
Shaping of Magnetic Field Coils in Fusion Reactors using Bayesian Optimisation

Timothy Nunn

UK Atomic Energy Authority
Oxford, OX14 3EB
timothy.nunn@ukaea.uk

Vignesh Gopakumar

UK Atomic Energy Authority
Oxford, OX14 3EB
vignesh.gopakumar@ukaea.uk

Sebastien Kahn

UK Atomic Energy Authority
Oxford, OX14 3EB
sebastien.kahn@ukaea.uk

Abstract

Nuclear fusion using magnetic confinement holds promise as a viable method for sustainable energy. However, most fusion devices have been experimental and as we move towards energy reactors, we are entering into a new paradigm of engineering. Curating a design for a fusion reactor is a high-dimensional multi-output optimisation process. Through this work we demonstrate a proof-of-concept of an AI-driven strategy to help explore the design search space and identify optimum parameters. By utilising a Multi-Output Bayesian Optimisation scheme, our strategy is capable of identifying the Pareto front associated with the optimisation of the toroidal field coil shape of a tokamak. The optimisation helps to identify design parameters that would minimise the costs incurred while maximising the plasma stability by way of minimising magnetic ripples.

1 Introduction

A tokamak is plasma magnetic confinement system laid out in a toroidal structure [1]. The goal is to confine the plasma at sufficiently high temperature to generate fusion reactions, attempting to generate electricity. To efficiently design a tokamak, a wide variety of physics and engineering challenges must be resolved within a limited amount of space. Hence, the process governing such design often involves various simulation codes and multiple stakeholders, often drawn out for months, while only evaluating a limited section of the design space [2]. The tokamak design code Bluemira, resulting from the merge between the BLUEPRINT [3] and the MIRA [4] codes, integrates various aspects of a fusion power plant design in the same framework, significantly reducing the design optimization time, allowing a wider design exploration. A major component of a tokamak is its toroidal field (TF) coils, responsible for the generation of the magnetic pressure confining the plasma. We invite the reader to view Appendix A for a schematic diagram of a Tokamak including the typical location of the TF coils. This work proposes a first use of an AI-based exploration and identification of optimum fusion reactor component design. We deploy a Multi-Objective Bayesian Optimisation algorithm (MOBO) to optimise the shape of the TF coils. For a further reading about TF coil optimization, we invite the reader to read the following thesis [5]. The MOBO is deployed over Bluemira and benchmarked against conventional optimizers to demonstrate the efficacy of our approach, laying the foundation to extend this proof-of-concept to actively search more complex design spaces.

2 Problem

The TF coil shape influences various aspects of the tokamak design, from the distribution of the electromechanical forces (structural design) to the size of the reactor (cost) and the toroidal field variations (ripple) at the plasma boundaries driving its performance [6, 7]. The problem considers optimising across the parameterisation of the 3D shape of the TF coils conductors (number, thickness, height, radii) as outlined in appendix B and illustrated in figure 4. The considered coil conductors have a rounded rectangular shape, while its cross-section is taken to be rectangular for simplicity. The optimisation is performed with the objective metrics summarised in table 1, that outline the volume of the magnets and its induced ripple effects. The magnetic field is calculated using an analytical 3D magnetostatic formulation, allowing one to quickly and accurately calculate the magnetic field at an arbitrary point in space, using a series of current filaments to approximate the conductor geometry of the TF coils. For a given point in the poloidal plane, the ripple is defined as the relative difference between the maximum magnetic field (on the TF coil leg) and the minimum magnetic field (between two TF coil legs) as described in [5, p. 46]. The coil size is simply defined as the perimeter length multiplied by the conductor cross-section. The plasma shape is fixed for this experiment as defined by the geometric parameterisation described in [8, p. 310-312]. The number of points at which the ripple is evaluated along the flux surface (ripple points) is fixed, ensuring that ripple evaluations occur at the same point on the plasma for each design exploration.

3 Optimisation

The optimisation pipeline was built using the BoTorch [9] framework integrated with GpyTorch[10]. The implementation of the design search space exploration using a Bayesian framework was done using the Ax API [11] built on top of BoTorch and GpyTorch.

3.1 Setup

In order to optimise the shape of TF coils using a Bayesian strategy, the problem is framed as a multi-objective optimisation scenario. The optimisation requires searching an 8-dimensional parameter hyperspace to find ideal points on the Pareto front that would minimise the magnetic ripple effects as well as the volumetric size of the TF coils as shown in Table 1. The parameters to be optimised for and their domain ranges can be found in Appendix B. Within the Bayesian optimisation process we have chosen to avoid adding additional constraints and have instead instilled the geometric constraints in the parameter bounds.

Metric	Description
Ripple	Maximum value of magnetic ripple encountered while evaluating the <i>ripple</i> function
Size	Volumetric sum of the TF Coils.

Table 1: The metric name and description for each of the metrics being minimised.

3.2 Implementation

The optimisation is split into two portions: initial quasi-random sampling within the domain followed by multi-objective Bayesian optimisation (MOBO) across the design space. The initial sampling was done using the Sobol [12] quasi-random sampling strategy. The MOBO loop will run for a prescribed number of iterations. Each iteration, a series of Gaussian Process (GP) surrogates are trained on all the currently observed data [13]. We construct an individual GP for each objective using Exact GPs, as the ripple function from Bluemira is treated to be as ground truth without any noise. Each GP is defined using a Gaussian likelihood and a prior with constant mean and a Matern Kernel [13] using a length scale prior. GPs are fitted by minimising the sum marginal log likelihood across them. Inputs to the GPs are each mapped independently onto the unit hypercube such that any design point $X \in [0, 1]^n$ where this problem has n inputs. Outputs Y are standardised to have a mean 0 and variance 1. Within the Bayesian framework, newer design points are sampled using the expected hypervolume improvement (EHVI) [14] acquisition function across the fitted GPs, in line with the mentioned objectives.

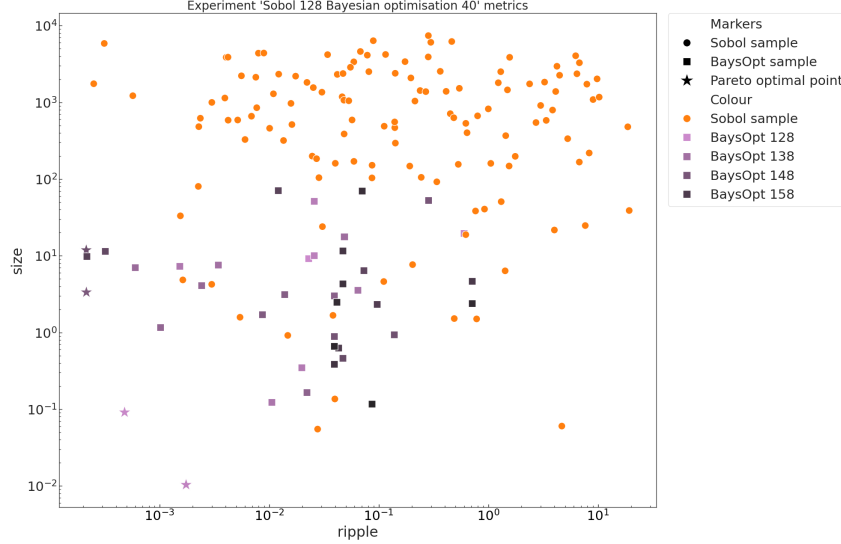


Figure 1: Log-log plot showing the resulting metrics of design points sampled in the Sobol stage (orange dots) and the Bayesian optimisation stage (purple squares) where the darker the colour of a Bayesian point, the later the iteration that design point corresponds to. Stars indicate the non-dominated points—the Pareto optimal design points.

4 Results

4.1 Comparison of designs

Because we are minimising both *size* and *ripple* of the design, there is no objective best design point. Instead we get a Pareto frontier of design points. The feasible design points are ordered by *size* and, generally, the design with the smallest coil is selected.

Expt #	Sobol Samples	Bayesian Samples	Size	Ripple
1	64	20	0.46	0.018
2	128	40	0.091	4.5e-4

Table 2: Comparison of the optimal design solution(s) achieved on different experiments utilising the Sobol + Bayesian optimisation method.

The Gaussian process is able to understand the general profile of the two metrics relatively quickly. Some GP fits associated with each objective can be seen in Appendix E. As seen in Table 2, experiment 2 finds better quality points, by several orders of magnitude, than experiment 1. Generally, as more samples of the *ripple* function are evaluated, points closer to the “true” Pareto front are found.

Figure 1 shows the *size* and *ripple* values observed in experiment 2, it can be seen how a majority of the Sobol samples are poor quality design points. The Bayesian optimisation is immediately able to identify regional proximity of the the Pareto front, and descends towards it.

The results in Table 2 and Figure 1 are in line with the physical description of the problem as stated in section 2, and the identified design points are well within the expected range. By definition, a larger TF coil has a lower ripple, and vice versa. The parameterisations for experiment 2 and the corresponding designs are discussed further in Appendix D.

4.2 Comparison with other methods

We benchmarked the Bayesian optimisation over a Gaussian Process surrogate approach against a purely quasi-random sampling across the search domain using Sobol, and a multi-objective evolutionary algorithm, specifically the Non-dominated Sorting Genetic Algorithm (NSGA-II) [15]. The NSGA-II is the de-facto optimisation toolkit used within Bluemira and is often the primary choice

for determining design choices. We will compare each optimisation method with respect to a baseline experiment, experiment 2 in Table 2. The comparison will be between the quality of the best Pareto optimal point(s) and the time taken to find said points for a fixed number of design point evaluations (168). All tests were run on a single core of an Intel(R) Xeon(R) Silver 4108 CPU @ 1.80GHz with 4GB of memory available.

Method	size	ripple	time (s)
Sobol + Bayesian optimisation (baseline)	0.012	4.5e-4	872
Sobol	0.055, 0.43	0.028, 5.9e-3	691
NSGA-II	89.1	0.19	786

Table 3: Comparison of the optimal point(s) for several multi-objective optimisation methods after allowing 168 *ripple* function evaluations. Where multiple Pareto optimal points are deemed to be of the same quality, they are presented comma delimited.

The Sobol samples in Table 3 are decent quality design points, but do not recover the Pareto frontier accurately. The best Sobol design points are still orders of magnitude larger than the best point in the baseline experiment. We can qualitatively verify this via Figure 1 where the Pareto frontier comprises only MOBO samples. Further, only a few of the Sobol samples are better than even the worst MOBO samples, with a majority of the Sobol samples orders of magnitude larger the MOBO samples. Overall, the best Sobol design points are of similar magnitude and quality to experiment 1 in Table 2, which required half the number of function evaluations.

We conclude that the 23% increase in the run-time of the baseline MOBO with respect to the Sobol is acceptable, considering the significant improvement in design quality. The difference in time between Sobol and the baseline in Table 3 can be attributed to the overhead of fitting the Gaussian Process and optimisation of the acquisition function. In order to achieve comparable design quality to the baseline experiment, more Sobol samples would need to be taken. Therefore, there exists a point where the overhead of fitting the GP will be equal to, or less than, the time taken to evaluate the additional Sobol samples required to achieve the same design quality. This situation becomes more likely if the cost of evaluating the *ripple* is increased by increasing the fidelity of the model and evaluating more ripple points on the plasma surface. We therefore find that reducing the number of function evaluations is significantly more important than the minor overhead of training.

NSGA-II significantly under performs compared to both the baseline and Sobol. This is likely due to the curse of dimensionality hampering the evolutionary nature of NSGA-II. In order to efficiently explore the search space for the genetic algorithm, we would require a lot more iterations and evaluations before obtaining comparable values. The method used just 8 Monte Carlo points with 20 evaluations, rendering the optimiser unable to efficiently minimise across both the objectives.

5 Conclusion and Future Work

Through this work, we demonstrate a proof-of-concept for a multi-objective Bayesian optimisation strategy deployed for determining better design strategies for a fusion reactor. Our work demonstrates that through MOBO we are able to better capture the Pareto front associated with the optimisation efficiently. The performance of the MOBO is further elucidated by benchmarking against a quasi-random exploration scheme (Sobol) and an evolutionary algorithm (NSGA-II). Though the optimised parameters lead to unrealistic design parameters for actual reactor construction, they accurately represent the ideal parameters as prescribed by Bluemira for the problem setup that we have provided, validating our optimisation strategy. As the next phase of our research, we would be extending the optimisation problem to include the magnetic equilibrium generated by the currents within the TF coils, optimising across both design and control elements of the reactor.

References

- [1] J. Wesson and D.J. Campbell. *Tokamaks*. International series of monographs on physics. Clarendon Press, 2004.
- [2] G. Federici, C. Bachmann, L. Barucca, W. Biel, L. Boccaccini, R. Brown, C. Bustreo, S. Ciattaglia, F. Cismonti, M. Coleman, V. Corato, C. Day, E. Diegele, U. Fischer, T. Franke, C. Gliss, A. Ibarra, R. Kembleton, A. Loving, F. Maviglia, B. Meszaros, G. Pintsuk, N. Taylor, M.Q. Tran, C. Vorpahl, R. Wenninger, and J.H. You. Demo design activity in europe: Progress and updates. *Fusion Engineering and Design*, 136:729–741, 2018. Special Issue: Proceedings of the 13th International Symposium on Fusion Nuclear Technology (ISFNT-13).
- [3] M. Coleman and S. McIntosh. Blueprint: A novel approach to fusion reactor design. *Fusion Engineering and Design*, 139:26–38, 2019.
- [4] F. Franza, L.V. Boccaccini, E. Fable, I. Landman, I.A. Maione, S. Petschanyi, R. Stieglitz, and H. Zohm. MIRA: a multi-physics approach to designing a fusion power plant. *Nuclear Fusion*, 62(7):076042, jun 2022.
- [5] M. Coleman. *An integrated design framework for future nuclear fusion power reactors*. PhD thesis, University of Cambridge, September 2021.
- [6] N A Uckan, K T Tsang, and J D Callen. Toroidal field ripple effects in large tokamaks. 1 1975.
- [7] J. Narl Davidson. Effect of toroidal field ripple on particle and energy transport in a tokamak. *Nuclear Fusion*, 16(5):731–742, nov 1976.
- [8] Johnner Jean. Helios: A zero-dimensional tool for next step and reactor studies. *Fusion Science and Technology*, 59(2):308–349, 2011.
- [9] Maximilian Balandat, Brian Karrer, Daniel R. Jiang, Samuel Daulton, Benjamin Letham, Andrew Gordon Wilson, and Eytan Bakshy. Botorch: A framework for efficient monte-carlo bayesian optimization. 2019.
- [10] Jacob R. Gardner, Geoff Pleiss, David Bindel, Kilian Q. Weinberger, and Andrew Gordon Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, page 7587–7597, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [11] Eytan Bakshy, Lili Dworkin, Brian Karrer, Konstantin Kashin, Benjamin Letham, Ashwin Murthy, and Shaun Singh. Ae: A domain-agnostic platform for adaptive experimentation. 2018.
- [12] I.M Sobol’. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, 1967.
- [13] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning series. MIT Press, 2005.
- [14] M.T.M. Emmerich, K.C. Giannakoglou, and B. Naujoks. Single- and multiobjective evolutionary optimization assisted by gaussian random field metamodels. *IEEE Transactions on Evolutionary Computation*, 10(4):421–439, 2006.
- [15] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [16] Shunjie Li, H. Jiang, Zhigang Ren, and C. Xu. Optimal tracking for a divergent-type parabolic pde system in current profile control. *Abstract and Applied Analysis*, 2014:1–8, 06 2014. CC BY 4.0 <https://creativecommons.org/licenses/by/4.0>, via Wikimedia Commons.

A Tokamak Illustration

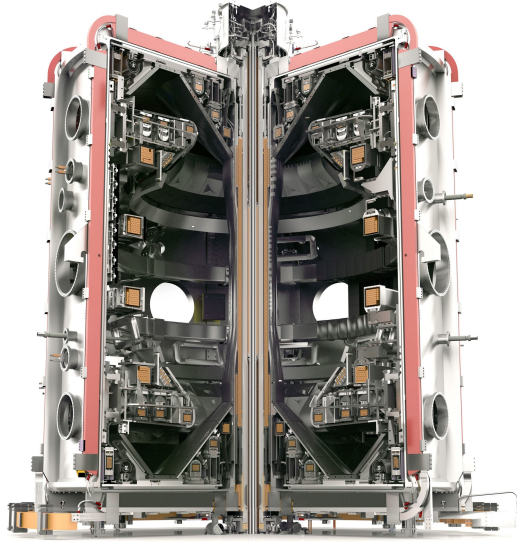


Figure 2: A cutaway illustration of MAST Upgrade, one of the Tokamaks at Culham Science Centre. The Toroidal Field coils (red) have a picture frame shape. Credit to UK Atomic Energy Authority.

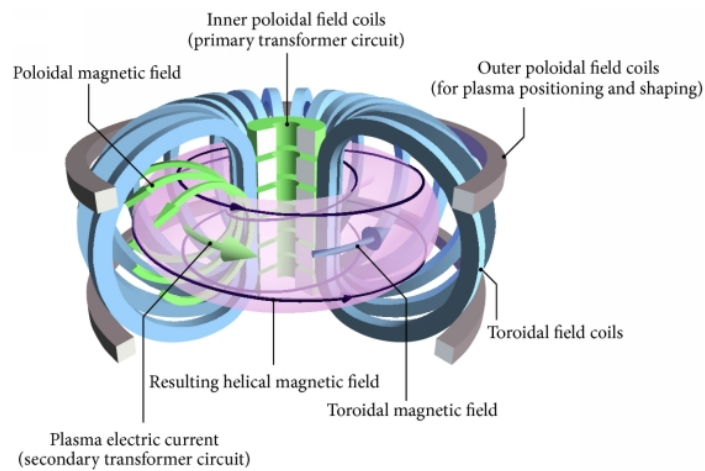


Figure 3: [16] A simple schematic diagram of a generic Tokamak with all of the main magnetic components and fields shown. The TF coil magnets (blue) are that which this work aimed to optimise for.

The vacuum vessel is the empty torus-shaped area at the center of the device where the plasma is magnetically confined. The toroidal magnetic field is produced by the Toroidal Field coils and is partially responsible for magnetic confinement of the plasma.

B Parameter description and ranges

This appendix summarises the optimization variables and their ranges listed in Table 4. The parameters setting the TF coil dimensions, $\{x_1, x_2, z_1, z_2\}$ are illustrated with the green arrows in figure 4, the corner rounding are parametrized by the light blue arrows $\{r_i, r_o\}$. The conductor cross-section is shown as the blue area in figure 5 showing the inboard mid-plane cross-section of a TF coil with a typical support structure (steel casing, grey). The rectangular cross-section is centred on the coil illustrated by the blue line in figure 4. Its radial width is taken as ΔR_{WP} , while its thickness in the toroidal direction $\Delta R_{\text{t,WP}}^{\text{int}}$ is parametrized with $f_{\Delta y_{\text{wp}}}$ defined as

$$f_{\Delta y_{\text{wp}}} = \frac{\Delta R_{\text{t,WP}}^{\text{int}}}{\left(x_1 - \frac{\Delta R_{\text{t,WP}}^{\text{int}}}{2}\right) \tan\left(\frac{\pi}{N_{\text{TF}}}\right)}$$

to avoid overlapping conductors in the inboard regions¹.

Input	Description	range
N_{TF}	Number of TF coils	[16, 24]
x_1	Inboard mid-plane TF coil centreline radius	[0.0, -4.5] m
x_2	Outboard mid-plane TF coil centreline radius	[13.0, 20.0] m
z_1	Top TF coil centreline height	[6.0, 15.0] m
z_2	Bottom TF coil centreline height	[-15, -6.0] m
r_{out}	Top/bottom outboard corners bending radii	[0.0, 5.0] m
r_{in}	Top/bottom inboard corners bending radii	[10 ⁻⁵ , 0.14] m
ΔR_{WP}	Thickness of the conductor (winding pack) in the poloidal plane	[0.1, 0.7] m
$f_{\Delta y_{\text{wp}}}$	Thickness of the conductor (winding pack) in the toroidal direction as a fraction of the maximum allowable value	[0.1, 1.0]

Table 4: Description of the input parameters for the optimization.

¹If $f_{\Delta y_{\text{wp}}}$ is below 1, no overlap happens by definition.

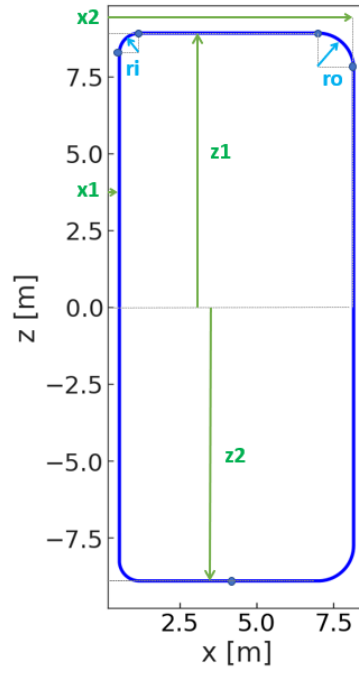


Figure 4: Illustration of the TF coil centreline shape parametrization.

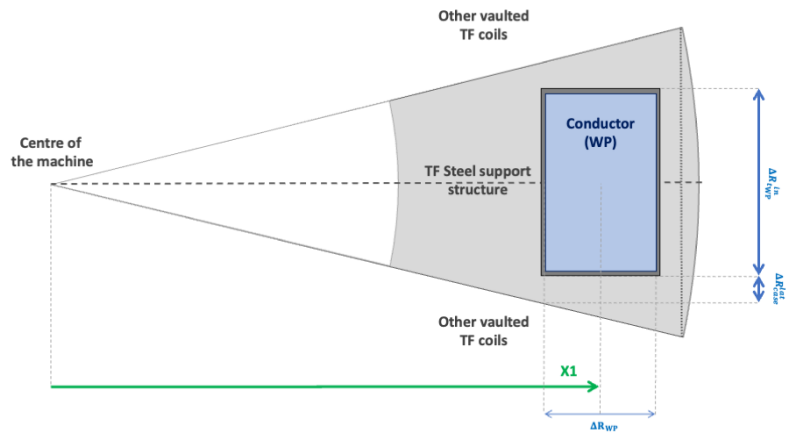


Figure 5: Illustration of the TF coil centreline shape parametrization.

C Additional Pareto frontiers

This appendix includes various Pareto frontier plots for experiments in Tables 2 & 3.

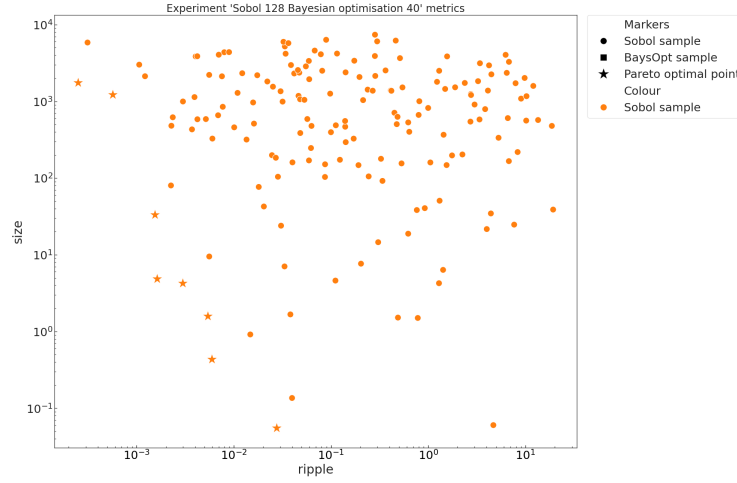


Figure 6: shows the Pareto frontier of the Sobol method from Table 3. When compared with Figure 1 it can be seen that the Pareto frontier exists several orders of magnitude above where it does for the baseline in Table 3

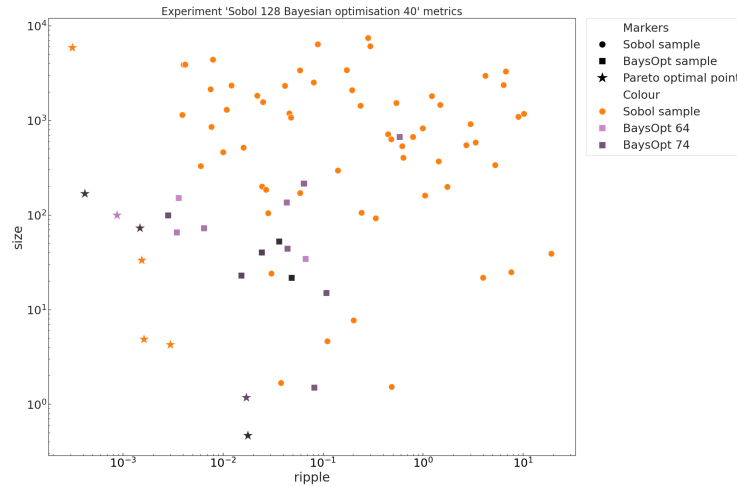


Figure 7: shows the Pareto frontier and distribution of metrics for experiment 1 of Table 2. We can see that the frontier is similar to that shown in 6 despite this frontier needing half the number of function evaluations to reach.

D Experiment 'Sobol 128 + Bayesian optimisation 40' optimal design point

Below we briefly discuss the design point deemed to be of the best quality of the four Pareto optimal design points in experiment 2 in Table 2, sample 129 of 168. We include the parameter values of the design point and include a picture of the resulting coil.

Parameter	Value
dr wp	0.1m
f y wp	0.659
number of coils	25
x1	0.0348m
x2	19.35m
z	15.0m
ri	0.0931m
ro	5.0m

Table 5: parameterisation of design point 19 for experiment 2

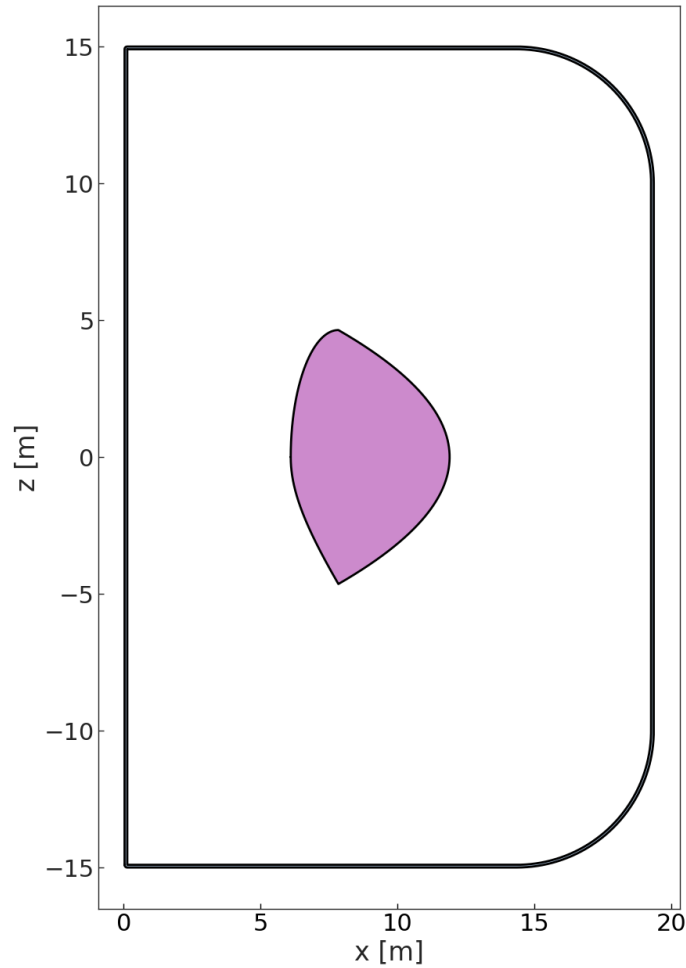


Figure 8: resulting TF coil (blue) from the parameterisation in Table 5 with the fixed plasma shape (purple) shown for reference.

This coil is rather large, ensuring that the distance $x_2 - x_1$ is maximised as this reduces the *ripple* the most. This means x_1 is close to its lower bound, and x_2 is close to its upper bound. The algorithm

manages to achieve a small volumetric size of the TF coil by making the winding pack extremely thin to the point where this coil would likely be physically unstable. Future work, discussed in Section 5, will aim to resolve this by enforcing constraints on stresses experienced by the coil structure.

E Experiment 'Sobol 128 + Bayesian optimisation 40' Gaussian Process Fits

Predictions for a 1-d slice of the parameter space

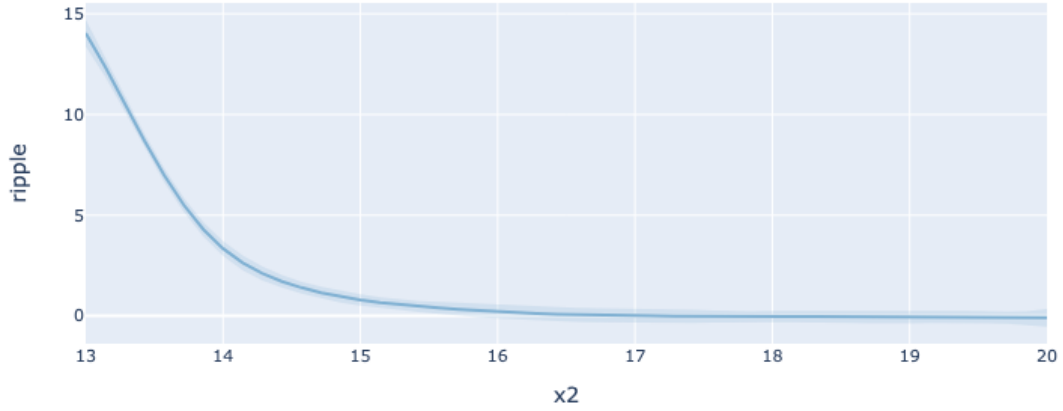


Figure 9: shows the fit of the Gaussian Process for parameter x_2 on the *ripple* metric given that the other variables have been fixed at their midpoints.

Predictions for a 1-d slice of the parameter space

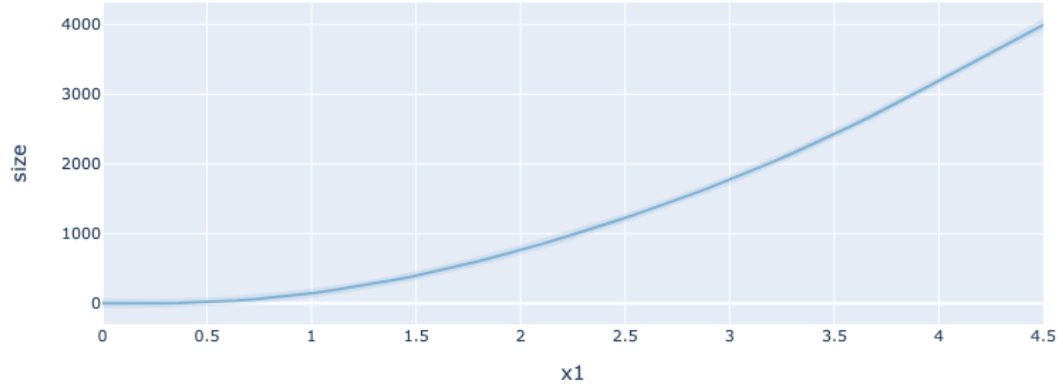


Figure 10: shows the fit of the Gaussian Process for parameter x_1 on the *size* metric given that the other variables have been fixed at their midpoints.

Figures 9 & 10 shows two of the Gaussian Process fits for x_2 on the *ripple* metric and for x_1 on the *size* metric, respectively. We can see that the Gaussian Process' are extremely confident in both metrics given these two physically relevant parameters.