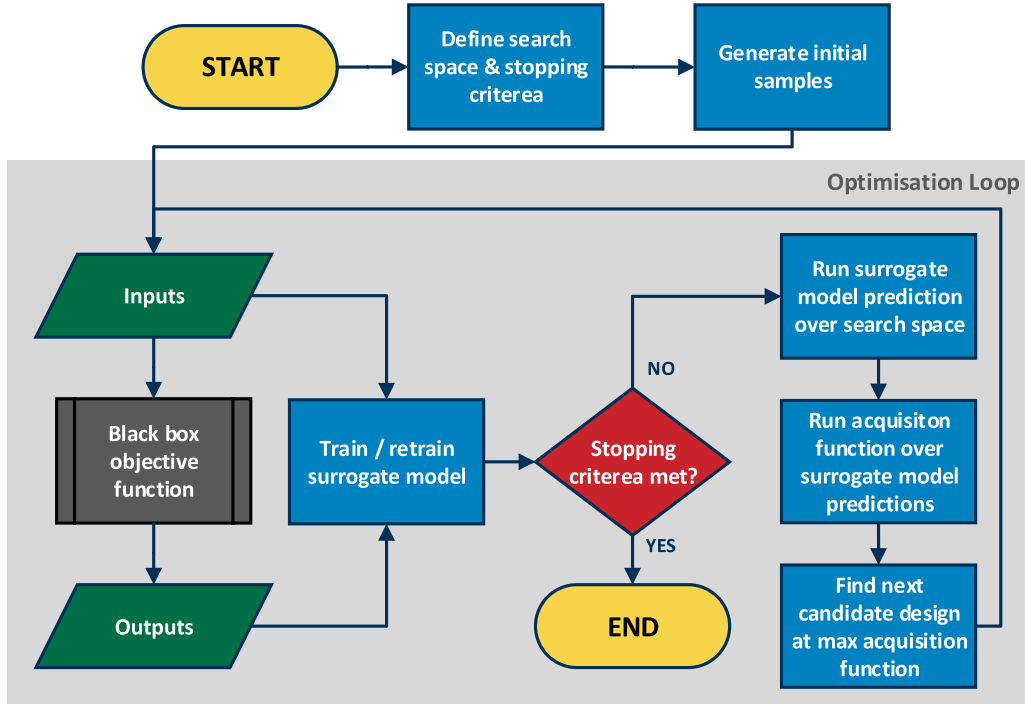


Figure 1. Flow chart describing the Bayesian optimisation workflow.

evaluations per dimension is $N = x^n$ via parameter sweep. Thus, for components with high dimensionality, fully exploring the design space is neither feasible nor even desirable. By contrast, in sequential optimisation, designs are evaluated one at a time with the performance of each design informing the selection of the next candidate design. One strategy for sequentially iterating on a component design is Bayesian optimisation, where the simulation is treated as a black-box function, and the aim is to optimise the design with as few design evaluations (i.e. expensive HPC simulations) as possible. When performing sequential optimisation, the computational engineer must:

- (i) Describe the design parametrically.
- (ii) Select the optimisation metric(s).
- (iii) Seek the values for the chosen parameters which optimise the chosen metric(s).

Breaking this down: the first step is to describe the design parametrically. The number of parameters and their bounds will vary between design problems. Appropriate parameters with sensible bounds should be used to avoid wasting resource exploring non-feasible regions of the design space. This step might also include defining input constraints: conditional relationships between inputs which must be met for a design candidate to be considered valid. This ensures invalid designs are not selected as candidates for simulation.

The second step is to define the optimisation metric(s). If there are multiple metrics, weights should be assigned and a weighted sum of metrics used to assess a design's quality. (Alternatively, especially if the weights are unknown, a

multi-objective algorithm may be chosen.) This step might also include defining output constraints: conditions on simulation outputs which must be met for the design to be considered valid. For example, ensuring a component remains within operational temperatures during the simulation.

The third step is the optimisation loop, as shown in figure 1. In Bayesian optimisation, the loop is first initialised by evaluating a set of samples which give a good amount of coverage throughout the design space (chosen according to a sample plan) and training a surrogate model on these data. The surrogate can interpolate to make predictions about the performance of unevaluated design points. Next, an acquisition function is run over the design space to determine the next candidate design. The acquisition function is highest for a design point whose predicted results maximises a trade-off between exploration and exploitation (i.e. exploring the design space versus improving on the best design so far). Finally, a simulation is run for the selected design point, and the surrogate model is updated with the results.

This optimisation loop is repeated until chosen stopping criteria are met, such as a maximum number of iterations or a sufficiently low value on the acquisition function. Note that training the surrogate model is not the goal of Bayesian optimisation but rather a means to an end for seeking an optimised design. Thus, it is perfectly acceptable for the surrogate to have high uncertainty within large regions of the design space at the end of the process, so long as the stopping criteria are satisfied.

Following its academic origins in the 1970s–1990s [3, 4], Bayesian optimisation has been gaining popularity in many academic and industrial fields. A 2016 review of Bayesian optimisation by Shahriari *et al* [5] describes its application in

sensor set selection [6], robotics [7], and drug discovery [8], among a diverse range of others. This trend has continued in recent years, seeing further industry application in fields such as chemical engineering [9] and railway engineering [10]. In a 2021 paper, Sobes *et al* [11] describe the machine learning (ML)-based design optimisation of a nuclear (fission) reactor core using HPC multiphysics simulation, demonstrating Bayesian optimisation of a continuously variable parameterised geometry. Though the focus was on nuclear fission, the methodology described is highly applicable to fusion design problems.

In a 2020 paper, Domptail *et al* [12] describe the parametric design optimisation of a monoblock divertor target with a thermal break interlayer (a component design developed by Fursdon *et al* in 2017 [13]). Domptail *et al* point out that such design optimisation methods are underutilised in fusion, suggesting that the same techniques could be applied to many other novel design challenges in the industry.

Motivated by this underutilisation of design optimisation in fusion, this study facilitates the first stages of development of a scalable open-source tool: the sequential learning engineering design optimiser (SLEDO), whose purpose is to enable engineers to run design optimisation on HPC resources without restrictive licenses.

The next section presents the methods used in this study, starting with the software tools developed (SLEDO), followed by details on the optimisation methods used, and then the modelling approach for the monoblock design being optimised. The results of three design optimisation runs of increasing complexity are presented and discussed in section 3, leading on to potential future work discussed in section 4. Finally, section 5 provides a summary and conclusion.

2. Methods

2.1. SLEDO

A new code for running design optimisation, SLEDO, was developed as part of the Aurora-multiphysics suite of open source fusion engineering software (available online [14]). SLEDO's primary dependencies are the Multiphysics Object-Oriented Simulation Environment (MOOSE) [15] and BoTorch [16] (a python package for Bayesian optimisation, implemented in SLEDO using Ax [17]).

The code's primary workflow is that shown in figure 1, with the MOOSE simulation acting as the black box objective function. The code is comprised primarily of two python classes: optimiser and simulation. The optimiser class handles the Bayesian optimisation loop, taking a parametric MOOSE input file and a search space, generating a modified input file for each design point, and passing this to the simulation class. The simulation class handles the simulation input–output (I/O): initiating the MOOSE simulation, reading the output, and passing the relevant metrics back to optimiser.

During the first few iterations, candidate designs are selected by an initial sample plan. At each subsequent iteration,

optimiser retrains the surrogate model and runs the acquisition function in order to determine the next candidate design. This loop continues until the stopping criteria are met. For this study, Sobol sampling (the default option in Ax) is used as the sample plan and each loop was allowed to run through 50 total design iterations (including the initial samples).

2.2. Surrogate model and acquisition function

In Bayesian optimisation, candidates for evaluation are selected by a trade-off between ‘exploration’ and ‘exploitation’, where the most useful candidates are those which provide information about uncertain regions of the search space (‘exploration’) while also being predicted to perform well based on the information gathered so far (‘exploitation’). This trade-off is quantified by an acquisition function, which is run over surrogate model predictions throughout the search space.

A surrogate model is a statistical model trained on a dataset of I/O pairs. A trained model is fit to the training data and can be queried to quickly provide a predicted output for an unknown point in the input space. Crucially, some models can also provide a measure of variance for the prediction based on proximity to training data, thus allowing the acquisition function to compute the exploration value.

In Bayesian optimisation, the process is as follows: first a surrogate model is trained on existing I/O data and used to provide uncertainty-quantified predictions throughout the search space, then an acquisition function takes these predictions and trades off between the exploration and exploitation value for each, finally, the maximum value of the acquisition function is found and the corresponding inputs are taken as the next candidates for evaluation.

For this study, the surrogate model used was a Gaussian processor (GP), and the acquisition function used was expected improvement (EI). This combination, often abbreviated as GPEI, is a common default option and a good starting point [16], however it may not be the best choice for every problem. Future work will compare the performance of different models for various problems.

EI is defined as:

$$EI(x) = \begin{cases} [\mu(x) - f(x^+) - \xi] \Phi(Z) + \sigma(x) \phi(Z) & \text{if } \sigma(x) > 0 \\ 0 & \text{if } \sigma(x) = 0 \end{cases} \quad (1)$$

$$Z = \begin{cases} [\mu(x) - f(x^+) - \xi] / \sigma(x) & \text{if } \sigma(x) > 0 \\ 0 & \text{if } \sigma(x) = 0 \end{cases} \quad (2)$$

where x is the design parameter, $\mu(x)$ and $\sigma(x)$ respectively are the mean and standard deviation of the surrogate model prediction at x , x^+ and $f(x^+)$ respectively are the parameters and value of the best design so far, Φ and ϕ respectively are the cumulative distribution function and probability density function of the normal distribution. ξ is an exploration hyperparameter, by default zero, which can be used to adjust the trade

off between the exploitation term $[\mu(x) - f(x^+) - \xi]\Phi(Z)$ and the exploitation term $\sigma(x)\phi(Z)$.

2.3. Divertor monoblock

The component to be optimised in this study is a divertor monoblock, a modular cooling component used as a divertor target in tokamaks. You *et al* [18] describes the monoblock design for the European DEMO; the geometry used in this study are based upon this design. Each monoblock is comprised of a cuboid block of tungsten armour surrounding a copper–chromium–zirconium pipe conveying coolant, with a softer copper interlayer between. These layers are bonded together without gaps; under operating temperatures each material expands differently due to differing coefficients of thermal expansion. This causes a stress at the interface of each layer; the role of the relatively soft interlayer is to reduce the chance of component failure by fracture. The incoming steady-state heat flux on the top surface of the monoblock is on the order of 10 MW m^{-2} , though off-normal conditions can increase the heat load significantly so an effective monoblock design should be able sustain heat fluxes up to 20 MW m^{-2} [18]. The optimisation objective for the component is to minimise the peak thermally-induced stress experienced under operational conditions. This component was chosen as it features a relatively straightforward parametric geometry to modify while also featuring the extreme thermal conditions that typify fusion thermomechanics and thermohydraulics modelling.

The geometry and meshing were done natively in MOOSE, making use of meshing tools from the reactor module by Shemon *et al* [19]. The model is remeshed for each design point. In order to facilitate geometric design space exploration, the number of mesh divisions for each geometric section is calculated by a formula which rounds to the nearest integer value in-line with a nominal mesh refinement of one division per millimetre. Note that this means that larger geometries contain more finite elements and thus take longer to simulate. Certain measurements (in particular the pipe and interlayer radii) are also given a minimum number of divisions to ensure a sufficient number of elements to model the temperature gradient over those regions even when the geometry becomes thin. Conservative estimates for mesh refinement minima were used for this study but, for a full scale design optimisation, a mesh refinement study should be used to determine suitable minima, as overly conservative values will lengthen time to solution with only a negligible increase to accuracy.

The simulation was carried out using the MOOSE application Proteus [20, 21], whose focus is on fusion digital twins involving fluid dynamics. The base monoblock models used in this study are available as examples in the Proteus repository [22]. Note that at this early stage, fluid dynamics have not been included in the model, but Proteus is used in anticipation of running a full-physics simulation of the monoblock in

future work, as this would require coupled fluid dynamics to model the flow of coolant.

2.3.1. Simplified model. A simplified model of the monoblock was developed first. This model is based upon an experimental model designed to investigate the thermal expansion of the interlayer when the block is heated to a constant temperature above the stress-free temperature; the simulation has been set up to replicate this experiment. Note that an operational monoblock experiences directional (not uniform) heating, though uniform heating will occur during manufacturing processes, which could lead to pre-stresses in the component before it is installed.

The model, as shown in figures 2(a) and 3, is comprised of a solid copper cylinder surrounded by a cuboid layer of tungsten armour. Neither the internal CuCrZr pipe nor the coolant channel are included in this model. Materials are assumed to be linear elastic isotropic using temperature-dependent material properties for thermal conductivity, coefficient of thermal expansion, density, elastic modulus, and specific heat capacity, applied via linear interpolation from available data [23, 24]. These material data are shown in tables A2 and A3. The mesh uses first order elements with a nominal mesh refinement of one division per millimetre. The boundary conditions are the stress-free temperature and the block temperature to which the block is uniformly heated. Zero displacement boundary conditions are present on the following nodes on the base: $\Delta x = 0$ for the centreline in x (on the y – z midplane), $\Delta z = 0$ for the centreline in z (on the y – z midplane), and $\Delta y = 0$ for all base nodes. The solution is steady state, outputting temperature, displacement, and stress fields. Table 1 shows the geometry and simulation parameters alongside their default values.

Note that while the terms block and armour are used somewhat interchangeably, the parameter ‘monoblock armour height’ refers to the dimensions of the additional armour on the top (plasma-facing) surface compared to the other sides, as shown in figure 3.

The simplicity of this design allowed for quick time-to-solution for a given design point, making it a useful model for testing and development of the optimisation tools without the need for HPC or acceleration. This model was optimised over a two-dimensional (2D) search space of the block thickness ($17 \text{ mm} < x_{\text{block}} < 34 \text{ mm}$) and armour height ($1 \text{ mm} < y_{\text{armour}} < 20 \text{ mm}$), assuming a fixed pipe/interlayer diameter.

2.3.2. Representative model. Following the simplified model, a more representative model of the monoblock was developed. This model, as shown in figures 2(b) and 4, builds upon the simple model by inclusion of the CuCrZr pipe and coolant channel (modelled as void) as well as heat-flux boundary conditions. Material properties are again assumed to be linear elastic isotropic, but now also include CuCrZr; material data are shown in table A1. The mesh for this model uses




Figure 2. The geometries, meshes, and materials of the base divertor monoblock models.




Figure 3. Geometric parameterisation of the simple monoblock model.

Table 1. Parameters of the simplified divertor monoblock model thermomechanical simulation, alongside their default values.

Parameter	Symbol/function	Default value
Pipe diameter	D_{pipe}	17 mm
Monoblock width	x_{block}	23 mm
Monoblock armour height	y_{armour}	8 mm
Monoblock height	$y_{\text{block}} = x_{\text{block}} + y_{\text{armour}}$	31 mm
Monoblock depth	z_{block}	12 mm
Stress free temperature	T_{SF}	20 °C
Steady state temperature	T	100 °C




Figure 4. Geometric parameterisation of the representative monoblock model.

second order elements; the nominal mesh refinement remains at one division per millimetre. The boundary conditions are the stress-free temperature, the incoming heat flux on the top surface of the armour, and the coolant water temperature. Zero displacement boundary conditions are present on several nodes at the base edges: $\Delta x = 0$ for the frontmost and backmost nodes of the centreline in x (on the y - z midplane), $\Delta z = 0$ for the leftmost and rightmost nodes of the centreline in z (on the x - y midplane), and $\Delta y = 0$ for all base nodes. The solution is steady state, outputting temperature, displacement, and stress fields. Table 2 shows the geometry and simulation parameters alongside their default values.

The outgoing convective heat flux on the internal pipe surface is a function of wall temperature, taken by linear interpolation of values approximated by the Sieder-Tate correlation [25] assuming a mean flow velocity of 16 m s^{-1} and pressure of 5 MPa as is considered for DEMO [26]. These values are shown in table A4. Note that these approximations neglect the inclusion of a swirl tape, which is likely to be used in DEMO monoblocks to increase mixing and improve heat transfer. In order to better simulate the heat transfer when a swirl tape is included, one could use a modified function for the heat transfer coefficient based on experimental data, such as in Li *et al* [27].

Note that since this model's geometry contains three distinct blocks (CuCrZr pipe, copper interlayer, and tungsten armour) it was necessary to redefine the parameterisation of the design to better facilitate exploration of the design space. Defining the block width as an independent absolute variable as in the simple model would lead to problematic areas of the search space where the diameter of the interlayer exceeded the width of the armour. One option to circumvent this issue is to set input constraints on the optimiser, only allowing valid designs to be evaluated. However it proved simpler to redefine

Table 2. Parameters of the divertor monoblock model thermomechanical simulation, alongside their default values.

Parameter	Symbol/function	Default value
Pipe thickness	t_{pipe}	1.5 mm
Pipe internal diameter	$D_{\text{int, pipe}}$	12 mm
Pipe external diameter	$D_{\text{ext, pipe}} = D_{\text{int, pipe}} + 2t_{\text{pipe}}$	15 mm
Interlayer thickness	$t_{\text{interlayer}}$	1 mm
Interlayer internal diameter	$D_{\text{int, interlayer}} = D_{\text{ext, pipe}}$	15 mm
Interlayer external diameter	$D_{\text{ext, interlayer}} + 2t_{\text{interlayer}}$	17 mm
Monoblock thickness	t_{block}	3 mm
Monoblock width	$x_{\text{block}} = D_{\text{ext, interlayer}} + 2t_{\text{block}}$	23 mm
Monoblock armour height	y_{armour}	8 mm
Monoblock height	$y_{\text{block}} = x_{\text{block}} + y_{\text{armour}}$	31 mm
Monoblock depth	z_{block}	12 mm
Stress free temperature	T_{SF}	20 °C
Coolant water temperature	T_w	150 °C
Heat flux in	Q	10 MW m ⁻²

the parameterisation in terms of the thickness of each layer, allowing the monoblock width x_{block} to become a derived value from these thicknesses. Thus, in this model, the monoblock thickness parameter t_{block} represents the apothem distance from the outer circumference of the interlayer to the outer edge of the monoblock (as shown in figure 4).




Figure 5. The steady-state stress fields of the base divertor monoblock models. Displacements are exaggerated by a scale factor of 10.

This model was first optimised over a four-dimensional (4D) search space of the following parameters: pipe thickness, interlayer thickness, block thickness, and armour height. In this first run, the pipe's internal diameter (and therefore the size of the flow channel) remains fixed. Secondly the same model was optimised over a six-dimensional (6D) search space of the same parameters plus the pipe internal diameter and block depth.

3. Results and discussion

Figure 5 shows the steady state stress fields for the original monoblock design (both simplified and representative models). These designs serve as a baseline point of comparison; it should be noted that, while these designs do exist within the bounded search spaces for each optimisation run, they are not included in initial samples. This is to allow the optimisation algorithm an unbiased view of the search space without explicit inclusion of any one design.

3.1. Simplified model, 2D search space

Figure 6 shows the optimised component design for the simple monoblock model, according to the 2D search space described in table 3. The block width has been increased significantly, landing on a value close to (but notably less than) the maximum bound, while the armour has been completely stripped away, having reduced to the minimum bound of 1.0 mm.

As can be seen by the optimisation trace in figure 7, the optimiser was able to reduce the max stress from 397 MPa to 336 MPa within 50 iterations. The initial sample plan comprised the first five trials and the best design found was on the 26th iteration, after which no better design was found.

Trial 26




Figure 6. Steady-state stress field for the simplified monoblock as optimised over a 2D search space. Displacements are exaggerated by a scale factor of 10.

Table 3. Comparison of original versus optimised values for the simplified monoblock model parameters optimised over a bounded 2D search space.

Parameter	Bounds	Original value	Optimised value
Block width	17.0–34.0 mm	23.0 mm	32.9 mm
Armour height	1.0–20.0 mm	8.0 mm	1.0 mm
Max stress		397 MPa	336 MPa

This result clearly does not represent the perfect monoblock design, but it does serve as an example of how ML techniques will only solve the design question as posed, and so




Figure 7. Optimisation trace for the simple monoblock model (2D input space). The thick blue line traces the performance of the best design found so far, while the dotted line shows the performance of each trial design. The orange dashed and green dash-dotted lines show the performance of the original and optimised designs respectively. The vertical dashed line indicates the model change from initial samples to GPEI optimisation loop. The y-axis has been cropped for increased visibility of the trace.

it is important to pose design questions effectively to generate a useful result. In this case, the armour is stripped away since it serves no purpose in a uniform heating case and only serves to add a directionality to the thermal expansion which causes stress to build up in a concentrated area. By removing the armour, the problem tends towards increased radial symmetry, and the stress becomes more evenly distributed. In the following representative case, which includes directional heat fluxes, the armour's importance to the thermal gradient becomes relevant again.

It is also of interest that the block width stopped short of the maximum bound. Using the extreme values $x_{\text{block}} = 34.0\text{mm}$, $y_{\text{armour}} = 1.0\text{mm}$, trial 10 achieved a max stress of 339MPa, slightly above the final optimum value. Given no obvious reason for nonlinearity in the design space, it is possible that this local optimum is caused by a discretisation artefact in the relatively coarse mesh. The low gradient of maximum stress values across the design space compared to the representative model supports this, as the more variant the output space, the less significant discretisation artefacts would be to the result.

3.2. Representative model, 4D search space

Figure 8 shows the optimised component design for the representative monoblock model, according to the 4D search space described in table 4. The design was reached by reducing the thicknesses of the pipe and interlayer while increasing the thickness of the tungsten block and maximising the height of additional armour on the top section.

As shown in figure 9, the max stress was reduced from 1090MPa to 668MPa within 50 iterations. The initial sample plan comprised the first eight trials and the best design found

Trial 16





Figure 8. Steady-state stress field for the representative monoblock as optimised over a 4D search space. Displacements are exaggerated by a scale factor of 10.

was on the 16th iteration. Compared with the simple model and 2D search space, more evaluations were required during the Bayesian optimisation loop following the initial sample plan but overall the optimal design was found in fewer iterations. This difference is likely due to the increased search space size in 4D, meaning the initial sample plan (though containing slightly more samples) leaves far more unexplored design space.

The contour plots in figure 10 reveal which parameters contributed most significantly to the max stress, with figure 10(a) demonstrating a steep linear gradient whose minimum lies

Table 4. Comparison of original versus optimised values for the representative monoblock model parameters optimised over a bounded 4D search space.

Parameter	Bounds	Original value	Optimised value
Pipe thickness	1.0–6.0 mm	1.5 mm	1.0 mm
Interlayer thickness	1.0–6.0 mm	1.0 mm	1.0 mm
Block thickness	1.0–6.0 mm	3.0 mm	5.38 mm
Armour height	1.0–16.0 mm	8.0 mm	16.0 mm
Max stress		1090 MPa	668 MPa

**Figure 9.** Optimisation trace (as figure 7) for the representative monoblock model optimised over a 4D search space.

in the lower-left corner. Comparing this to the relatively low variation of stress over the axes for block thickness and armour height in figures 10(b)–(e), it appears that the pipe and interlayer dimensions are the more significant parameters. Interestingly, $t_{\text{block}} = 5.38 \text{ mm}$ was the only parameter not to hit an upper or lower bound, hinting at a potential non-linearity (as can be seen in the saddle contour in figure 10(f)).

3.3. Representative model, 6D search space

Figure 11 shows the optimised component design for the representative monoblock model, according to the 6D search space described in table 5. The internal pipe diameter was maximised to the upper bound of 24.0 mm, while the block depth was reduced slightly to 10.5 mm. The block thickness is lower here than the previous design, but still higher than the original design. This design otherwise closely matches the 4D optimisation: both have minimised pipe and interlayer thickness and maximised armour height.

Note that allowing the internal pipe diameter to vary while assuming a consistent heat transfer coefficient is somewhat unphysical since the heat transfer coefficient is dependent on the flow conditions [25]. In theory, the flow velocity could be altered to control the heat transfer coefficient, though this will impact the required pumping power. Such an investigation

would best be carried out as a multi-objective optimisation with the additional objective of minimised pumping power.

As shown in figure 12, the max stress was reduced from 1090 MPa to 620 MPa within 50 iterations. The initial sample plan comprised the first 12 trials and the best design found was on the 31st iteration. Compared to the previous 4D optimisation, this model achieved a slightly lower maximum stress though significantly more iterations were required to find this optimum. This indicates that for components with a high number of design parameters, it may be beneficial to run a parameter sensitivity study and construct a reduced dimensionality search space which can still efficiently determine the optimal design.

Notably, the saddle contour from figure 10(f) is not present in the equivalent figure C1(m), indicating that the space's true shape likely differs from either GP prediction. This serves as a reminder that the surrogate models generated during Bayesian optimisation are not generally accurate throughout the design space due to the clustering of training data (and associated model accuracy) towards regions selected by the acquisition function. As mentioned in section 1, this is not an issue for Bayesian optimisation, but the uneven distribution of model error should be kept in mind when using the surrogates generated by the optimisation loop to make more general inferences about the design space.




Figure 10. Contour plots for the representative monoblock model displaying the GP mean posterior prediction for maximum stress over the 4D search space. Parameters not included in a given plot are set to the middle of their range.




Figure 11. Steady-state stress field for the representative monoblock as optimised over a 6D search space. Displacements are exaggerated by a scale factor of 10.

Table 5. Comparison of original versus optimised values for the representative monoblock model parameters optimised over a bounded 6D search space.

Parameter	Bounds	Original value	Optimised value
Pipe int. diameter	6.0–24.0 mm	12.0 mm	24.0 mm
Pipe thickness	1.0–6.0 mm	1.5 mm	1.0 mm
Interlayer thickness	1.0–6.0 mm	1.0 mm	1.0 mm
Block thickness	1.0–6.0 mm	3.0 mm	4.35 mm
Armour height	1.0–16.0 mm	8.0 mm	16.0 mm
Block depth	10.0–14.0 mm	12.0 mm	10.5 mm
Max stress		1090 MPa	620 MPa




Figure 12. Optimisation trace (as figure 7) for the representative monoblock model optimised over a 6D search space.

4. Future work

The parameter spaces in this study were on the whole linear, containing no instances of local minima which might trap a simplistic gradient descender. For this reason, the problem as posed fails to present a sufficient stress test for the optimisation tools to evaluate their performance in complex fusion engineering. Indeed, the simplicity of the models also led to optimised designs whose real world performance would be questionable. For the monoblock design, the complexity of the modelling could be increased by inclusion of various features such as:

- Variable convective heat transfer coefficient as a function of the internal pipe geometry, (potentially incorporating the effects of swirl tape mixing as mentioned in section 2.3.2).
- Plastic deformation of the copper interlayer, which acts to reduce the stress in the CuCrZr pipe and tungsten armour [28].
- Multiple failure criteria such as melting, creep, fatigue, fracture, plastic collapse, and/or ductile failure under different stress triaxiality [27–29]
- Irradiation effects on material properties, such as neutron damage and/or embrittlement [30].
- Residual stresses from manufacturing processes [31].
- Transient evaluation of the monoblock's resistance to thermal shock loading from a plasma disruption [32].
- Transient evaluation of the effectiveness of plasma strike point sweeping [33].
- Erosion of the armour layer due to plasma surface interaction [34].
- Changes in tungsten material properties due to heating above the recrystallisation temperature [35].

The complexity of the optimisation problem could also be increased by using a multi-objective cost function including the following objectives:

- Minimise the pumping power required.
- Minimise the cost of materials used.

Future work will iterate on the problem's complexity to introduce non-linearity in the search space, and compare the performance of different ML techniques at optimising the design in as few design evaluations as possible. Future work will also tackle components with more complex geometries, such as tokamak blankets, to evaluate the performance of different ML techniques in higher dimensional search spaces.

From a software development perspective, future versions of SLEDO will implement dimensionality reduction features and visualisation tools, as well as the ability to choose different optimisation techniques besides GPEI Bayesian optimisation. In order to lower the barrier to entry for users unfamiliar with MOOSE, future development will also automatically generate MOOSE input files for a given optimisation problem without requiring a base input file to be supplied.

5. Conclusion

In order to test and develop a scalable optimisation methodology for fusion problems, the European DEMO divertor monoblock design was taken as a representative fusion component. The design was optimised using a new code, SLEDO, to minimise the maximum stress experienced in the block due to thermal expansion under operational conditions. The component was modelled at two levels of fidelity: a simplified model heated to a uniform temperature and a representative model including operational heat-flux boundary conditions; in each case the solution was steady-state.

The simplified model was successfully optimised over a 2D search space. The max stress was reduced from 397 MPa to 336 MPa in 26 iterations, though the resulting design would not be viable in practice (due to the simplicity of the model). In particular the uniform heating does not reflect the conditions experienced by an *in-situ* monoblock.

The representative model was optimised first over a 4D search space and then over a 6D search space. The resulting designs reduced the max stress from 1090 MPa to 668 MPa and 620 MPa respectively, over 16 and 31 iterations respectively. Increasing the number of dimensions thus showed diminishing gains for significantly higher computational cost, indicating that the problem could benefit from an input sensitivity study and dimensionality reduction.

ML techniques offer the computational engineer a new set of tools when it comes to component design, but these tools must be used correctly to achieve a viable optimised design. Future work will continue the development of SLEDO and explore the use of different ML techniques for more complex models and higher-dimensional components.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.14468/77qp-vn81>.

Acknowledgments

The lead author would like to thank their colleagues in the Advanced Engineering Simulation Group at UKAEA. In particular, many thanks to Alex Blair, Lead HPC Computational Engineer, for their assistance in writing the parametric MOOSE input files used in this study. Finally, a special thanks to ParaView Expert Engineers Mathieu Westphal and Julien Fausty at Kitware who helped improve the monoblock visualisations generated by SLEDO and used in this paper.

Appendix A. Material properties

The values for temperature-dependent material properties used for the thermomechanical models in this study are shown in tables A1–A3. Additionally, Poisson's ratio was set to 0.33

Table A1. Temperature-dependent material properties for the CuCrZr pipe, where T is the temperature, K is the thermal conductivity, α is the coefficient of thermal expansion, ρ is the density, λ is the elastic modulus, and c is the specific heat capacity.

T ($^{\circ}\text{C}$)	K ($\text{W}(\text{m}^{\circ}\text{C})^{-1}$)	α ($\mu\text{m}(\text{m}^{\circ}\text{C})^{-1}$)	ρ (kg m^{-3})	λ (GPa)	c ($\text{J}(\text{kg}^{\circ}\text{C})^{-1}$)
20	318	16.7	8900	128	390
50	324	17	8886	127	393
100	333	17.3	8863	127	398
150	339	17.5	8840	125	402
200	343	17.7	8816	123	407
250	345	17.8	8791	121	412
300	346	18	8797	118	417
350	347	18	8742	116	422
400	347	18.1	8716	113	427
450	346	18.2	8691	110	432
500	346	18.4	8665	106	437
550	—	18.5	—	100	442
600	—	18.6	—	95	447
650	—	—	—	90	452
700	—	—	—	86	458

Table A2. Temperature-dependent material properties for the copper interlayer, where T is the temperature, K is the thermal conductivity, α is the coefficient of thermal expansion, ρ is the density, λ is the elastic modulus, and c is the specific heat capacity.

T ($^{\circ}\text{C}$)	K ($\text{W}(\text{m}^{\circ}\text{C})^{-1}$)	α ($\mu\text{m}(\text{m}^{\circ}\text{C})^{-1}$)	ρ (kg m^{-3})	λ (GPa)	c ($\text{J}(\text{kg}^{\circ}\text{C})^{-1}$)
20	401	16.7	8940	117	388
50	398	17.0	8926	116	390
100	395	17.2	8903	114	394
150	391	17.5	8879	112	398
200	388	17.7	8854	110	401
250	384	17.8	8829	108	406
300	381	18.0	8802	105	410
350	378	18.1	8774	102	415
400	374	18.2	8744	98	419
450	371	18.4	8713	—	424
500	367	18.5	8681	—	430
550	364	18.7	8647	—	435
600	360	18.8	8612	—	441
650	357	19.0	8575	—	447
700	354	19.1	8536	—	453
750	350	19.3	8495	—	459
800	347	19.6	8453	—	466
850	344	19.8	8409	—	472
900	340	20.1	8363	—	479
950	337	—	—	—	487
1000	334	—	—	—	494

Table A3. Temperature-dependent material properties for the tungsten armour, where T is the temperature, K is the thermal conductivity, α is the coefficient of thermal expansion, ρ is the density, λ is the elastic modulus, and c is the specific heat capacity.

T ($^{\circ}\text{C}$)	K ($\text{W (m }^{\circ}\text{C)}^{-1}$)	α ($\mu\text{m (m }^{\circ}\text{C)}^{-1}$)	ρ (kg m^{-3})	λ (GPa)	c ($\text{J (kg }^{\circ}\text{C)}^{-1}$)
20	173	4.50	19 300	398	129
50	170	—	19 290	398	130
100	165	4.50	19 280	397	132
150	160	—	19 270	397	133
200	156	4.53	19 250	396	135
250	151	—	19 240	396	136
300	147	4.58	19 230	395	138
350	143	—	19 220	394	139
400	140	4.63	19 200	393	141
450	136	—	19 190	391	142
500	133	4.68	19 180	390	144
550	130	—	19 170	388	145
600	127	4.72	19 150	387	147
650	125	—	19 140	385	148
700	122	4.76	19 130	383	150
750	120	—	19 110	381	151
800	118	4.81	19 100	379	152
850	116	—	19 080	376	154
900	114	4.85	19 070	374	155
950	112	—	19 060	371	156
1000	110	4.89	19 040	368	158
1100	108	—	19 010	362	160
1200	105	4.98	18 990	356	163
1400	—	5.08	—	—	—
1600	—	5.18	—	—	—
1800	—	5.30	—	—	—
2000	—	5.43	—	—	—
2200	—	5.57	—	—	—
2400	—	5.74	—	—	—
2600	—	5.93	—	—	—
2800	—	6.15	—	—	—
3000	—	6.40	—	—	—
3200	—	6.67	—	—	—

Table A4. Temperature-dependent heat transfer coefficient for the coolant water, where T is the temperature and h is the heat transfer coefficient. Values approximated by the Sieder–Tate correlation [25] assuming a mean flow velocity of 16 m s^{-1} and pressure of 5 MPa as is considered for DEMO [26].

T ($^{\circ}\text{C}$)	h ($\text{W (m}^2\text{ }^{\circ}\text{C)}^{-1}$)
1	4.0
100	109 100.0
150	115 900.0
200	121 010.0
250	128 800.0
295	208 200.0

for the CuCrZr pipe and copper interlayer, and 0.29 for the tungsten armour. All material data are taken from the ITER structural design criteria for in-vessel components [23, 24].

Table A4 shows the temperature-dependent values for heat transfer coefficient used to model the outgoing heat flux on the internal pipe surface in the representative model.

Appendix B. Equations used

The MOOSE tensor mechanics module was used to solve for displacement and stress due to thermal expansion. Materials were modelled as linear elastic isotropic, with elastic modulus supplied as a function of temperature (as shown in appendix A).

The `ComputeInstantaneousThermalExpansionFunctionEigenstrain` material block was used to calculate the eigenstrain tensor for a given timestep as:

$$\epsilon_{t+1}^{\text{th}} = \epsilon_t^{\text{th}} + \frac{1}{2} (\alpha_{(T_t)} + \alpha_{(T_{t+1})}) (T_{t+1} - T_t) I \quad (\text{B.1})$$

where ϵ^{th} is the eigenstrain tensor, $t+1$ and t denote the current and previous timesteps respectively, T is the temperature, $\alpha_{(T)}$ is the instantaneous thermal expansion at a given temperature, and I is the identity matrix. Note that for the first timestep, the stress-free temperature is used for the previous timestep.

The `ComputeFiniteStrainElasticStress` material block was used to calculate the elastic stress tensor as:

$$\sigma_{ij} = C_{ijkl} \Delta \epsilon_{kl} \quad (\text{B.2})$$

where σ_{ij} are the components of the elastic stress tensor, C_{ijkl} are the components of the elasticity tensor, and $\Delta \epsilon_{kl}$ are the components of the strain increment tensor.

The steady state solution was reached by solving for heat conduction using the `HeatConduction` kernel, which uses the following equation:

$$0 = \nabla K(t, \vec{x}) \nabla T \quad \text{for } \vec{x} \in \Omega$$

where T is the temperature, t is time, K is the thermal conductivity, $\vec{x}$ is the vector of spatial coordinates, and Ω is the solution domain.

Appendix C. Additional contour plots

The 15 contour plots for the monoblock optimisation run over a six-dimensional search space are shown in figure [C1](#).





Figure C1. Contour plots for the representative monoblock model displaying the GP mean posterior prediction for maximum stress over the 6D search space. Parameters not included in a given plot are set to the middle of their range.

**Figure C1.** (Continued.)




Figure C1. (Continued.)

ORCID iD

L R Humphrey  <https://orcid.org/0000-0001-5066-6036>

References

- [1] Davis A, Dubas A and Otin R 2020 *EPJ Web Conf.* **245** 09001
- [2] Brooks H and Davis A 2022 *Plasma Phys. Control. Fusion* **65** 024002
- [3] Moćkus J 1975 On Bayesian Methods for Seeking the Extremum *Optimization Techniques IFIP Technical Conf. (Lecture Notes in Computer Science) (Novosibirsk, 1–7 July 1974)* ed G I Marchuk (Springer) pp 400–4
- [4] Jones D R, Schonlau M and Welch W J 1998 *J. Global Opt.* **13** 455–92
- [5] Shahriari B, Swersky K, Wang Z, Adams R P and De Freitas N 2016 *Proc. IEEE* **104** 148–75
- [6] Garnett R, Osborne M A and Roberts S J 2010 Bayesian optimization for sensor set selection *Proc. 9th ACM/IEEE Int. Conf. on Information Processing in Sensor Networks (IPSN '10)* (Association for Computing Machinery) pp 209–19
- [7] Lizotte D J, Wang T, Bowling M H and Schuurmans D 2007 Automatic gait optimization with Gaussian process regression *IJCAI* vol 7 pp 944–9
- [8] Negoescu D M, Frazier P I and Powell W B 2011 *Inform. J. Comput.* **23** 346–63
- [9] Hernández-Lobato J M, Requeima J, Pyzer-Knapp E O and Aspuru-Guzik A 2017 Parallel and distributed thompson sampling for large-scale accelerated exploration of chemical *Space Proc. 34th Int. Conf. on Machine Learning (PMLR)* pp 1470–9
- [10] Singh A, Gu Z, Hou X, Liu Y and Hughes D J 2022 *Compos. Struct.* **282** 115107
- [11] Sobes V, Hiscox B, Popov E, Archibald R, Hauck C, Betzler B and Terrani K 2021 *Sci. Rep.* **11** 19646
- [12] Domptail F, Barrett T R, Fursdon M, Lukenskas A and You J H 2020 *Fusion Eng. Des.* **154** 111497
- [13] Fursdon M et al 2017 *Phys. Scr.* **2017** 014042
- [14] Humphrey L et al 2023 (available at: <https://github.com/aurora-multiphysics/sledo>) (Accessed 4 September 2023)
- [15] Lindsay A D et al 2022 *SoftwareX* **20** 101202
- [16] Balandat M, Karrer B, Jiang D R, Daulton S, Letham B, Wilson A G and Bakshy E 2019 arXiv:1910.06403

- [17] Meta Open Source 2023 (available at: <https://ax.dev/>) (Accessed 4 September 2023)
- [18] You J H *et al* 2022 *Fusion Eng. Des.* **175** 113010
- [19] Shemon E, Miao Y, Kumar S, Mo K, Jung Y S, Oaks A, Richards S, Giudicelli G, Harbour L and Stogner R 2023 *Nucl. Sci. Eng.* **197** 1–25
- [20] Dubas A *et al* 2023 (available at: <https://github.com/aurora-multiphysics/proteus/>) (Accessed 4 September 2023)
- [21] Peterson J W, Lindsay A D and Kong F 2018 *Adv. Eng. Softw.* **119** 68–92
- [22] Dubas A *et al* 2023 (available at: <https://github.com/aurora-multiphysics/proteus/tree/main/examples/monoblock>) (Accessed 4 September 2023)
- [23] ITER Organization 2012 ITER structural design criteria for in-vessel components (SDC-IC)
- [24] ITER Organization 2013 ITER SDC-IC appendix A, materials design limit data
- [25] Sieder E N and Tate G E 1936 *Ind. Eng. Chem.* **28** 1429–35
- [26] You J H, Mazzone G, Visca E, Bachmann C, Autissier E, Barrett T, Cocilovo V, Crescenzi F, Domalapally P K and Dongiovanni D 2016 *Fusion Eng. Des.* **109** 1598–603
- [27] Li M, Werner E and You J H 2014 *Fusion Eng. Des.* **89** 2716–25
- [28] You J H, Li M and Zhang K 2021 *Fusion Eng. Des.* **164** 112203
- [29] Panayotis S, Hirai T, Barabash V, Amzallag C, Escourbiac F, Durocher A, Komarov V, Martinez J and Merola M 2017 *Fusion Eng. Des.* **125** 256–62
- [30] Gilbert M R, Dudarev S L, Zheng S, Packer L W and Sublet J C 2012 *Nucl. Fusion* **52** 083019
- [31] Evans L *et al* 2018 *Fusion Eng. Des.* **134** 97–108
- [32] Wang S, Li J, Wang Y, Zhang X, Wang R, Wang Y and Cao J 2020 *Sci. Rep.* **10** 1359
- [33] Maviglia F *et al* 2016 *Fusion Eng. Des.* **109–111** 1067–71
- [34] Ueda Y, Coenen J, De Temmerman G, Doerner R, Linke J, Philipps V and Tsitrone E 2014 *Fusion Eng. Des.* **89** 901–6
- [35] Jin Y Z, Liu X, Lian Y Y and Song J P 2022 *Tungsten* **4** 194–202