frontiers | Frontiers in Nuclear Engineering

TYPE Original Research
PUBLISHED 02 February 2026
DOI 10.3389/fnuen.2025.1694684

OPEN ACCESS

EDITED BY
Stefano Segantin,
Massachusetts Institute of Technology,
United States

REVIEWED BY
Davide Pettinari,
Polytechnic University of Turin, Italy
Marco De Pietri,
Massachusetts Institute of Technology,
United States

*CORRESPONDENCE
Luke Humphrey,
luke.humphrey@ukaea.uk

RECEIVED 28 August 2025
REVISED 03 November 2025
ACCEPTED 21 November 2025
PUBLISHED 02 February 2026

CITATION
Humphrey L, Brooks H, Mungale S, Davis A and Foster D (2026) Accelerating the engineering design of breeder blankets with parametric optimisation and sequential learning. *Front. Nucl. Eng.* 4:1694684. doi: 10.3389/fnuen.2025.1694684

Accelerating the engineering design of breeder blankets with parametric optimisation and sequential learning

Luke Humphrey*, Helen Brooks, Siddharth Mungale, Andrew Davis and David Foster

United Kingdom Atomic Energy Authority, Culham Campus, Abingdon, Oxfordshire, United Kingdom

The competing requirements of fusion breeder blankets and the high dimensionality of their design space necessitate a systematic treatment to map the variations in performance against given objective metrics and to understand the operational envelope. In this endeavour, a digital engineering pipeline for design evaluation and optimisation has been developed. The tools involved are Hypnos for parametric breeder blanket geometry instantiation, OpenMC for neutronics analysis, MOOSE for thermal hydraulics analysis, and SLEDO for design space sampling, sensitivity analysis, and optimisation. An optimisation of the baseline design for a solid ceramic breeder mock-up that is relevant to the Lithium Breeding Tritium Innovation (LIBRTI) program is performed. Two optimisation studies are performed, the first involving only neutronics, while the second includes the impact of thermal hydraulics. The figures of merit are taken to be the tritium breeding ratio (TBR) and the pressure drop of the outer coolant (combined in a weighted sum for the second analysis). In the first study, for the same acquisition function (taken to be expected improvement), two different values are selected for the hyperparameter that controls the trade-off between exploration and exploitation. In the second study, with the inclusion of thermal hydraulics, a larger parameter space was explored to assess the performance of the method in a higher dimensionality setting. In both cases, the selected figures of merit were improved over the baseline design. Finally, we discuss extensions of the procedure to include a more thorough multi-physics analysis and a more sophisticated treatment of multiple objectives.

KEYWORDS

breeder blankets, neutronics, optimisation, sequential learning, OpenMC, multi-physics object oriented simulation environment, sequential learning engineering design optimiser, hypnos

1 Introduction

Breeder blankets encapsulate the challenges faced more broadly by in-vessel components within fusion pilot plants: high heat fluxes, ionising radiation, and electromagnetic and gravitational loading. Blanket systems therefore constitute a worthwhile focal point for the application of novel design methodologies. Competition arises due to these components having multiple requirements, whose fulfilment may result in diverse and potentially opposing changes, ultimately implying trade-offs in the spatial allocation of functional materials.

The challenges and functional requirements for breeder blankets have long been recognised (Abdou et al., 1985). Some of the key requirements may be stipulated as follows.

- The breeder blanket must shield the rest of the device from damaging irradiation. The device must have structural integrity, and any structural materials must be resilient to high neutron fluxes and temperatures.
- The breeder blanket must be actively cooled to ensure that safe operational temperatures are maintained and, simultaneously, that deposited heat is transported for the ultimate purpose of power generation.
- The breeder blanket must breed sufficient tritium. For commercial viability and sustainability of fusion as an energy source, the tritium fuel cycle must be self-sustaining: for every neutron emitted, there must be one or more subsequent nuclear reactions to produce tritium.

These requirements each entail conflicting design choices from which trade-offs naturally emerge. Designs that favour shielding and structural considerations would increase the volume of non-breeder, non-coolant materials, which nonetheless require cooling and which absorb neutrons without contributing to tritium production. Alternatively, designs that favour effective thermal management would reduce the volume of material in need of cooling (including both structural and breeding material); these designs would also seek to reduce required pumping power by having fewer channels across which to distribute the coolant. Finally, designs favouring high tritium breeding capability would minimise neutron loss by reducing the volume of structural and coolant materials in favour of the breeder and multiplier. In addition, these designs would distribute materials as homogeneously as possible, necessitating an increase in the number of coolant channels—a preference which is discussed further in Section 3.

Design challenges of this nature can be posed mathematically, commonly referred to as "multidisciplinary design analysis optimisation" (MDAO). Suppose that our design point is defined by a finite number of configurable parameters $\mathbf{x} = \{x_1, x_2, \ldots, x_n\}$. The ranges of those parameters define some design space $\mathbf{\Omega}$. Figure(s) of merit, such as the tritium breeding ratio (TBR), can be evaluated as a function of the design point, $f(\mathbf{x})$. Similarly, any constraints may be expressed as $g(\mathbf{x}) > 0$, which is equivalent to restriction to one or more subdomains $\mathbf{\Omega}'$ on $\mathbf{\Omega}$. The optimisation problem then is simply to maximise $f(\mathbf{x})$ on $\mathbf{\Omega}'$.

Given that the dimensionality of such a problem is typically large, it is hard to intuitively comprehend the subspace within which optimal solutions exist. Indeed, improperly posed, there is no guarantee of the existence of any solution at all; such a scenario would arise if constraints are mutually exclusive and $\mathbf{\Omega}'$ becomes a null space. Alternatively, optima may be unstable to small perturbations in the design space, a scenario which translates to rapid degradation in performance due to (for example) natural variation in the manufacturing process or small changes in the specification of requirements.

Both scenarios outlined above are best avoided; they would likely necessitate a complete design overhaul, at significant expense. Instead, it is more efficient to systematically explore the design space, as well as evaluating sensitivities to parameters away from any optima. Here, the dimensionality of the problem again poses a challenge; a brute force exploration may become computationally prohibitive. However, if a surrogate model of the figure of merit can be constructed on the fly (and progressively refined with more model evaluations), an understanding of the design space and any associated trade-offs in parameters could be attained more rapidly.

A schematic of such a workflow is shown in Figure 1. Given the problem's definition, candidate designs are generated in an initial sample, and a figure of merit is evaluated using physical models. The initial pool of results will inform a more targeted sampling in subsequent cycles of this *outer loop*. After a sufficient number of iterations, the optimal configuration(s) may be identified and their robustness assessed. This process, however, is predicated on the existence of a so-called *design-point evaluator*. In other words, it must be possible to (i) construct a design from some parameterisation $\mathbf{x}$, (ii) stipulate the models that govern the behaviour of physical variables, and (iii) simulate these numerically, computing in turn the figure of merit from those variables. Only with such an implementation in place is it possible to effectively leverage optimisation algorithms in order to accelerate the engineering design process.

As part of the digital workstream of the UK Atomic Energy Authority's Lithium Breeding Tritium Innovation program (LIBRTI) (Gilbert et al., 2025), an implementation of such a digital engineering workflow has been developed for the specific case of parametric geometry optimisation. While eventually the intention is to improve the designs of breeder blankets, the short-term focus is to develop breeder mock-up experiments. These can be tested at the planned LIBRTI facility, which will host a 14 MeV neutron source and will enable the qualification of both tritium breeding technology and predictive modelling capability.

The outline of this paper is as follows: the methodology is presented in Section 2 where in Section 2.1, we describe the full tool-chain employed to implement the workflow of Figure 1. First, a novel software tool, Hypnos (Mungale, 2024), for parametric blanket geometry preparation is presented in Section 2.1.1, and the reference breeder mock-up design is described (this tool provides the functionality for the "instantiate model" step). Next, in Section 2.1.2 and Section 2.1.3 the analysis set-up for neutronics with OpenMC (Romano and Forget, 2013; Romano et al., 2015) and thermal hydraulics with MOOSE (Gaston et al., 2009; Permann et al., 2020; Giudicelli et al., 2024) is described. Although the complex nature of the breeder blanket should necessitate a tightly coupled treatment of physics, in this initial demonstration of the optimisation set-up, these analyses are treated separately and are used to define two separate figures of merit. Finally, for the provision of outer loop functions—sampling and statistical aggregation for the purpose of optimisation—we review SLEDO software (Humphrey et al., 2024) in Section 2.1.4.

In Section 2.2, a number of different optimisation campaigns are defined. The first considers neutronics in isolation and includes an investigation into the impact of varying the search strategy through a hyperparameter. In the second campaign, both neutronics and thermal hydraulics analyses are performed, permitting an investigation into the trade-off between competing objectives.

Having established the methodology, both the baseline results and the results of the optimisation campaigns are presented in

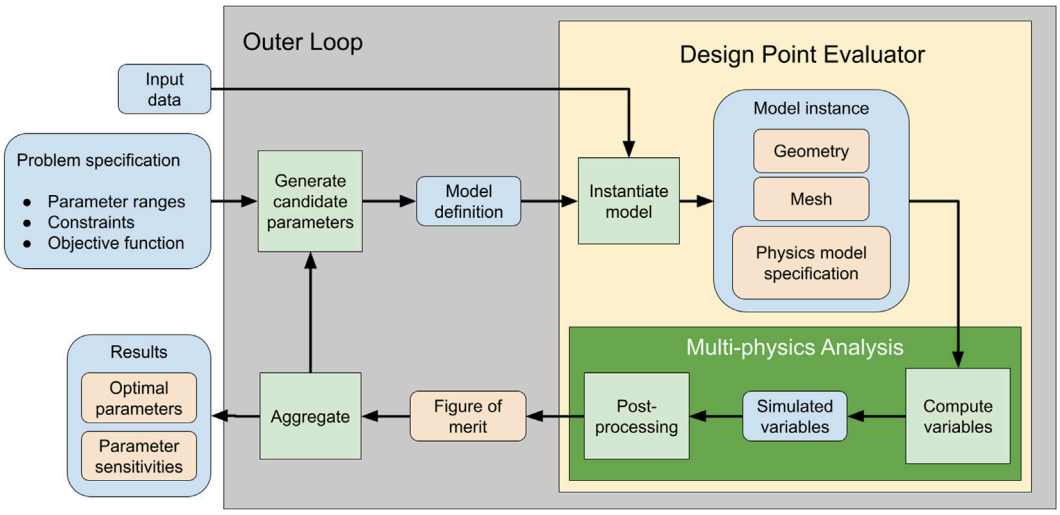

FIGURE 1
Generalised digital engineering workflow for optimisation and uncertainty quantification.

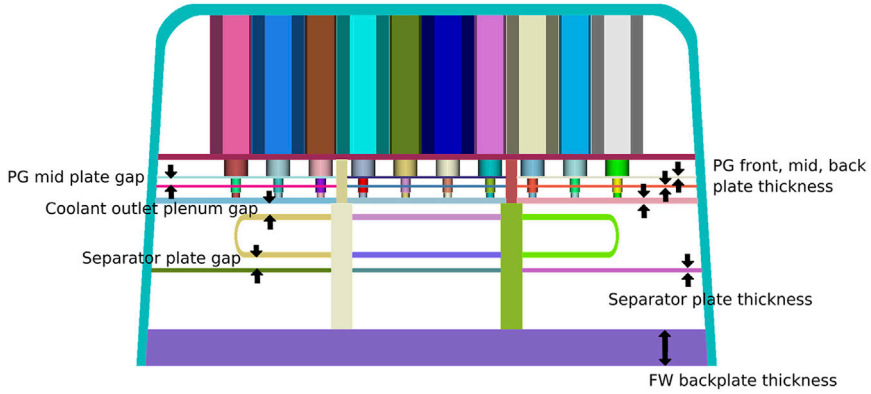

FIGURE 2
Illustrative view of HCPB blanket parametric geometry distributed with Hypnos. Several parameters are labelled, with the full set available described in the online documentation.

Section 3 and discussed in Section 4. Finally, the outlook and conclusions are presented in Section 5 and Section 6, respectively.

2 Methodology

2.1 Software tools and simulation set-up

2.1.1 Parametric breeder geometry instantiation with Hypnos

Hypnos (Mungale, 2024) is a new open-source parametric geometry engine used to model fusion-relevant components. Written in Python, it uses Coreform Cubit's (Coreform LLC, 2024) API to instantiate CAD from geometric parameters and generate meshes in multiple file formats, including surface mesh data such as DAGMC (.h5m) (Tautges et al., 2009) for neutron transport and unstructured meshes such as ExodusII (Schoof and Yarberry, 1994). Mesh metadata like sideset, block, and material names can be assigned and queried by component, material, and interface.

The decision to utilise Coreform Cubit as the provider of the underlying CAD technology, as opposed to alternative open-source alternatives such as Open Cascade (Open CASCADE SAS, 2019) or FreeCAD (Jürgen et al., 2024), was influenced by several factors, including its robust automatic meshing, existing integration with DAGMC, and its compatibility with many file formats. Moreover, Coreform Cubit is developing support for isogeometric analysis (Cottrell et al., 2009), which in future may provide a more efficient alternative to finite element analysis methods, thus making it a strategic selection. It should be noted that Hypnos itself is open source; provided that they have a Cubit licence, users may benefit from a number of convenience functions above a generic CAD library, to be detailed as follows.

10.3389/fnuen.2025.1694684

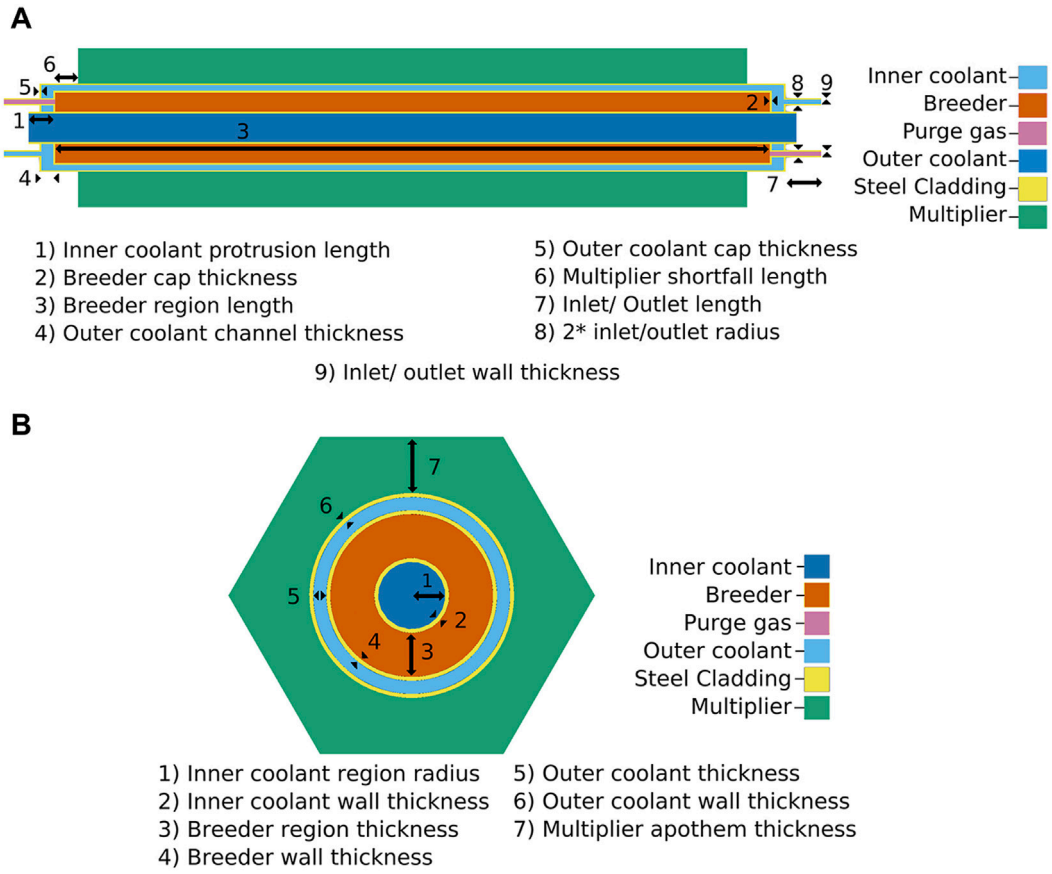

FIGURE 3
Parameterisation of the LIBRTI pin-cell geometry in Hypnos. The geometry is composed of three hollow steel cylindrical layers: an inner coolant channel, a breeder layer (containing solid breeder pellets), and an outer coolant channel; this is all surrounded by a hexagonal prism of multiplier material. Both coolant and purge gas annular volumes have smaller cylindrical inlets and outlets aligned with the z-axis. **(A)** Axial pin-cell parameterisation. **(B)** Radial pin-cell parameterisation.

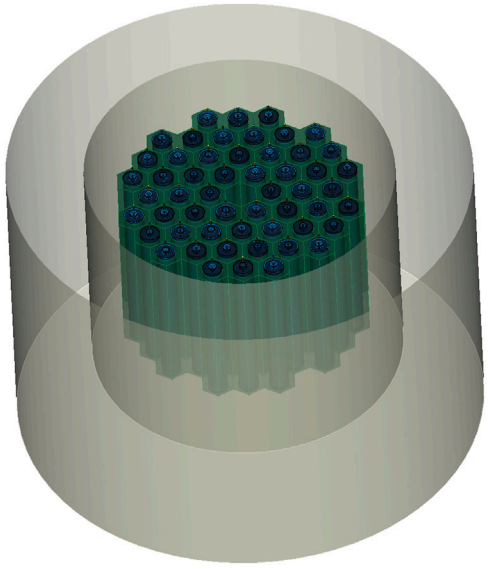

FIGURE 4
Visualisation of the full CAD geometry for the LIBRTI pin-cell assembly generated by Hypnos.

Distributed as part of the package is an initial implementation of a helium-cooled pebble bed (HCPB) breeder blanket (Zhou et al., 2023) which includes individual breeder pins, as well as channels and manifolds for the coolant and purge gas to flow through the system (Figure 2). For such existing geometries, Hypnos may be run as a standalone executable from the command line, with component parameter values supplied by JSON files and other behaviour controlled with a configuration file. Further detail is available in the online documentation[1].

Furthermore, Hypnos is designed to be modular and extensible, with component classes implemented hierarchically to reflect ontological relations. Users can implement custom classes for the creation and assembly of these components, inheriting from base classes that wrap Coreform Cubit functionality. This may be achieved using Hypnos as a Python library, which also makes available a suite of functions to help manipulate and create geometries.

1 Available at https://aurora-multiphysics.github.io/hypnos/

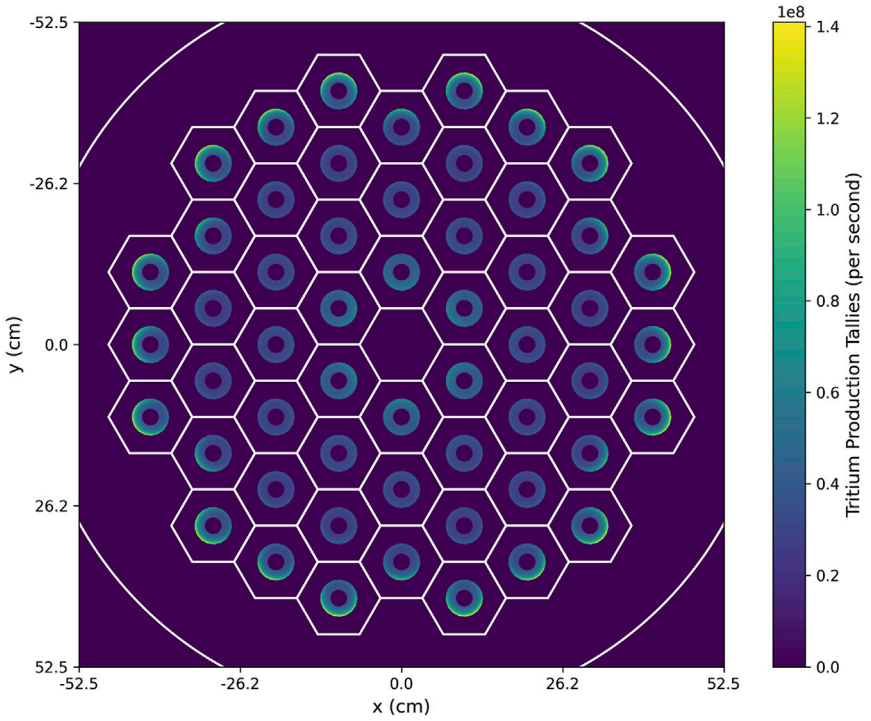

FIGURE 5
Spatial distribution of tritium production in the x–y plane for the baseline design using helium coolant with a temperature of 873.15 K. Tallies were scored by OpenMC on a regular mesh with a summation over contributions along the z-direction. The edge of each pin cell and the reflector's inner wall are overlaid in white.

The baseline model used as a reference point for the optimisation study in Section 3 is the LIBRTI straw-person conceptual design for a solid breeder mock-up experiment, originally presented in Gilbert et al. (2025). A parameterisation of the geometry was implemented using Hypnos (Figure 3). This parameterisation was chosen so that parameter bounds may be set independently; all lengths are set relative to the breeder region, while the radial build is codified using thicknesses. The structure of a single pin cell consists of hollow cylinders made of steel. The innermost cylinder is a channel for coolant (inner coolant channel). This is surrounded by a chamber containing the solid ceramic breeder pebbles, which additionally has an inlet and outlet at either end to direct purge gas through to extract produced tritium. This is further surrounded by another channel through which coolant is passed (outer coolant channel). Surrounding the outer coolant is a hexagonal prism of multiplier material. An arrangement of pins can be tessellated around an axially aligned neutron source (Figure 4). Finally, the assembly is surrounded by a volume acting as a neutron reflector.

2.1.2 Neutronics modelling with OpenMC

OpenMC (Romano and Forget, 2013; Romano et al., 2015) is selected for neutronics analysis. Not only is this a scalable open-source Monte Carlo neutral particle transport code, but it also integrates within the MOOSE (Multi-physics Object Oriented Simulation Environment) framework (Gaston et al., 2009; Permann et al., 2020; Giudicelli et al., 2024) via the application Cardinal (Novak et al., 2022). As noted earlier, the dependency of blanket requirements on interconnected physics domains requires a coupled multi-physics treatment. While, in this initial demonstration of our workflow, we remain restricted to a neutronics analysis, we intend to extend to multiple physics domains in the future, so it is prudent to have a strategy for this extension.

The material definitions assumed for the study are as follows. The structural material is taken to be 304 stainless steel. The neutron multiplier material is lead. The tritium breeding material is KALOS (KArlsruhe Lithium OrthoSilicate) ceramic pebbles (Leys et al., 2021), which contain a mixture of 65% Li_4SiO_4 and 35% Li_2TiO_3 homogenised over the pebbles.

The coolant was initially selected to be water, which was used for the standalone neutronics optimisation; in the second optimisation, the coolant was to be helium doped with 0.1% weight hydrogen for consistency with the thermal hydraulics analysis. In the first study, the temperature was set at 296 K (=22.85 °C); in the second study, the temperature was set at 873.15 K (=600 °C), again for consistency with the thermal hydraulics analysis. Finally, the pin-cell assembly is housed within a neutron reflector (Figure 4), with the material selected to be graphite.

It should be noted that the geometry is considered in isolation to its surroundings, such that the impact of the surrounding environment is not considered, with the reflector acting as an over-simplified proxy. As such, the results should not be

TABLE 1 Operating conditions and helium material properties used in thermal hydraulics calculation.

Parameter name	Symbol	Value/unit	Reference
Temperature	T	600 C	Gilbert et al. (2025)
Outlet pressure	p	8 MPa	Zhou et al. (2023)
Inlet mass flow rate	$\dot{m}$	0.045 kg/s	Zhou et al. (2020)
Surface roughness	ε	300 μm	Zhou et al. (2020)
Ratio of specific heat	γ	1.67	Petersen (1970)
Thermal conductivity	k	$2.682 \times 10^{-3}\,(1 + 1.123 \times 10^{-3}\,p)T^{0.71\,(1-2\times10^{-4}\,p)}$ W/(mK)	Petersen (1970)
Dynamic viscosity	μ	$3.674 \times 10^{-7} T^{0.7}$ kg/(m s)	Petersen (1970)

considered indicative of actual blanket designs. Notwithstanding this limitation, provided the results are interpreted in a relative sense (rather than as an absolute prediction of TBR), the analysis is nonetheless informative in a comparative study.

OpenMC was run with a fixed source using a continuous energy treatment of cross-sections and with photon transport enabled. Cross-sections were taken from the ENDF/B-VIII.0 library (Brown et al., 2018). The neutron source is an approximation of that found in commercially available neutron generators; Kulcinski et al. (2016) describe it as a line source, having a spatial variation in intensity along the axial z-direction. The energy distribution corresponds to mono-energetic 14 MeV neutrons. The resultant neutron flux distribution, and tritium production rate are scored on Cartesian grids (Figure 5). Here, the mesh resolution used was 1000 subdivisions in each direction; the number of histories was ten batches of 100 million particles. Radially, the peaks of tritium production are observed at the innermost and outermost points in the assembly, corresponding to breeder material closest to the source and to the reflector's inner wall, respectively.

In the study presented in Section 3, the figure of merit used is the global TBR. This is computed in OpenMC using the `H3-production` score tallied over the entire geometry; this returns the total number of tritium nuclei produced per source neutron. Calculating the global TBR does not rely on tracking the spatial distribution of tritium production, so a coarser mesh was used (and only for the purpose of interpreting results). Attaining a suitable level of convergence (standard deviation for global TBR on the order of 10^{-4}) required fewer histories—five batches of 10 million particles.

2.1.3 Thermal hydraulics

MOOSE is employed to perform thermal hydraulics modelling of the coolant. The approach uses the Compressible Euler Flow Model; the original implementation for RELAP7 is described in Berry et al. (2016). The governing equations, applicable to one-dimensional single-phase flow, are denoted by Equation 1-3:

$$\frac{\partial(\rho A)}{\partial t} + \frac{\partial(\rho A u)}{\partial x} = 0, \tag{1}$$

$$\frac{\partial(\rho u A)}{\partial t} + \frac{\partial(\rho u^2 + p)A}{\partial x} = p\frac{\partial A}{\partial x} - FA + \rho g A, \tag{2}$$

$$\frac{\partial(\rho E A)}{\partial t} + \frac{\partial u(\rho E + p)A}{\partial x} = \rho u g A + q''' A. \tag{3}$$

These equations correspond to balance of mass, momentum, and energy, respectively. Here, A is the cross-sectional area, ρ is density, u is the axial velocity, E is the specific total energy, p is pressure, F is the viscous drag force, g is the gravitational acceleration projected along the axial direction, and q''' is the volumetric heat source.

In the baseline model, a given breeder pin is cooled by an inner cylinder and an outer annulus of helium coolant. As the most simplistic extension of the previous study, only the outer channel is considered in this analysis. This should not be considered representative of full thermal hydraulics but rather a first test of a simple multi-objective problem.

The helium coolant is treated as an ideal gas with approximately constant viscosity, thermal conductivity, and specific heat. The pressure drop of the coolant flowing through the annulus is computed assuming a fixed operational temperature, mass flow rate, and outlet pressure. This may be treated using a single thermal hydraulic flow channel of fixed length and hydraulic diameter given by $D_H = D_o - D_i$, which is the difference between the inner and outer diameters. A Churchill friction factor is assumed to apply at the walls of the channel. The assumed values of all operational and helium material input parameters are summarised in Table 1.

Whilst the approach described here constitutes a fairly simplistic treatment, in the context of the optimisation study to be presented in Section 3, it should nevertheless introduce some penalty that prevents the minimisation of the coolant volume. Even if restricted to 1D thermal hydraulics, there are a number of improvements which could be considered in future; these are discussed in Section 5.

2.1.4 Optimisation with SLEDO

SLEDO (Sequential Learning Engineering Design Optimiser) (Humphrey et al., 2024) is a software tool which performs Bayesian optimisation, selecting candidate designs from a parametric search space and evaluating them by deploying a "black-box" function. Developed with the deployment of computationally expensive multi-physics simulation workflows in mind, the scalable parameter search framework Ray Tune (Liaw et al., 2018) was recently integrated to handle HPC deployment. A range of Bayesian optimisation libraries are available to drive candidate

TABLE 2 Parameters varied and their bounds for the neutronics optimisation.

Parameter name	Lower bound	Upper bound	Reference value
Breeder region thickness	4.0625 mm	65.0 mm	16.25 mm
Multiplier apothem thickness	5.225 mm	83.6 mm	20.9 mm

TABLE 3 Sampling strategies employed in neutronics-only optimisation.

Strategy	ξ	Sobol sample size n
1	2.5	32
2	0.01	8

selection; for this study, Ray Tune's implementation of the BayesOpt package (Nogueira, 2014) was used to provide functionality for the search algorithm.

The procedure employed by SLEDO is as follows.

1. (Optional) An initial population of candidate designs is selected quasi-randomly by a sample plan run over a bounded search space $\mathbf{x}$.
2. Candidate designs are evaluated to arrive at a scalar-valued figure of merit $f(\mathbf{x})$.
3. A surrogate model is trained, from which posteriors for the mean $\tilde{f}(\mathbf{x})$ and standard deviation $\sigma(\mathbf{x})$ of f are obtained.
4. An acquisition function $A(\mathbf{x})$ is computed and maximised to generate subsequent candidates for evaluation.
5. Candidate designs are generated and evaluated until a fixed number of trials have concluded.

The role of the surrogate model is to use the information gained from all previous trials to predict the performance of a proposed design candidate. This statistical model should be relatively cheap to train and evaluate, such that in each iteration it may be updated and the acquisition function run over the entire search space. In this study, a Gaussian process (GP) is used—a generalisation of a Gaussian distribution to arbitrary dimensions (Williams and Rasmussen, 2006).

The role of the acquisition function is to generate candidates in a sequence which most efficiently converges on an optimised design. To achieve this, the function must be able to balance exploitation (sampling where the design is predicted to perform well) and exploration (sampling where its performance is uncertain). Many acquisition functions are available; in this study, the expected improvement (EI) is utilised, defined by Equation 4, 5:

$$A(\mathbf{x}) = \left(\tilde{f}(\mathbf{x}) - f(\mathbf{x}^+) - \xi\right)\Phi(Z) + \sigma(\mathbf{x})\phi(Z), \tag{4}$$

$$Z = \tilde{f}(\mathbf{x}) - f(\mathbf{x}^+) - \xi. \tag{5}$$

Here, $\mathbf{x}^+$ is the current set of best design parameters; Φ and ϕ are the cumulative distribution function and probability density function of the normal distribution, respectively; ξ is a tuneable hyperparameter which balances the trade-off between exploitation and exploration.

While optional, the role of the initial sample plan is to provide an initial set of candidate designs which achieve good coverage across $\mathbf{x}$. This provides a reasonable prior to the surrogate model at the onset of the Bayesian optimisation sequence. Without initial data, $A(\mathbf{x})$ returns a flat distribution, so the first candidate would be randomly selected. For the following iteration, only the exploration term would vary across $\mathbf{x}$, so the most distant point would be selected. Recognising that exploration of the search space is most valuable in these early iterations, while the identification of local maxima is preferable to convergence on the global maximum, one advantage of using a sample plan is that it allows the acquisition function itself to be tuned more in favour of exploitation. Another advantage is that the initial candidates may be evaluated in parallel, since they do not depend on each other. This is particularly useful for applications where time-to-solution is valuable, though for simplicity of deployment, designs were evaluated sequentially in this study. The sample plan used in this study is a Sobol sequence (SI, 1967), using the SciPy implementation (Virtanen et al., 2020) from which a number of samples must be drawn to achieve a balanced spread through the search space. The ideal number of Sobol samples to be drawn before switching to Bayesian optimisation, which is set to be a power of 2 $n = 2^m$, depends on the problem's dimensionality and expected degree of non-linearity; in this study, the number was varied between 8 and 32.

2.2 Definition of optimisation campaigns

2.2.1 Definition of neutronics optimisation campaign

As an initial test of the overall pipeline, a standalone neutronics analysis was performed, and the objective function was set to

TABLE 4 Parameters varied and their bounds for the combined neutronics and thermal hydraulics optimisation.

Parameter name	Lower bound	Upper bound	Reference value
Breeder region length	460.5 mm	767.5 mm	614 mm
Breeder region thickness	5.4167 mm	48.75 mm	16.25 mm
Multiplier apothem thickness	1.6167 mm	14.55 mm	20.9 mm
Outer coolant region thickness	6.966 mm	62.69 mm	4.85 mm

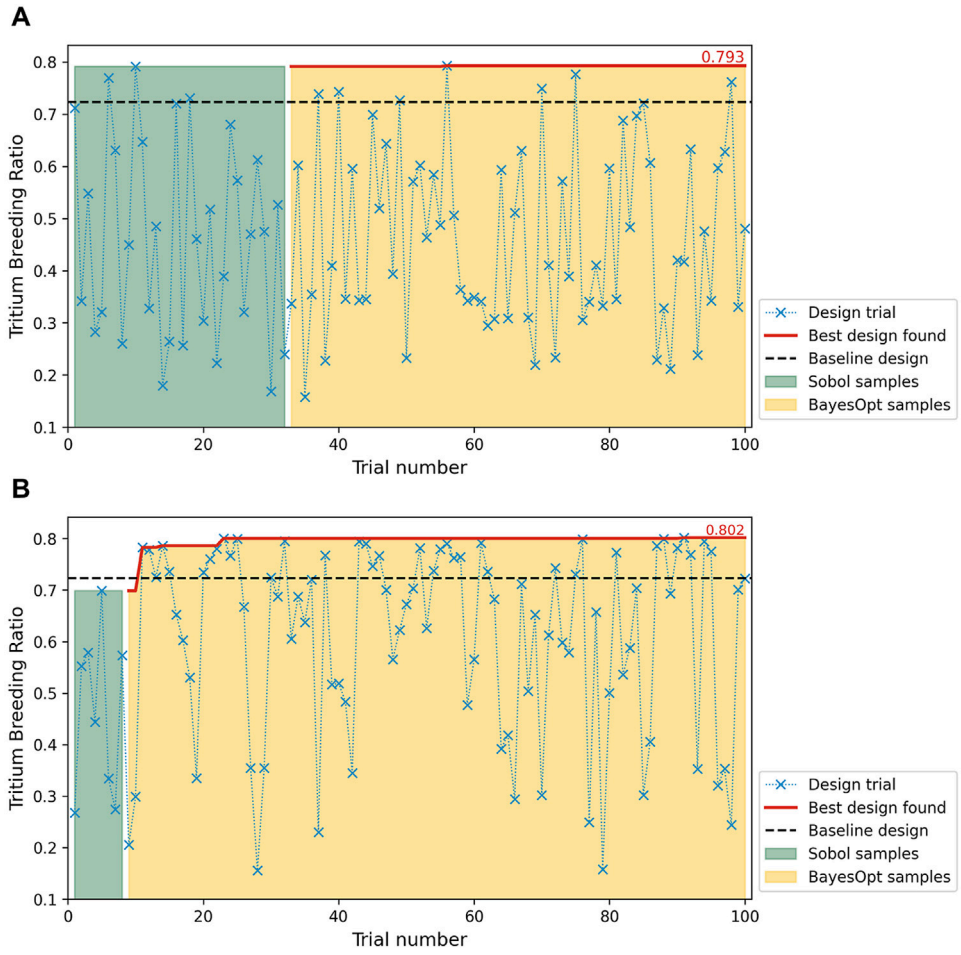

FIGURE 6
Figure of merit (TBR) for the neutronics-only optimisation for each trial, plotted in the order in which they were evaluated. Solid red line shows the evolution of the best design found thus far at each step of the sequence. Dashed black line shows the performance of the baseline design. Candidates selected by Sobol sequence are shown with a green background on the left; those selected by Bayesian optimisation are shown with a yellow background on the right. **(A)** Figure of merit for strategy 1 with 32 Sobol samples and $\xi = 2.5$. **(B)** Figure of merit for strategy 2 with 8 Sobol samples and $\xi = 0.01$.

maximise a single figure of merit: the TBR. A disadvantage of geometric optimisation performed in isolation from other physics is that components with functions that do not pertain to increasing this metric would be minimised unnaturally. Therefore, in this initial analysis only two parameters with a function directly pertaining to tritium breeding were allowed to vary; these, along with their ranges, are shown in Table 2. In both cases, the upper and lower limits are obtained by multiplying the reference value used in the baseline design (also shown in the table) by a factor of $\frac{1}{4}$ and 4, respectively. Both of these parameters control radial thicknesses within each pin (Figure 3B). In addition, the total number of pins in the assembly is not fixed but is computed as the maximum number that may be fully contained within the inner radius of the reflector (additionally stipulating a minimum gap of 75 mm between this and the outermost pin). The factors of $\frac{1}{4}$ and 4 were chosen somewhat arbitrarily to allow designs to significantly vary the number of pin cells while avoiding designs featuring sub-millimetre scale thicknesses at one extreme and pin cells too large for the reflector at the other. The baseline design is taken as the midpoint (logarithmically) to avoid biasing the search in either direction over the other.

The surrogate model employed was a Gaussian process, while the acquisition function was "expected improvement," as defined by Equation 4. Two optimisation campaigns were performed which employed different sampling strategies; these are summarised in Table 3. Each campaign used a fixed simulation budget of 100 trials but differed in how candidate designs were selected. In the first campaign, a large value of ξ was used, corresponding to a preference for exploration over exploitation, and a relatively high proportion of Sobol samples were generated. In the second campaign, a lower value of ξ was used, corresponding to increased preference for exploitation, and a relatively low proportion of Sobol samples were generated.

2.2.2 Definition of neutronics and thermal hydraulics optimisation campaign

As noted earlier, optimising a geometry when considering a single physics domain in isolation would result in certain volumes being minimised where they serve no role in augmenting the given performance metric. For example, as the coolant does not act to increase TBR, is not reasonable to optimise those parameters which

10.3389/fnuen.2025.1694684

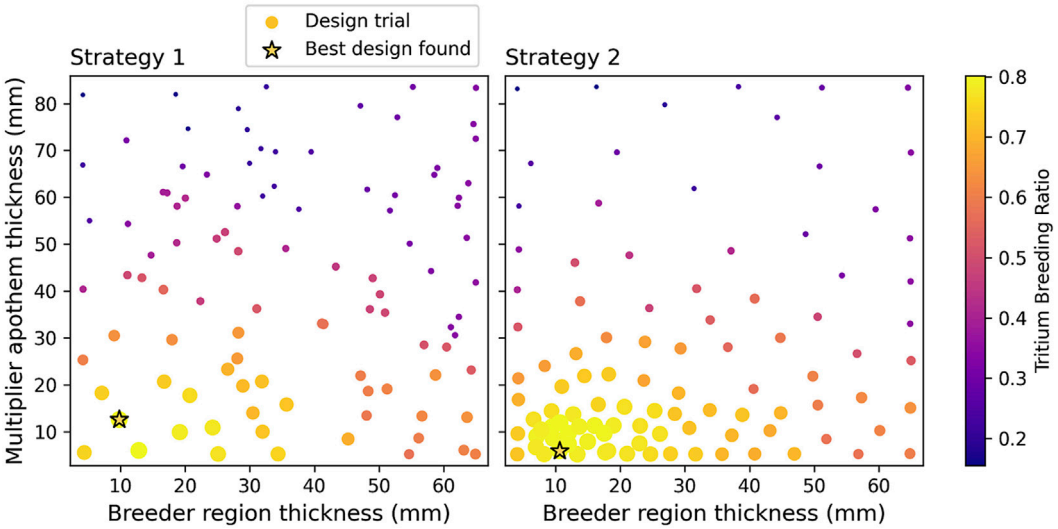

FIGURE 7
Distribution of samples within the two-dimensional parameter space for the neutronics optimisation. The two panes correspond to two different values of the hyperparameter: $\xi = 2.5$ (left) and $\xi = 0.01$ (right). The colour scale and size of the marker indicate the corresponding value of TBR attained for the design point. The best design is indicated by a star marker.

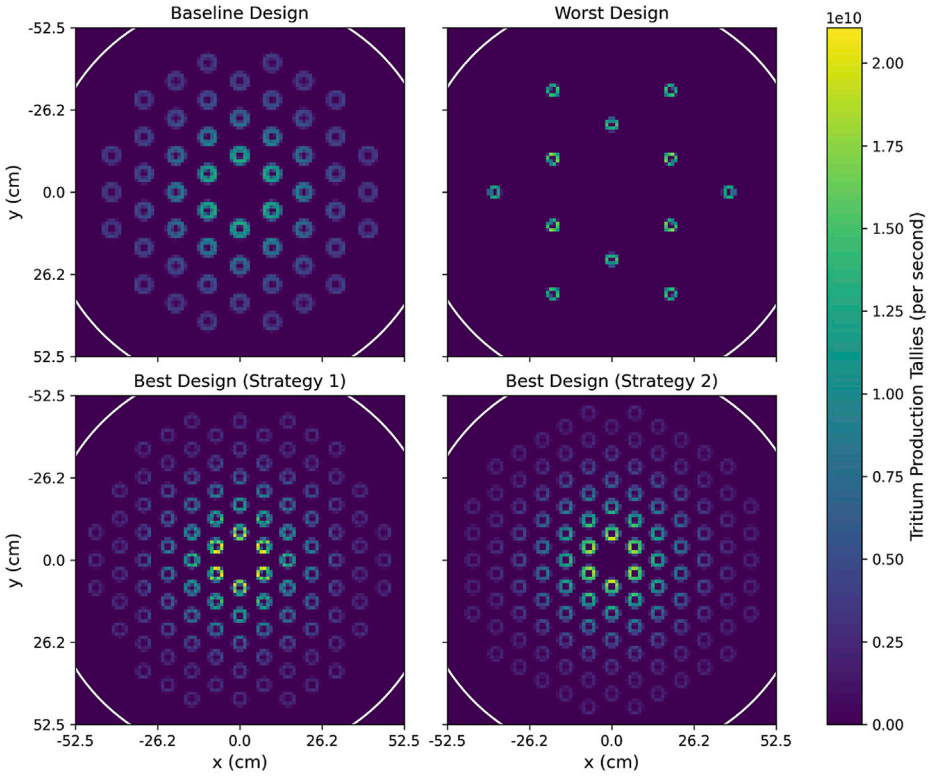

FIGURE 8
Spatial distribution of tritium production in the x–y plane for the baseline, worst, and best designs arising from the neutronics-only optimisation. Tallies were scored by OpenMC on a regular mesh, with a summation over contributions along the z-direction.

TABLE 5 Input and output parameters for the optimised design given to three significant figures alongside the values for the baseline design.

Parameter name	Optimised value	Baseline value
Breeder region length	761 mm	614 mm
Breeder region thickness	34.5 mm	16.25 mm
Multiplier apothem thickness	9.37 mm	20.9 mm
Outer coolant region thickness	1.62 mm	4.85 mm
Tritium breeding ratio	0.729	0.474
$\mathrm{TBR}/\mathrm{TBR}_{\mathrm{ref}}$	1.04	0.676
Pressure drop	2.08×10^4 Pa	8.42×10^2 Pa
$1 - (p_{in} - p_{out})/p_{in}$	0.997	1.00
Figure of merit	2.04	1.68

control the coolant if simulating only neutronics. Further exploration of the methodology within a higher dimensional parameter space dictates a more complete treatment of physics. Therefore, a second study evaluated the performance of each design by computing metrics arising from both neutronics and thermal hydraulics simulations. The set-up of the modelling was described in detail in Sections 2.1.2 and 2.1.3.

The two metrics used were the TBR and the pressure drop across the pin (computed as the difference in pressure between inlet and outlet), which should be maximised and minimised respectively. The objective function is then defined as a weighted sum of two metrics:

$$f = \frac{\mathrm{TBR}}{\mathrm{TBR}_{\mathrm{ref}}} + \left(1 - \frac{p_{in} - p_{out}}{p_{in}}\right). \tag{6}$$

In the first term, the evaluated TBR is normalised to a reference TBR value, taken to be $\mathrm{TBR}_{\mathrm{ref}} = 0.7$. In the second term, the pressure drop is normalised to the inlet pressure; as this relative quantity cannot exceed unity, subtracting this from 1 results in a positive-definite quantity. Moreover, maximising this quantity is equivalent to minimising the pressure drop, and thus may be additively combined with the first metric. In principle, the relative importance of each of these terms could be controlled by multiplying the second term by an optional hyperparameter; however, since the domain of each term is of order unity, this was deemed to be a suitable function. Nevertheless, it is worth emphasising that the ambiguity in the definition of "objective function" results in some arbitrariness in the definition of "optimal". This matter is discussed further in Section 5.

The parameters varied, with their upper and lower limits shown in Table 4. Note that since only the outer coolant loop is considered by the model, the geometry of the inner coolant channel is fixed in this campaign. The baseline design is used as the basis for each set of bounds: all three radial parameters were varied between a factor of $\frac{1}{3}$ to 3 relative to the reference value. The rationale for these limits, as in the previous study, is to focus the search around the baseline design and avoid extreme values. In this case, the factors are reduced in magnitude due to the fact that an additional radial thickness is being varied, which otherwise increases the range of overall pin-cell diameters allowed by the bounds. However, note that these factors remain only arbitrary estimations used *in lieu* of precise requirements. The breeder region length was varied by factors of 0.75–1.25 relative to the reference value. Unlike the radial build, in which the number of pin cells is adjusted to make use of the space available, the pin-cell length is directly constrained by the reflector height, so that the upper limit is set accordingly (allowing some degree of clearance above and below, as a precise clearance requirement is not yet defined). With these additional parameters, it is possible to explore the trade-off between increasing the TBR though increased breeder and multiplier volumes or reducing pressure drop through increased coolant region thickness and decreased pin-cell length.

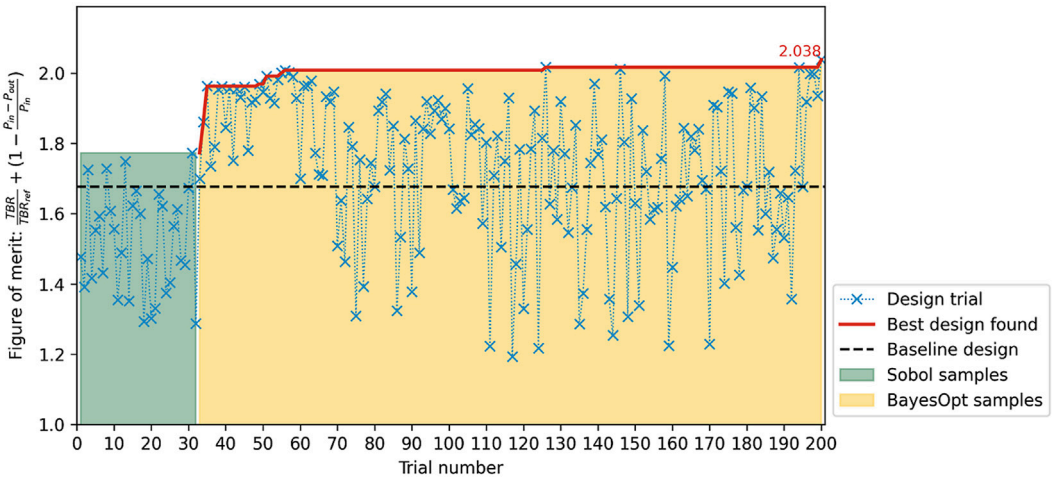

FIGURE 9
Figure of merit for the neutronics and thermal hydraulics optimisation for each trial, plotted in the order in which they were evaluated. Solid red line shows the evolution of the best design found thus far at each step of the sequence. Dashed black line shows the performance of the baseline design. Candidates selected by Sobol sequence are shown with a green background on the left; those selected by the Bayesian optimisation are shown with a yellow background on the right.

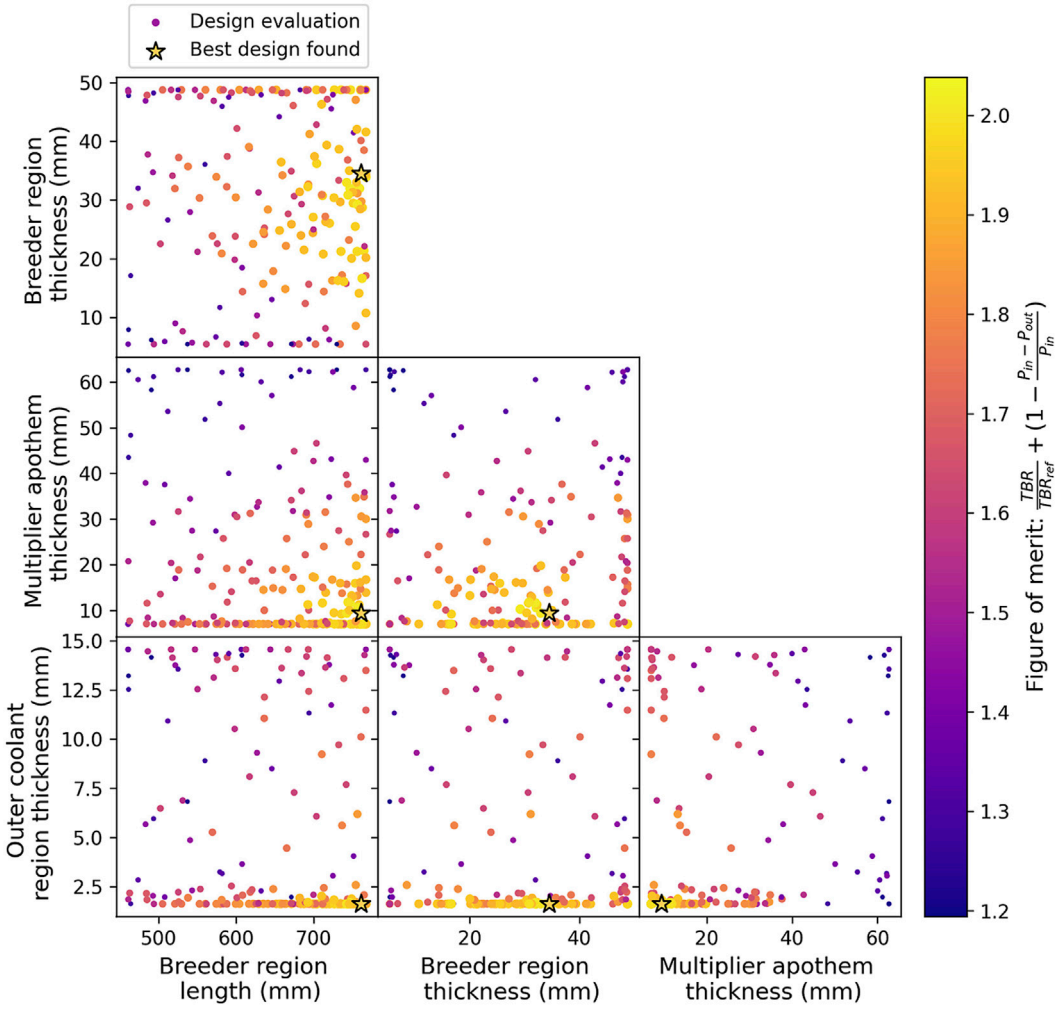

FIGURE 10
Distribution of samples within the four-dimensional parameter space for the combined neutronics and thermal hydraulics optimisation. Each pane corresponds to a unique projection of the space onto two of the parameters. The colour scale and size of the marker indicate the corresponding value of the figure of merit attained for the design point. The best design is indicated by a star marker.

Following strategy 2 from the neutronics-only optimisation, the surrogate model is a GP and the acquisition function is EI, with exploration parameter ξ kept at 0.01. Given the increased dimensionality, the number of initial Sobol samples was increased to 32.

3 Results

3.1 Baseline results

The global TBR for the baseline design was found to be 0.723 for the first study. Since this application corresponds to an experiment, while it is desirable to maximise the TBR (as this would ultimately improve the signal to noise ratio of any measurements), it is not strictly required to attain a target TBR in excess of unity (as would be required for a self-sustaining fuel cycle). Indeed, this would require higher coverage in a solid angle around the neutron source, whereas this conceptual experimental design is intended to enable modularity.

In the second study, where the coolant was changed from water coolant at 296 K to helium coolant at 873.15 K, the TBR changed to 0.474. The baseline pressure drop in the outer coolant was 8.42×10^2 Pa.

3.2 Neutronics optimisation

Here, we present the results of the optimisation campaign defined in Section 2.2.1.

The trace of the figure of merit as a function of the trial number for the two sampling strategies of Table 3 is represented in Figures 6A,B, respectively. For strategy 1, the TBR of the baseline design was exceeded within the initial Sobol sequence, but it was not significantly improved upon during the Bayesian optimisation sequence; for strategy 2, the reference TBR was exceeded only a

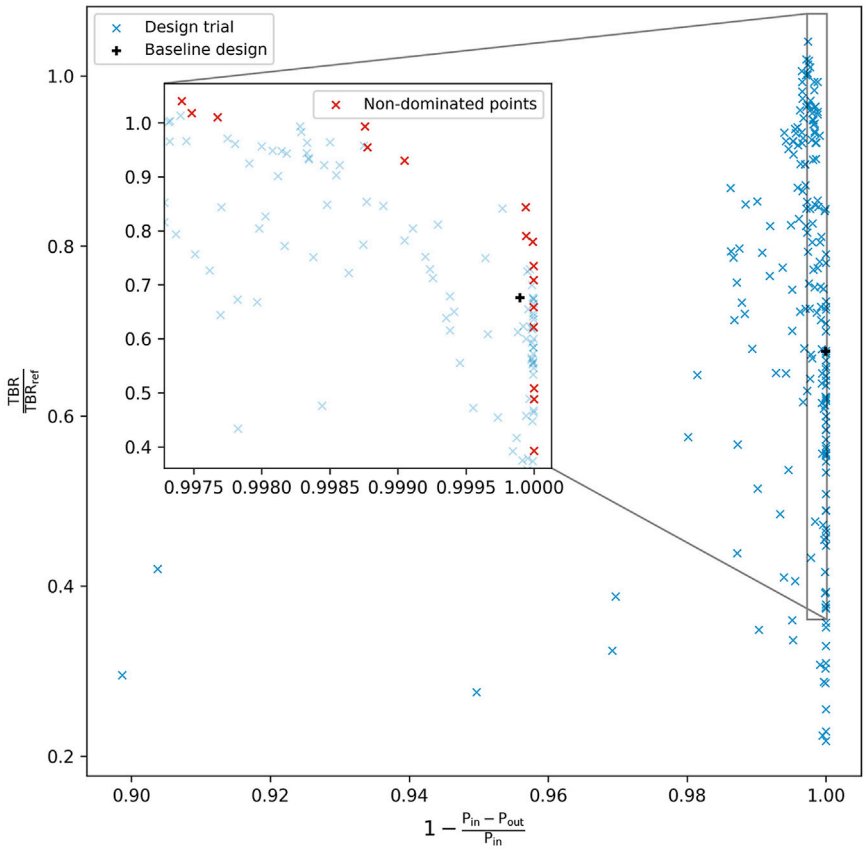

FIGURE 11
Comparison of trial performance in each term of the figure of merit. Design trials are shown with blue crosses, while the baseline design is shown as a black plus. The inset highlights non-dominated points as red crosses.

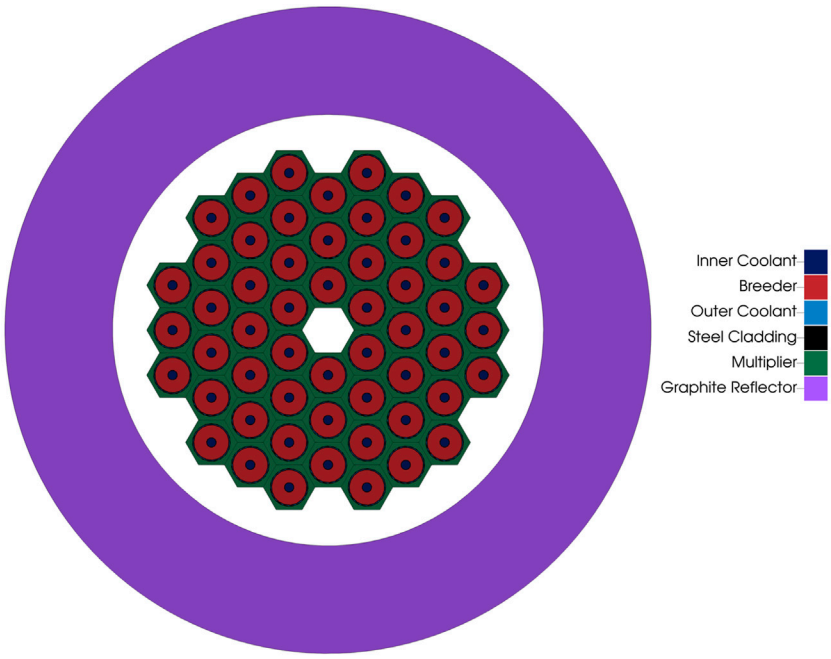

FIGURE 12
Cross-section geometry for best-performing pin-cell assembly design arising from the combined neutronics and thermal hydraulics optimisation.

few trials into the Bayesian optimisation sequence. The best design found by strategy 1 achieved a TBR of 0.791, which was further improved to 0.802 by strategy 2.

Figure 7 shows the distribution of selected candidates from each strategy. Strategy 1 features a good spread of points throughout the space but fails to cluster in the performant region, while strategy 2 shows good clustering with sufficient spread throughout the space; this difference corresponds to the choice of hyperparameter ξ, with the low-value in strategy 2 reducing the weight of uncertainty in the acquisition function's choice of candidate. For this two-dimensional search space and fairly linear response surface, strategy 2 (favouring exploitation) was preferable. However, it should be noted that as the dimensionality is increased, the size of the search space increases exponentially, and thus the importance of exploration will become more pronounced.

The optimised designs for each strategy are shown in Figure 8, alongside the baseline design and the worst-performing design as a point of comparison. These results are discussed further in Section 4.1.

3.3 Neutronics and thermal hydraulics optimisation

Here, we present the results of the optimisation campaign defined in Section 2.2.2.

The trace of the figure of merit (now defined by Equation 6) is shown in Figure 9; it improves rapidly during the first few tens of Bayesian optimisation trials, then exploration dominates candidate selection until the final few trials, where it is further improves. Indeed, the best design found was the final trial.

The distribution of trials over input space is shown in Figure 10. As before, the exploration–exploitation trade-off achieves a balanced spread of points with clustering around performant regions. Note that compared to Figure 7, the increased dimensionality of the space means that nearby points in one subplot are nearby in only two dimensions but not necessarily close in value of the other two dimensions not shown.

Compared with the baseline design, the figure of merit was increased from 1.68 to 2.04, corresponding to an improvement in TBR (which increased from 0.474 to 0.729) but a worsened pressure drop (which increased from 8.42×10^2 Pa to 2.08×10^4 Pa). Whether this is an acceptable trade-off should be considered alongside appropriate domain expertise. More generally, the set of designs which cannot improve in one metric without reducing performance in another is known as the *Pareto front*. The non-dominated points, which provide an approximation to this front, are shown in Figure 11. An alternative approach to multi-objective optimisation is to seek these Pareto-optimal designs—discussed further in Section 5.2.1.

The best design found is visualised in Figure 12. The input parameters and evaluation metrics are shown in Table 5. These results are discussed further in Section 4.2.

4 Discussion

4.1 Neutronics optimisation

Many smaller pin cells were preferred over fewer larger ones, indicating a preference for the breeder and multiplier to be distributed as homogeneously as possible. This is in line with the expectation for fixed source problems, where the optimal result is typically one where material boundaries are fully blurred and there is a closer-to-equal probability of interacting with any material in the problem. This maximises the probability of sequential moderation and absorption, where in this context "moderation" refers to the scattering of a neutron into the thermal region of the Li_6 cross-section and "absorption" refers to a Li_6 tritium production reaction.

4.2 Neutronics and thermal hydraulics optimisation

Compared with the neutronics-only results, the best design found now uses the same number of pin cells as the baseline design, albeit with a different radial build. A thicker breeder region is used at the expense of multiplier and coolant. However, a larger breeder zone implicitly increases the cross-sectional area of the annular flow channel for a fixed thickness. Therefore, the inclusion of the pressure drop in the figure of merit implicitly placed a penalty on smaller pins. This finding reinforces the notion that trade-offs are often subtle, and a holistic optimisation is essential.

Within the bounds of the search space, two parameters reached extreme values, the outer coolant region thickness being minimised and the breeder length being maximised to their lower and upper bounds, respectively. This suggests that the true optimum for these parameters lies outside the search-space bounds. The preference for increasing the pin length suggests that the cost of increasing the coolant length is not as detrimental to increasing the pressure drop as it is favourable to increasing the volume of breeder material (and hence TBR). Therefore, the limiting factor is simply the overall size, as here constrained by a fixed reflector volume. Similarly, the cost of reducing the coolant hydraulic diameter (also acting to increase pressure drop) is not so severe as to prevent the preference for increased volume for breeder and multiplier regions. We note that the efficacy of cooling was not taken into consideration; were this to be included as an additional objective of the optimisation, it would further penalise a reduction in hydraulic diameter. In addition, there are also likely to be practical constraints upon the minimum coolant diameter arising from manufacturing considerations, as well as structural considerations for the steel cladding.

Furthermore, it should be understood that the locations of these true optima are dependent on the relative weights of the metrics in the figure of merit (as defined by Equation 6), which were selected somewhat arbitrarily. Thus, it is possible that with different weightings, the optimal values currently outside the search space could move within the search space and *vice versa* for those currently within the space. That is to say, the optima arrived at here are only those for this particular choice of objective function.

The impact upon sampling arising from the choice of objective function is also apparent in the distribution of trials over the two performance indicators (Figure 11). The x and y coordinates here correspond to the first and second terms in Equation 6, and their sum evaluates to the figure of merit. Therefore, optimal designs are to be found in top-right corner of the plot. The TBR is fairly sensitive to the sampling strategy (as shown by the broader distribution of points); meanwhile, the range for the normalised pressure drop is

extremely narrow and sparsely sampled in the tail. Thus, while the metrics' absolute values are of similar magnitude, their variance is not. A repeat of the optimisation would suggest that a change of variables for the pressure drop performance indicator within the figure of merit to flatten this distribution might be desirable.

5 Outlook

In this study, we have provided an indicative analysis for the purpose of demonstrating the potential of an integrated workflow for parametric design optimisation of a fusion breeder experiment. However, a number of limitations exist in regard to both the modelling and optimisation methodologies that were presented in Section 2.1, and here we describe a number of potential improvements.

5.1 Modelling improvements

5.1.1 Neutronics modelling improvements

From the perspective of neutronics analysis, it was already acknowledged that the geometry of the breeder mock-up technology was modelled independently from the room (although a graphite reflector was included to account for re-entrant thermalised neutrons). As information on the intended facility becomes available, it would be desirable to understand the impact of the room and shielding and assess any asymmetries introduced as a result. While such high-fidelity modelling might be impractical from the standpoint of an optimisation study with many iterations, nevertheless, depending on compute budget and appetite, even a one-off evaluation would provide context for the comparative results yielded by the optimisation.

5.1.2 Thermal hydraulics modelling improvements

Thermal hydraulics were included to introduce a penalty arising from either reducing the coolant cross-sectional area or from increasing the length of the pin (both acting to undesirably increase the pressure drop). However, the treatment was particularly simplistic, and several improvements could be made. Rather than consider the entire fluid manifold through the pin assembly, we considered flow through a single pin in isolation, assuming this to be sufficiently representative of all pins in the assembly and neglecting any spatial variations in operational conditions. Somewhat unphysically, the mass flow rate per pin was considered fixed; in reality, for a fixed pumping power or total mass flow rate, flow per pin would also depend on the number of pins, a quantity that in turn depends on their diameters. In addition, only the pressure drop corresponding to the annular outer coolant was computed; the inner coolant was omitted from consideration, its diameter being considered fixed. Assuming a single coolant loop, these two flow channels should not be treated independently. In future research, we propose to model the complete flow manifold (including the impact of pumping) as this would enable a more holistic examination of the space of geometric configurations of the coolant.

5.1.3 Multi-physics modelling improvements

For both neutronics and thermal hydraulics, the temperature was considered fixed, nor were any feedback mechanisms included. The impact of neutronic heating for an experiment that is anticipated to have comparatively low neutron fluxes (which should be on the order of $10^{13}s^{-1}$ within the LIBRTI facility) is expected to be fairly minor; in full blanket systems, however, this heating is likely to be considerable, and a coupled treatment would be necessary. An integration of OpenMC within MOOSE via Cardinal (Novak et al., 2022) is already available, so the inclusion of these effects is already feasible. More important for the current context is the impact of conjugate heat transfer to the solid structures, which was not considered. It is intended that the breeder experiment be exposed to heating, with the role of the coolant being to maintain operational temperature. As the coolant parameters would certainly impact the efficacy of this heat extraction and may result in gradients in temperature, it would be worthwhile repeating the optimisation with a full thermal analysis. This would also allow the efficacy of cooling to be added as a figure of merit, which may ultimately impose a stronger penalty on coolant reduction. Finally, the bounds of geometrical parameters should be informed by practical manufacturing and structural limits.

Another limitation of this initial study was its selection of tritium breeding ratio as a performance indicator. The metric only accounts for the tritium produced *in situ*, making no attempt to model parasitic losses of tritium via permeation into, and retention within, solid structures. Therefore, a more informative and holistic indicator of the breeder technology is the tritium removed from the system within the purge gas, which requires a proper treatment of tritium transport within the system. As the migration of tritium is highly sensitive to temperature, this would also involve a coupled multi-physics approach to modelling. Within the MOOSE ecosystem of software, a suitable and well-validated implementation already exists in TMAP8 (Simon et al., 2025). Employing TMAP8 in an integrated manner to perform the optimisation using a more realistic performance indicator would be a natural extension to this research.

5.1.4 Cost modelling

Another practical consideration relevant to the design of breeder blanket technology not yet explored here is the cost of materials. Depending on the specific choice of structural, multiplier, and breeder materials, increasing the volume of certain components may correspond to a dramatically increased cost. For a given cost model (for example, as implemented within the systems code Process (Kovari et al., 2016; Morris et al., 2023)), it should be straightforward to include cost as a figure of merit in the optimisation procedure demonstrated here. However, this would be predicated on the availability of data to estimate the price of those materials, and for many advanced materials employed in fusion environments, these are generally unknown for production at scale. Capital costs are influenced by a variety of factors, including natural abundance, safety of handling, maturity of manufacturing techniques, and availability of supply chain; as noted by Lux et al. (2024) and Chapman (2025), estimates of capital costs often carry large uncertainties that sometimes exceed the difference in cost

between materials. Therefore, additional research would be required to meaningfully incorporate the impact of cost.

5.2 Optimisation methodology improvements

5.2.1 Multi-objective optimisation

It should be apparent from the discussion thus far that there is, in general, more than one indicator of performance for assessing breeder blanket technology. The approach taken in this study, known as linear scalarisation, combines figures of merit as a weighted sum of metrics and optimises for this single scalar quantity. As can be seen in Figure 11, candidates chosen in this fashion tend to favour optimisation of that metric which dominates that sum. Provided that the relative weights are chosen appropriately, this may appear as unproblematic or even desired behaviour. In practice, however, even with domain expertise, the selection of weightings *a priori* can be challenging and presents a significant source of risk.

An alternative approach to multi-objective optimisation (MOO) is to actively seek the set of design candidates on the Pareto frontier, which removes the need to weigh metrics. Rather than returning a single optimised design, a number of such Pareto-optimised designs are returned, each of which offers alternative trade-offs between objectives. By providing the full spectrum of possibilities within the trade-off space, a more informed decision can be made on which Pareto-optimised design(s) to progress.

This approach to MOO has seen recent adoption in fusion engineering applications. In stellarator design, Bindel et al. (2023) describe the use of continuation methods for magnetic field coil design, exploring the trade-offs between plasma quality and coil complexity, while Packman et al. (2025) describe, using Bayesian methods, coil optimisation with simultaneous magnetic and mechanical objectives. In tokamaks, Nunn et al. (2023) explored the use of Bayesian methods for coil design using the acquisition function Expected Hypervolume Improvement (EHVI) (Daulton et al., 2020), which is the multi-objective equivalent to EI using in this study. The noise-aware variant (q-Noisy EHVI) has also been used by UKAEA to investigate the optimisation of plasma current drive profiles (Brown et al., 2023).

There are two main options for implementing such an approach into SLEDO. Given that the latest version of SLEDO uses Ray Tune for deployment on HPC, the clearest approach would be to utilise existing support for multi-objective within Ray Tune. However, at the time of writing, support for this feature is limited; only one of the Bayesian optimisation tools natively integrated with Ray Tune, Optuna (Akiba et al., 2019), features multi-objective optimisation. If additional customisation control is desired over how the optimisation is performed than is offered out of the box by Ray Tune and Optuna, an alternative approach would be to implement the optimisation loop directly using a lower-level Bayesian optimisation framework such as BoTorch (Balandat et al., 2019), which was used in earlier versions of SLEDO.

5.2.2 High-dimensionality optimisation

The scope addressed in this study considered only geometric parameters as the space of engineering design choices. In principle, provided that there exists a suitable "model instantiator" for the corresponding input space (cf. Figure 1), the optimization techniques demonstrated here could be used for a wider array of design choices. Of particular relevance for the neutronics analysis presented above is the composition of materials. One mechanism to incorporate this seamlessly into the current workflow might involve standardized material naming within a property database, incorporating these into the metadata. Leveraging Hypnos' querying capability, it would then be possible to codify fetching from and subsequently perturbing the corresponding property data automatically.

Addressing a wider and more generic set of input parameters, especially for the design problem of breeder blankets, naturally inflates the dimensionality of the problem. The scalability of Bayesian methods is problem-specific; while they are theoretically scalable to an arbitrary number of degrees of freedom d, they are typically considered to only remain effective up to $d \leq 20$ for most problems (Frazier, 2018). As full blanket designs are likely to present search spaces of $d > 20$, it would be advisable to employ one or more high-dimensional Bayesian optimisation (HDBO) techniques. HDBO is an active area of research, with an overview given by Bindel et al. (2023). Options such as dimensionality reduction (Wang et al., 2016), sparse axis-aligned subspaces (SAAS) (Eriksson and Jankowiak, 2021), or trust-region Bayesian optimization (TuRBO) (Eriksson et al., 2020) may prove effective.

6 Conclusion

In this study, we present a digital engineering pipeline for generating, evaluating, and optimising fusion breeder technology. This involved a number of open-source software packages. A new tool, Hypnos, was implemented to instantiate geometry from input parameters and encapsulate the design point. OpenMC and MOOSE were employed to perform neutronics and thermal hydraulics analyses, respectively. Finally, SLEDO was used to perform sampling of the design space and optimise the stipulated figure of merit.

The target application selected for optimisation is a conceptual design for the solid breeder experiment at the planned LIBRTI facility, the purpose of which is to provide qualification of the tritium breeding technology and validation of fundamental tritium transport modelling capability.

Two optimisation studies were performed. In the first, the analysis only considered neutronics, and the selected figure of merit was the TBR, which was improved from the baseline value of 0.723–0.802. The optimised design was found to have a greater number of smaller pins, indicating a preference for a more homogenised arrangement.

In the second study, the analysis was extended to include a simple treatment of thermal hydraulics for a representative pin. The figure of merit was updated to be a weighted sum formulated from the TBR and pressure drop over the outer coolant. Compared with the baseline design, the figure of merit was increased from 1.68 to 2.04, corresponding to a TBR increased from 0.474 to 0.729 at the cost of pressure drop, which increased from 8.42×10^2 to 2.08×10^4 Pa. In this case, the optimised design did not change the number of pins relative to the baseline design but, rather, increased the breeder

thickness and reduced coolant thickness (albeit with a larger coolant inner radius).

Through these studies, it was shown that in combination with an evaluation pipeline, Bayesian methods are effective for improving the intended figure of merit in a computationally expensive engineering design problem relevant to fusion breeding technology. Nevertheless, the approach also exhibited some limitations, and several potential improvements were identified:

- improved fidelity of neutronics and thermal hydraulics treatment;
- incorporation of coupled multi-physics;
- incorporation of cost modelling;
- extension of the procedure to Pareto optimisation;
- extension to high-dimensionality Bayesian optimisation techniques.

Notwithstanding those limitations already discussed, this study demonstrates the potential of parametric design optimisation to accelerate the design of fusion breeder technology. With further development, this will empower decision-makers to comprehend a highly complex design space and identify improved configurations more efficiently.

Data availability statement

The datasets presented in this article are not readily available because the LIBRTI Programme involves collaboration with third parties and the data may be commercial. Requests to access the datasets should be directed to luke.humphrey@ukaea.uk.

Author contributions

LH: Formal Analysis, Investigation, Methodology, Software, Visualization, Writing – original draft. HB: Conceptualization, Methodology, Project administration, Software, Supervision, Writing – original draft. SM: Software, Visualization, Writing – review and editing. AD: Resources, Supervision, Writing – review and editing. DF: Project administration, Writing – review and editing.

Funding

The authors declare that financial support was received for the research and/or publication of this article. This work has been funded by the LIBRTI project which is part of the Fusion Futures Programme. As announced by the UK Government in October 2023, Fusion Futures aims to provide holistic support for the development of the fusion sector.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The authors declare that no Generative AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

References

Abdou, M. A., Gierszewski, P. J., Tillack, M. S., Taghavi, K., Kleefeldt, K., Bell, G., et al. (1985). A study of the issues and experiments for fusion nuclear technology. *Fusion Technol.* 8, 2595–2645. doi:10.13182/FST85-A24685

Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. (2019). "Optuna: a next-generation hyperparameter optimization framework," in Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2623–2631. doi:10.1145/3292500.3330701

Balandat, M., Karrer, B., Jiang, D. R., Daulton, S., Letham, B., Wilson, A. G., et al. (2019). BoTorch: programmable bayesian optimization in PyTorch. Corr. *Abs/1910*, 06403. doi:10.48550/arXiv.1910.06403

Berry, R. A., Zou, L., Zhao, H., Zhang, H., Peterson, J. W., Martineau, R. C., et al. (2016). *RELAP-7 theory manual. Tech. rep.* Idaho Falls, ID (United States): Idaho National Lab. INL. doi:10.2172/1262488

Bindel, D., Landreman, M., and Padidar, M. (2023). Understanding trade-offs in stellarator design with multi-objective optimization. *J. Plasma Phys.* 89, 905890503. doi:10.1017/S0022377823000788

Brown, D., Chadwick, M., Capote, R., Kahler, A., Trkov, A., Herman, M., et al. (2018). ENDF/B-VIII.0: the 8th major release of the nuclear reaction data library with CIELO-project cross sections, new standards and thermal scattering data. *Nucl. Data Sheets* 148, 1–142. doi:10.1016/j.nds.2018.02.001

Brown, T., Marsden, S., Gopakumar, V., Terenin, A., Ge, H., and Casson, F. (2023). Multi-objective Bayesian optimisation for design of Pareto-optimal current drive profiles in STEP ArXiv:2310.02669 [physics].

Chapman, R. (2025). Developing integrated cost models for fusion power plants. *J. Fusion Energy* 44, 43. doi:10.1007/s10894-025-00515-1

Coreform LLC, (2024). *Coreform cubit*. Computer software.Version 2024.3. Orem, UT.

Lux, H., Brown, C., Butcher, M., Chapman, R., Foster, J., and Nawal, N. (2024). Optimizing the cost of the STEP programme. *Philosophical Trans. R. Soc. A Math. Phys. Eng. Sci.* 382, 20230413. doi:10.1098/rsta.2023.0413

Cottrell, J., Hughes, T., and Bazilevs, Y. (2009). Isogeometric analysis: toward integration of CAD and FEA. doi:10.1002/9780470749081.ch7

Daulton, S., Balandat, M., and Bakshy, E. (2020). "Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization," in *Advances in neural information processing systems*. Editors H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Red Hook, NY: Curran Associates, Inc.), 33, 9851–9864.

Eriksson, D., and Jankowiak, M. High-dimensional bayesian optimization with sparse axis-aligned subspaces (2021) ArXiv:2103.00349 [cs]. doi:10.48550/arXiv.2103.00349

Eriksson, D., Pearce, M., Gardner, J. R., Turner, R., and Poloczek, M. (2020). Scalable global optimization via local bayesian optimization. ArXiv:1910.01739 [cs]. doi:10.48550/arXiv.1910.01739

Frazier, P. I. (2018). A tutorial on bayesian optimization. ArXiv:1807.02811. doi:10.48550/arXiv.1807.02811

Gaston, D., Newman, C., Hansen, G., and Lebrun-Grandié, D. (2009). MOOSE: a parallel computational framework for coupled systems of nonlinear equations. *Nucl. Eng. Des.* 239, 1768–1778. doi:10.1016/j.nucengdes.2009.05.021

Gilbert, M. R., Foster, D., Bradnam, S. C., Davis, A., Martis, V., and Lavrentiev, M. Y. (2025). "Modelling of a neutron source test bed for the fusion fuel cycle," in International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, New York, NY (American Nuclear Society).

Giudicelli, G., Lindsay, A., Harbour, L., Icenhour, C., Li, M., Hansel, J. E., et al. (2024). 3.0 - MOOSE: enabling massively parallel multiphysics simulations. *SoftwareX* 26, 101690. doi:10.1016/j.softx.2024.101690

Humphrey, L. R., Dubas, A. J., Fletcher, L. C., and Davis, A. (2024). Machine learning techniques for sequential learning engineering design optimisation. *Plasma Phys. Control. Fusion* 66, 025002. doi:10.1088/1361-6587/ad11fb

Jürgen, R., Mayer, W., and van Havre, Y. (2024). Freecad. Available online at: http://www.freecadweb.org.

Kovari, M., Fox, F., Harrington, C., Kembleton, R., Knight, P., Lux, H., et al. (2016). "PROCESS": a systems code for fusion power plants – part 2: engineering. *Fusion Eng. Des.* 104, 9–20. doi:10.1016/j.fusengdes.2016.01.007

Kulcinski, G. L., Radel, R. F., and Davis, A. (2016). "Near term, low cost, 14MeV fusion neutron irradiation facility for testing the viability of fusion structural materials," in Proceedings of the 12th International Symposium on Fusion Nuclear Technology-12 (ISFNT-12), 1072–1076. doi:10.1016/j.fusengdes.2016.01.022

Leys, O., Leys, J. M., and Knitter, R. (2021). Current status and future perspectives of EU ceramic breeder development. *Fusion Eng. Des.* 164, 112171. doi:10.1016/j.fusengdes.2020.112171

Liaw, R., Liang, E., Nishihara, R., Moritz, P., Gonzalez, J. E., and Stoica, I. (2018). Tune: a research platform for distributed model selection and training. *arXiv Preprint arXiv:1807.05118*. doi:10.48550/arXiv.1807.05118

Morris, J., Muldrew, S., Knight, P., Kovari, M., Pearce, A., Kahn, S., et al. (2023). PROCESS. doi:10.5281/zenodo.8338226

Mungale, S. (2024). Hypnos. doi:10.5281/zenodo.14277222

Nogueira, F. (2014). Bayesian optimization: open source constrained global optimization tool for python.

Novak, A., Andrs, D., Shriwise, P., Fang, J., Yuan, H., Shaver, D., et al. (2022). Coupled monte carlo and thermal-fluid modeling of high temperature gas reactors using cardinal. *Ann. Nucl. Energy* 177, 109310. doi:10.1016/j.anucene.2022.109310

Nunn, T., Gopakumar, V., and Kahn, S. (2023). Shaping of magnetic field coils in fusion reactors using bayesian optimisation doi:10.48550/arXiv.2310.01455ArXiv:2310.01455

Open CASCADE SAS, (2019). Open cascade technology. Version 7.4.0.Available online at: https://dev.opencascade.org/.

Packman, S., Riva, N., and Rodriguez-Fernandez, P. (2025). Bayesian methods for magnetic and mechanical optimization of superconducting magnets for fusion. *J. Fusion Energy* 44, 19. doi:10.1007/s10894-025-00486-3

Permann, C. J., Gaston, D. R., Andrš, D., Carlsen, R. W., Kong, F., Lindsay, A. D., et al. (2020). MOOSE: enabling massively parallel multiphysics simulation. *SoftwareX* 11, 100430. doi:10.1016/j.softx.2020.100430

Petersen, H. (1970). The properties of helium: Density, specific heats. Viscosity, and thermal conductivity at pressures from 1 to 100 bar and from room temperature to about 1800 K. *Riso National Laboratory. Denmark. Forskningscenter Risoe. Risoe-R No. 224.*

Romano, P. K., and Forget, B. (2013). The OpenMC monte carlo particle transport code. *Ann. Nucl. Energy* 51, 274–281. doi:10.1016/j.anucene.2012.06.040

Romano, P. K., Horelik, N. E., Herman, B. R., Nelson, A. G., Forget, B., and Smith, K. (2015). "OpenMC: a state-of-the-art monte carlo code for research and development," in Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013 (Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms), Berkley, CA: Elsevier. 90–97. doi:10.1016/j.anucene.2014.07.048

Schoof, L. A., and Yarberry, V. R. (1994). *Exodus II: a finite element data model*. Albuquerque, NM (United States): Sandia National Lab. SNL-NM. doi:10.2172/10102115

Si, M. (1967). On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Comput. Math. Math. Phys.* 7, 86–112. doi:10.1016/0041-5553(67)90144-9

Simon, P. C. A., Icenhour, C. T., Singh, G., Lindsay, A. D., Bhave, C. V., Yang, L., et al. (2025). MOOSE-based tritium migration analysis program, version 8 (TMAP8) for advanced open-source tritium transport and fuel cycle modeling. *Fusion Eng. Des.* 214, 114874. doi:10.1016/j.fusengdes.2025.114874

Tautges, T., Wilson, P., Kraftcheck, J., Smith, B., and Henderson, D. (2009). "Acceleration techniques for direct use of CAD-based geometries in monte carlo radiation transport," in *Proceedings of M&C*.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in python. *Nat. Methods* 17, 261–272. doi:10.1038/s41592-019-0686-2

Wang, Z., Hutter, F., Zoghi, M., Matheson, D., and Freitas, Nd (2016). Bayesian optimization in a billion dimensions via random embeddings, ArXiv:1301.1942. *J. Artif. Intell. Res.* 55, 361–387. doi:10.1613/jair.4806

Williams, C. K., and Rasmussen, C. E. (2006). *Gaussian processes for machine learning*, 2. Cambridge, MA: MIT press.

Zhou, G., Kang, Q., Hernández, F. A., D'Amico, S., and Kiss, B. (2020). Thermal hydraulics activities for consolidating HCPB breeding blanket of the european DEMO. *Nucl. Fusion* 60, 096008. doi:10.1088/1741-4326/ab96f2

Zhou, G., Hernández, F. A., Pereslavtsev, P., Kiss, B., Retheesh, A., Maqueda, L., et al. (2023). The european DEMO helium cooled pebble bed breeding blanket: design status at the conclusion of the pre-concept design phase. *Energies* 16, 5377. doi:10.3390/en16145377