

Direct stellarator coil design using global optimization: application to a comprehensive exploration of quasi-axisymmetric devices

Andrew Giuliani  

Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

(Received 29 October 2023; revised 8 March 2024; accepted 12 March 2024)

Many stellarator coil design problems are plagued by multiple minima, where the locally optimal coil sets can sometimes vary substantially in performance. As a result, solving a coil design problem a single time with a local optimization algorithm is usually insufficient and better optima likely do exist. To address this problem, we propose a global optimization algorithm for the design of stellarator coils and outline how to apply box constraints to the physical positions of the coils. The algorithm has a global exploration phase that searches for interesting regions of design space and is followed by three local optimization algorithms that search in these interesting regions (a ‘global-to-local’ approach). The first local algorithm (phase I), following the globalization phase, is based on near-axis expansions and finds stellarator coils that optimize for quasisymmetry in the neighbourhood of a magnetic axis. The second local algorithm (phase II) takes these coil sets and optimizes them for nested flux surfaces and quasisymmetry on a toroidal volume. The final local algorithm (phase III) polishes these configurations for an accurate approximation of quasisymmetry. Using our global algorithm, we study the trade-off between coil length, aspect ratio, rotational transform and quality of quasi-axisymmetry. The database of stellarators, which comprises approximately 200 000 coil sets, is available online and is called QUASR, for ‘quasi-symmetric stellarator repository’.

Key words: plasma confinement, plasma devices

1. Introduction

Stellarator coil design is typically framed as an optimization problem, where the objective function targets charged particle confinement, other physics properties and engineering requirements on the electromagnetic coils. Coil design problems can admit multiple local minima (Miner *et al.* 2001; Zhu *et al.* 2017) with large variability in performance of the discovered designs (Wechsung *et al.* 2022). One remedy of this problem is to use stochastic optimization (Glas *et al.* 2022; Wechsung *et al.* 2022), which appears to reduce the variability of the discovered minima. Stochastic optimization of coils has also been explored in Lobsien, Drevlak & Pedersen (2018); Lobsien *et al.* (2020). Another approach is to use global optimization algorithms that attempt to fully explore the range of allowable stellarator coil designs. Glas *et al.* (2022) combine both remedies by

† Email address for correspondence: agiuliani@flatironinstitute.org




FIGURE 1. Current workflow wraps phase I (the near-axis expansion algorithm, Giuliani *et al.* 2022a) in globalization subroutines (Eriksson *et al.* 2019) to generate interesting initial guesses for the subsequent volume QA optimizations in phase II (Giuliani *et al.* 2022b) and III (Giuliani *et al.* 2023).

searching for the global optimum of a stochastic objective function. Genetic algorithms have also been explored in Miner *et al.* (2001).

In this work, we propose a novel technique for applying global optimization algorithms to deterministic coil design problems plagued by many local minima. Our coil design workflow is decomposed into three phases, wrapped in a globalization procedure (figure 1) to automatically design a large database of vacuum field stellarators for various design targets. During phase I, the first algorithm finds initial coil sets with nested magnetic surfaces in the neighbourhood of a magnetic axis by using the near-axis expansion formalism (Giuliani *et al.* 2022a). During phase II, the second algorithm takes these coil sets and expands the region of nested flux surfaces with a good approximation of quasi-axisymmetry (QA) (Giuliani *et al.* 2023). Finally, during phase III, the third algorithm polishes these coil sets for precise QA (Giuliani *et al.* 2022b).

We compare two approaches to globalization. The first, somewhat naive approach attempts to find a global minimum of the objective by perturbing initial guesses provided to a local optimization algorithm. The disadvantage is that, if the perturbation is too large, then the optimizer might be sent to an uninteresting region of parameter space. Conversely, if the perturbation is too small, the optimizer might not sufficiently explore the design space. Despite this downside, it can yield good results with some tuning (Wechsung *et al.* 2022). A second, much less *ad hoc* approach, relies on the global optimization algorithm called trust region Bayesian optimization (TuRBO) (Eriksson *et al.* 2019). For this algorithm, the user must provide lower and upper bounds on the design variables, known as box constraints. This algorithm has been used before in coil design problems (Glas *et al.* 2022), but our approach is notably different as we reformulate our optimization problem to accept constraints on the geometry of the coils in physical, rather than Fourier space.

Since we are generating a large data set of stellarators, we also study the trade-offs between multiple competing design targets, which is the aim of multi-objective optimization. Recently, multi-objective optimization was applied to stellarator coil design (Bindel, Landreman & Padidar 2023), where a first-order continuation algorithm to construct a local Pareto front was developed. The algorithm was based on a Taylor expansion of the optimality condition that points on the Pareto front satisfy. In this work, we do not use this continuation approach but nevertheless study the trade-offs in stellarator coil design using our database.

To summarize our contributions, we propose a simple way to incorporate box constraints on the physical positions of stellarator coils so that global optimization algorithms can be applied to the first phase of the coil design procedure. Following this first phase are two volume QA algorithms which target nested flux surfaces and precise QA. Scanning over various physics design targets, we have compiled a comprehensive database of approximately 200 000 stellarator devices, and examined the trade-off between competing

design objectives (quality of quasisymmetry, total coil length, rotational transform and device aspect ratio). A subset of the devices have comparable quasisymmetry to the highly optimized configurations in Giuliani *et al.* (2023).

We would also like to point out some limitations of this work. First, the globalization is only applied to the near-axis coil design algorithm, and not to the volume QA phases of coil design. Because of this, we may miss out on performant stellarators should precise first-order near-axis QA in phase I not correlate with precise volume QA in phases II and III. Second, this database only contains curl-free stellarators with optimized QA, although our algorithms are generic and apply to other flavours of quasisymmetry. These vacuum-field devices are can also useful as they can be used as initial guesses during an optimization as plasma pressure is progressively increased (Boozer 2019). Third, there may be duplicate devices in the database and possible mechanisms for this are outlined in § 5. Finally, there are specialized algorithms for visualizing the Pareto front, which we do not use here.

2. Phase I: stellarators optimized for quasisymmetry on the magnetic axis

The goal of the first optimization problem in phase I is to find an initial set of coils that has nested flux surfaces and produces a good approximation of quasisymmetry, at least locally to the magnetic axis. To this end, we recall here the optimization problem in Giuliani *et al.* (2022a) that computes coils based on near-axis expansions

$$\min_{c, I, a, \sigma, \bar{\eta}, \iota_a} \hat{f}_{\text{axis}}(c, I, a, \sigma, \bar{\eta}, \iota_a), \tag{2.1a}$$

$$\text{subject to } \mathbf{g}(a, \sigma, \bar{\eta}, \iota_a) = 0 \text{ on axis,} \tag{2.1b}$$

$$\iota_a = \iota_{\text{target}}, \tag{2.1c}$$

$$L_i = L_{\text{target}}, \tag{2.1d}$$

$$\kappa_i \leq \kappa_{\text{max}}, \quad i = 1, \dots, N_c, \tag{2.1e}$$

$$\frac{1}{L_i} \int_{\Gamma^{(i)}} \kappa_i^2 dl \leq \kappa_{\text{msc}}, \quad i = 1, \dots, N_c, \tag{2.1f}$$

$$\|\mathbf{\Gamma}^{(i)} - \mathbf{\Gamma}^{(\text{axis})}\| \geq d_{\text{min}}, \tag{2.1g}$$

$$\|\mathbf{\Gamma}^{(i)} - \mathbf{\Gamma}^{(j)}\| \geq d_{\text{min}}, \quad \text{for } i \neq j, \tag{2.1h}$$

$$\|\mathbf{\Gamma}^{(i)}\| - L_i = 0, \quad \text{for } i = 1, \dots, N_c, \tag{2.1i}$$

where

$$\begin{aligned} \hat{f}_{\text{axis}}(c, I, a, \bar{\eta}, \sigma, \iota_a) = & \int_{\text{axis}} (\|\mathbf{B}(c, I, a)\| - B_{\text{QS}}(a))^2 dl \\ & + \int_{\text{axis}} \|\nabla \mathbf{B}(c, I, a) - \nabla B_{\text{QS}}(a, \bar{\eta}, \sigma, \iota_a)\|^2 dl, \end{aligned} \tag{2.2}$$

and \mathbf{B} , $\nabla \mathbf{B}$ are the magnetic field and its gradient generated by the coils, B_{QS} , ∇B_{QS} are the magnetic field and its gradient, which is quasisymmetric to first order for the given magnetic axis shape; $\mathbf{\Gamma}^{(\text{axis})}$, $\mathbf{\Gamma}^{(i)} \in \mathbb{R}^3$ are respectively the vector of the magnetic axis and i th coil's Cartesian coordinates. The vector of coil degrees of freedom are stored in $c \in \mathbb{R}^{3N_c(2N_{f,c}+1)}$. The x, y, z positions of the N_c base modular coils are each represented using a Fourier series with $N_{f,c}$ modes (Zhu *et al.* 2017). The electromagnetic coils also have an associated current stored in the vector $I \in \mathbb{R}^{N_c}$. The full device is obtained by applying

discrete rotational symmetry with n_{fp} field periods and stellarator symmetry to N_c base modular coils. The length of the i th coil is $L_i = \int_0^1 \|\mathbf{r}^{(i)}(\theta)\| d\theta$. The stellarator symmetric magnetic axis is represented using a Fourier series of the R and Z positions of the magnetic axis, represented in cylindrical coordinates. In total, this results in $2N_{f,a} + 1$ degrees of freedom associated with magnetic axis, stored in $\mathbf{a} \in \mathbb{R}^{2N_{f,a}}$, where the dimensionality is one less than the number of Fourier coefficients of the axis since we require that the mean radius of the magnetic axis be 1 m. Finally, $\iota_a \in \mathbb{R}$ is the on-axis rotational transform, and $\bar{\eta} \in \mathbb{R} \setminus \{0\}$ is a scalar that describes how much that magnetic field strength varies on surfaces in the neighbourhood of the axis. Given \mathbf{a} and $\bar{\eta}$, the vector $\boldsymbol{\sigma} \in \mathbb{R}^{2N_{f,a}}$ and on-axis rotational transform ι_a are fully determined via the constraint \mathbf{g} in (2.1b), which is the periodic Riccati equation

$$\frac{d\boldsymbol{\sigma}}{d\varphi} + \iota \left[\frac{\bar{\eta}}{\kappa^4} + 1 + \boldsymbol{\sigma}^2 \right] + 2\tau \frac{G_0 \bar{\eta}^2}{B_0 \kappa^2} = 0, \quad (2.3)$$

discretized using a Fourier collocation method (Giuliani *et al.* 2022a), where κ, τ are respectively the curvature and torsion on the magnetic axis and B_0 is the field strength on axis. Other variables in (2.3) such as φ, σ, G_0 are defined in Landreman & Sengupta (2018). The data $\iota_a, \bar{\eta}, \boldsymbol{\sigma}$ fully define the gradient of the target quasi-symmetric magnetic field on axis. After solving (2.3), then we can evaluate the near-axis gradient $\nabla \mathbf{B}_{QS}$.

Using the implicit function theorem, we minimize the reduced objective

$$f_{\text{axis}}(\mathbf{c}, \mathbf{I}, \mathbf{a}, \bar{\eta}) = \hat{f}_{\text{axis}}(\mathbf{c}, \mathbf{I}, \mathbf{a}, \bar{\eta}, \boldsymbol{\sigma}(\mathbf{a}, \bar{\eta}), \iota_a(\mathbf{a}, \bar{\eta})), \quad (2.4)$$

by eliminating $\boldsymbol{\sigma}, \iota_a$ via the constraint (2.1b). Thus, this constraint is satisfied exactly. The constraint in (2.1c) ensures that a stellarator is found with the target on-axis rotational transform ι_{target} . The constraint in (2.1d) ensures that each electromagnetic coil has the same length L_{target} . The constraints in (2.1e) and (2.1f) ensure that the maximum curvature and mean squared curvature do not exceed κ_{max} and κ_{msc} , respectively. The constraints in (2.1g), (2.1h) prevent the inter-coil and coil-to-axis distances from decreasing below d_{min} . Finally, the constraint in (2.1i) ensures that the coil incremental arclength stays uniform. Constraints (2.1c)–(2.1i) are enforced using a penalty method, and are satisfied to 0.1% precision. For each of these constraints, a quadratic penalty term appears in the objective function. For example, the quadratic penalty associated with the equality constraint (2.1c) is $\frac{1}{2}(\iota_a - \iota_{\text{target}})^2$. To ensure that the constraints are satisfied accurately enough, the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm is run multiple times with a fixed computational budget. At the end of each BFGS run, the weight in front of the quadratic penalty was increased by a factor of 10 if the constraint is violated by more than 0.1%.

A local minimum of (2.1) is a set of electromagnetic coils that produce a magnetic field that is close to the quasisymmetric magnetic field on a magnetic axis. This optimization problem was solved using local quasi-Newton optimization methods in Giuliani *et al.* (2022a) and Wechsung *et al.* (2022), where the analytical gradient of the reduced objective f_{axis} was obtained using both forward sensitivities and adjoint approaches. We always use the BFGS quasi-Newton method in phase I of the workflow.

In the next two sections, we both illustrate the need for globalization when solving (2.1) and outline two different techniques to accomplish this.

2.1. Naive globalization of phase I

A first naive approach to globalization is to perturb the geometry and currents of flat electromagnetic coils, as well as the parameters of the magnetic axis to generate a set

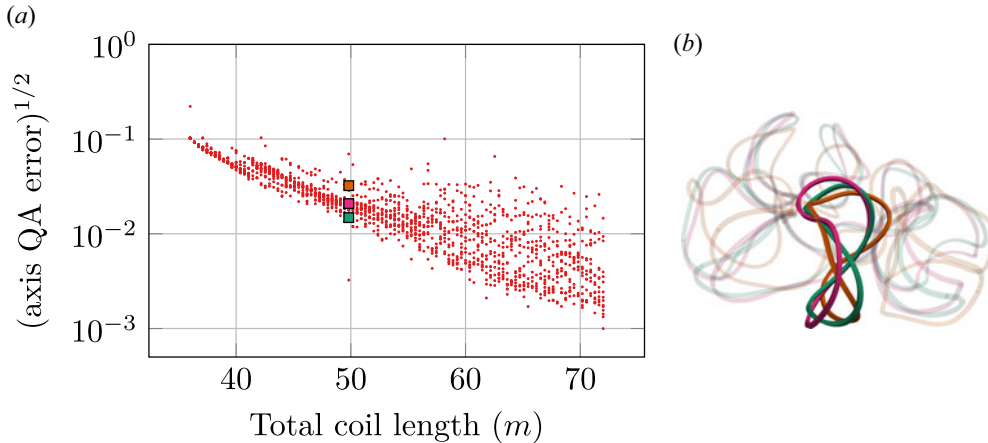


FIGURE 2. (a) We plot the on-axis quasisymmetry error with respect to the total coil length used, where we observe that there are multiple minima. (b) We plot four distinct local minima when the total coil length is 49.81 m, where the colours of the coils correspond to those of the squares on the left. The coil sets' associated magnetic axes are not plotted to avoid clutter in the image.

of initial starting points. These starting points are used to initialize phase I, where the local gradient-based optimization algorithm BFGS is used to find a local minimum of the objective in the neighbourhood of the perturbed initial guess. The physics and engineering quantities that we target are $\iota_{\text{target}} = 0.9$ using n_{fp} , $n_{\text{coils per hp}} = 2$, $d_{\text{min}} = 0.1$ m, $\kappa_{\text{max}} = 5 \text{ m}^{-1}$, $\kappa_{\text{msc}} = 5 \text{ m}^{-2}$, the target coil length L_{target} is varied between 40 and 80 m. For each value of coil length, we solve (2.1) with distinct initial guesses 16 times. The initial guesses are obtained by perturbing the Fourier coefficients of initially flat coils with normally distributed noise with zero mean and standard deviation ϵ . One has substantial freedom to decide many Fourier coefficients to perturb, as well as the standard deviation of the perturbation and how it decays. If ϵ is too large then the local optimizer might be sent into an uninteresting region of coil parameter space. If ϵ is too small, and it might not fully explore the set of feasible coil designs. For the experiment here, we perturb the currents, $\bar{\eta}$, and only the first two Fourier harmonics of the coils, and magnetic axis. Extensive (but tedious and computationally intensive) tuning revealed that the standard deviation of the noise that discovered devices with lowest quasi-symmetry error was $\epsilon = 0.01$ m. We find that the quality of the minima found depends strongly on the choice of ϵ , the number of Fourier modes that are perturbed, and that there is no way to know *a priori* that this value works well for other stellarator designs.

In figure 2, we plot the trade-off between the on-axis quasisymmetry error and total coil length. We find, as expected, that lower quasisymmetry errors are attainable with longer coils. It is clear that making no attempt at globalization to solve this optimization problem might result in a sub-optimal coil set. The problem of multiple minima becomes more pronounced for devices with longer coil lengths because there are more, intricate ways for the coils to arrange themselves. Finally, at around 50 m, there is a device that appears as an outlier with remarkably low near-axis quasisymmetry error. This hints that the naive approach is missing interesting solution branches.

2.2. The TuRBO globalization of phase I

An alternative to the naive approach is to use a global optimization algorithm. The coil set it finds is then given as the starting point in phase I to a local gradient-based optimization algorithm BFGS. This global-to-local approach was explored in Glas *et al.* (2022), where stochastic global optimization of coil geometries was done using the DTuRBO algorithm (Padidar *et al.* 2021), an extension of the gradient-free global optimization algorithm TuRBO (Eriksson *et al.* 2019) that takes into account gradients. The TuRBO algorithm solves the minimization problem

$$\left. \begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^N} f(\mathbf{x}) \\ \ell_i \leq x_i \leq u_i, \quad \text{for } i = 1 \dots N, \end{aligned} \right\} \quad (2.5)$$

where N is the dimension of the optimization problem, ℓ_i, u_i are respectively lower and upper bounds on the components of the control variables. That is, the algorithm attempts to find the global minimum $f(\mathbf{x})$ on an N -dimensional hyper-rectangle. The TuRBO algorithm completes a preliminary exploration phase by quasi-uniformly sampling the hyper-rectangle to find interesting areas of parameter space. Then, it uses trust-region-based Bayesian optimization algorithms to further refine the solution.

It is difficult to define sensible lower and upper bounds in (2.5) on the unknowns when they are the Fourier coefficients of the coils. There should be some decay in the size of the boxes with increasing mode number, but it unclear how this decay should be chosen. In Glas *et al.* (2022), manufacturing errors were projected onto a Fourier basis, from which reasonable bounds on the Fourier coefficients were inferred. In this work, we propose a different way of determining the box constraints for deterministic coil design problems, however, our technique is generic enough to be used in the stochastic context as well.

2.2.1. Bounding boxes

It is straightforward to constrain the coil currents to $-1 \leq \mu_0 I_i \leq 1$, where $\mu_0 = 4\pi \times 10^{-7}$ is the magnetic constant. The parameter $\bar{\eta}$ can also be constrained by $0 \leq \bar{\eta} \leq 2$.

We still use a Fourier series to represent the coils and axis, but change the degrees of freedom from Fourier coefficients to spatial coordinates, linked to each other by the discrete Fourier transform. In this way, physically meaningful box constraints can be provided to TuRBO. These positions, or anchor points, in cylindrical coordinates, (r_i, θ_i, z_i) , are constrained to the box $[0, 1 + R_{\text{minor}}] \times [\theta_i - \Delta\theta/2, \theta_i + \Delta\theta/2] \times [-R_{\text{minor}}, R_{\text{minor}}]$, where $R_{\text{minor}} = L_{\text{target}}/2\pi$, and $\Delta\theta = \rho(2\pi/2n_{\text{fp}}n_{\text{coils per hp}})$, $\theta_i = (2\pi/2n_{\text{fp}}n_{\text{coils per hp}})(i + 1/2)$, R_{minor} is the radius of the perfectly circular coil with length L_{target} , $\Delta\theta$ defines a cylindrical sector that the coil can occupy and the ρ multiplier ensures that the coil bounding boxes overlap somewhat. Randomly sampling $2N_{f,c} + 1$ points on the box will frequently generate coils with complex geometries (figure 3a). These geometries are not particularly useful and we would like TuRBO to avoid wasting time in these uninteresting areas of parameter space. These configurations can be avoided, while still using the same anchor points, by re-indexing them so that the points are ordered counterclockwise about their barycentre in the RZ plane (figure 3b). This unravels the coils. Finally, these coil coordinates are converted to Cartesian coordinates, projected onto a Fourier basis, then used to evaluate the near-axis objective (2.4).




FIGURE 3. The randomly generated coil in cylindrical coordinates in (a) is too complex. By *Re*-indexing the nodes by ordering them about their barycentre in the *RZ* plane, the coil untangles in (b). (c) A set of coils, the magnetic axis and their associated bounded boxes in physical space are illustrated. The full stellarator is not shown, but can be obtained by applying symmetries.

We also generate (R_i, ϕ_i, Z_i) positions, or anchor points, for the magnetic axis, where (R_i, Z_i) are constrained to the box $[1 - 1/(1 + n_{fp}^2), 1 + 1/(1 + n_{fp}^2)] \times [-0.2, 0.2]$, and ϕ_i is a uniformly spaced cylindrical angle in $2\pi/n_{fp}$. The bounding boxes for R_i and Z_i are informed by the fact that quasi-axisymmetric magnetic axes have zero helicity. Helicity is an integer associated with the axis geometry that measures how many poloidal transits the axis normal vector makes as the axis traces one toroidal revolution. It was shown in Rodriguez, Sengupta & Bhattacharjee (2022) that magnetic axes of the form

$$R(\phi) = 1 + a \cos(n_{fp}\phi), \tag{2.6}$$

$$Z(\phi) = b \sin(n_{fp}\phi), \tag{2.7}$$

have zero helicity when $a < 1/(1 + n_{fp}^2)$. The situation is more delicate if additional Fourier harmonics are used to represent the axis (as we do here), but in the case of a single dominant harmonic, this bound on a is a well-informed estimate. The box associated with Z_i was chosen to prevent large excursions from the $Z = 0$ plane, although the size of this box was somewhat arbitrary and other values might have been used.

There is one final detail that ensures that the resulting axis curve is stellarator symmetric. After sampling $N_{f,a} + 1$ points from the bounding box associated with the radial coordinates R_i , they are arranged in an $2N_{f,a} + 1$ -sized array as follows:

$$[R_0 - s, R_1 - s, \dots, R_{N_{f,a}} - s, R_{N_{f,a}} - s, \dots, R_1 - s], \tag{2.8}$$

where s is a shift to ensure the array has mean 1. After the shift, the radial positions of the magnetic axis may lie slightly outside the original bounding box. Similarly, we sample $N_{f,a}$ points from the bounding box associated with the vertical coordinates Z_i , and arrange them in a $2N_{f,a} + 1$ -sized array as follows:

$$[0, Z_1, \dots, Z_{N_{f,a}}, -Z_{N_{f,a}}, \dots, -Z_1]. \tag{2.9}$$

The stellarator symmetric R and Z harmonics of the magnetic axis are obtained by discrete Fourier transform of these arrays and provided to the near-axis objective (2.4).

The bounding boxes of different coils can overlap each other and all the bounding boxes of the coils necessarily overlap the one of the magnetic axis, since we do not search for windowpane coils. Thus, it is possible for the generated coils to be linked, or for

the coils to not be linked with the magnetic axis. In addition, it is still possible that the magnetic axis sampled from the bounding boxes to have non-zero helicity. It is for this reason that we add penalty terms to the objective f_{axis} that avoid these undesirable configurations

$$f(\mathbf{c}, \mathbf{I}, \mathbf{a}, \bar{\eta}) = f_{\text{axis}}(\mathbf{c}, \mathbf{I}, \mathbf{a}, \bar{\eta}) + w \left[\sum_i \sum_{j>i} l_{i,j}(\mathbf{c}) + \sum_i |l_{i,\text{axis}}(\mathbf{c}, \mathbf{a}) - 1| + h(\mathbf{a}) \right], \quad (2.10)$$

where we abuse notation here and use \mathbf{c}, \mathbf{a} to represent the vector of anchor points on the coils and axis, respectively. These positions can be converted to the Fourier coefficients by using a discrete Fourier transform. The value $l_{i,j}$ is the absolute value of the linking number between coils i and j , and $l_{i,\text{axis}}$ is the absolute value of the linking number between coil i and the magnetic axis. For the linking number calculation, we use the implementation in Bertolazzi, Ghiloni & Specogna (2019) since it returns a linking number with certified accuracy by tracking rounding errors in the calculation. Finally, $h(\mathbf{a})$ is the absolute value of the axis helicity. The weight $w = 10^7$ is chosen to make the terms that it multiplies greater than f_{axis} . The penalty terms multiplied by w are zero when the design variables describe a stellarator with unlinked coils that each wrap once poloidally around a magnetic axis with zero helicity. Since the objective is no longer differentiable due to the linking number and helicity calculation, we use TuRBO rather than DTuRBO.

We also designed stellarators using an axis bounding box independent of n_{fp} , $(R_i, Z_i) \in [0.85, 1.15] \times [-0.15, 0.15]$, and neglecting to include the helicity penalty. In a postprocessing step, devices that had an axis with non-zero helicity were discarded. This approach risks discarding devices and wasting computational resources but it is capable of producing comparable results.

2.3. Illustration of TuRBO globalization followed by phase I

To illustrate the global-to-local procedure, we examine the TuRBO globalization procedure followed by the phase I optimization for a single stellarator. We use a batch size of 100 and restrict TuRBO to a budget of 15 000 function evaluations, of which 1000 are used in the initial exploration phase; TuRBO's internal Gaussian process model is only trained using the most recent 1000 function evaluations for speed. Finally, the underlying Fourier discretization uses $N_{f,c} = N_{f,a} = 2$ for the coils and axis during this global exploration phase. The coil set and magnetic axis from the global search are used as the initial guess to the BFGS algorithm to polish the solution. During this phase, the number of Fourier modes used is increased to $N_{f,c} = N_{f,a} = 6$. Before BFGS obtains a good enough approximation to the inverse Hessian, especially during the first few iterations, the line search might accept a large enough step such that the coils become interlocked or unlinked with the magnetic axis. To avoid this, we use a linking number calculation to trigger the line search.

In figure 4 we provide the objective function evaluations over the course the TuRBO globalization then phase I of figure 1. The TuRBO phase of the run has large variations in the objective function, varying over several orders of magnitude. When the objective is larger than 10^7 , this likely corresponds to a device that has an axis with non-zero helicity, linked coils or coils unlinked with the axis; TuRBO quickly finds an interesting region of parameter space and it could have used a smaller budget. The notable jump in the objective function value around 11 000 function evaluations occurs when TuRBO no longer can

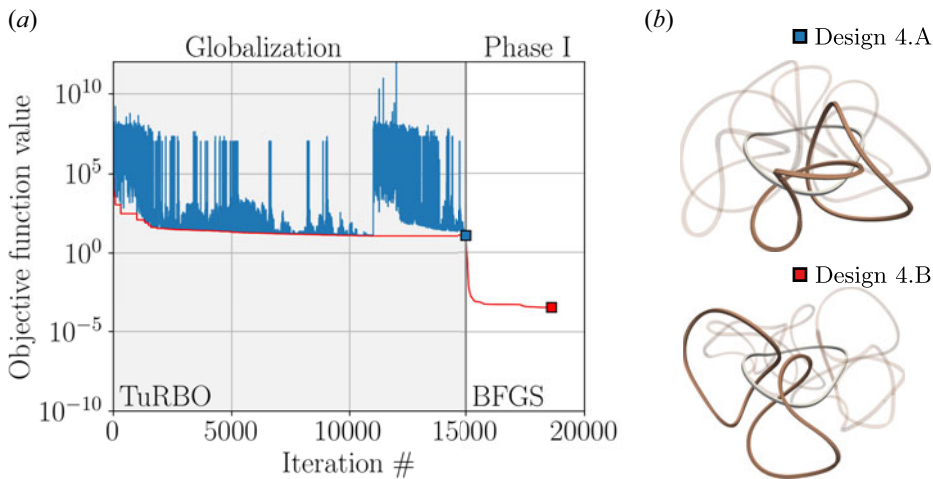


FIGURE 4. The value of the objective function when applying TuRBO then BFGS. For the TuRBO phase, the objective values are evaluations of (2.10), which includes the non-differentiable penalties and the ‘iteration #’ is actually number of function evaluations. For the BFGS phase, the objective values are evaluations of (2.2), and ‘iteration #’ refers to the objective function values after an accepted line search. The red line is the best function evaluation so far. (b) We plot the best stellarator found by TuRBO, given by Design 4.A, provided as the initial guess to BFGS in phase I. Then, the stellarator found by BFGS after increasing the number Fourier modes is given by Design 4.B.

make progress in the trust region, so it is discarded and another one is initialized. The best configuration found by TuRBO is provided and labelled Design 4.A. The coil set at the end of phase I (Design 4.B) has changed substantially from those obtained after the TuRBO phase.

2.4. Comparison of devices from TuRBO and naive globalization after phase I

In this section, we compare the performance of TuRBO with the more naive approach by solving (2.1) for many target coil length values when $\iota = 0.9$, n_{fp} , $n_{coils\ per\ hp} = 2$. To this end, we execute the TuRBO globalization followed by phase I, as illustrated in § 2.3 and figure 4. In both naive and TuRBO approaches, the workflow is executed 16 times per target coil length. The quasisymmetry error of the resulting configurations are provided in figure 5. TuRBO finds devices that outperform those found by the naive approach, and we plot two devices from each algorithm, called Designs 5.A and 5.B. We find that Design 5.B is from a genuinely different solution branch. To drive this home, we use it as an initial guess, we solve (2.1) for multiple different target coil lengths. This solution branch persists on both sides, and was evidently even missed by TuRBO. This highlights the difficulty of finding global minima.

With very modest input from the practitioner, TuRBO does a good job of localizing performant minima. We find that the choice of the overlap factor ρ only modestly affected the quality of the minimizers found. In contrast, one has to decide on how many Fourier modes to perturb and the size of each perturbation to use with the naive approach. This can have a large effect on the quality of the minimizers found.

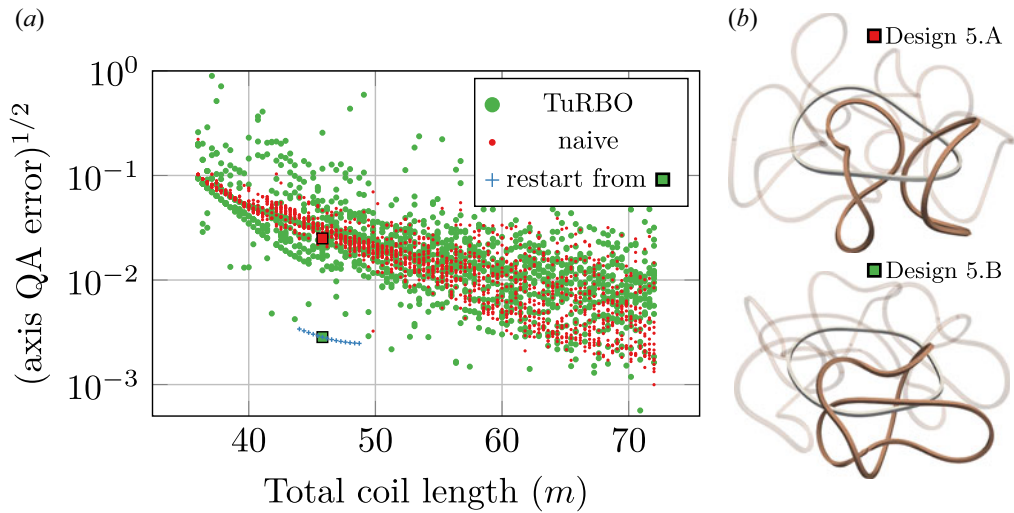


FIGURE 5. (a) We provide the trade-off between axis quasisymmetry error and total coil length used for local minima found using the naive approach and TuRBO with $\iota = 0.9$, and $n_{\text{fp}}, n_{\text{coils per hp}} = 2$. (b) We have plotted two coil sets, called Designs 5.A and 5.B, corresponding respectively to the red, and green square markers. Using Design 5.B as an initial guess, we solve (2.1) for different target coil lengths, resulting in the devices corresponding to the blue crosses. The axis QA error refers to the terms in (2.2).

3. The near-axis coil sets obtained after phase I

We solve (2.1) multiple times with combinations of the following design targets:

$$\left. \begin{aligned} \iota_{\text{target}} &= 0.1, 0.2, 0.3, \dots, 0.9, \\ L_{\text{target}} &= 4.5, 4.75, 5, \dots, 8.5, 8.75, 9, \\ n_{\text{coils per hp}} &= 1, 2, 3, 4, \dots, 13, \\ n_{\text{fp}} &= 1, 2, 3, 4, 5, \end{aligned} \right\} \quad (3.1)$$

using both the naive and TuRBO globalization approaches, then combining the data sets obtained. To avoid unrealistic configurations, we only use combinations of the above design targets such that the total coil length used $2n_{\text{fp}}n_{\text{coils per hp}}L_{\text{target}} \leq 120$ m. For all devices and physics targets, we use $d_{\text{min}} = 0.1$ m, $\kappa_{\text{max}} = 5 \text{ m}^{-1}$, $\kappa_{\text{msc}} = 5 \text{ m}^{-2}$. In figure 6, we provide the on-axis quasi-symmetry error for a subset of the stellarators, where we observe the expected trade-off between axis quasisymmetry and total coil length: as the total coil length gets longer, the minimum attainable quasisymmetry error decreases and levels off. If one fixes the total coil length used, but increases the number of coils per half-period, this can drastically improve the on-axis quasisymmetry error. This is illustrated in Designs A and B plotted on the bottom row of figure 6. Both stellarators use the same length of coil (54 m), but have 1 and 2 coils per half-period, respectively, where shorter, more numerous coils are preferred. The coils in Design B also interleave with one another, in a similar fashion to the coils in figure 5. It appears that we can find configurations with a good approximation of on-axis quasi-symmetry for all values of n_{fp} . Moreover, for a fixed coil length, increasing the number of field periods seems to improve the on-axis quasisymmetry error, although the coil-to-axis distance decreases. As

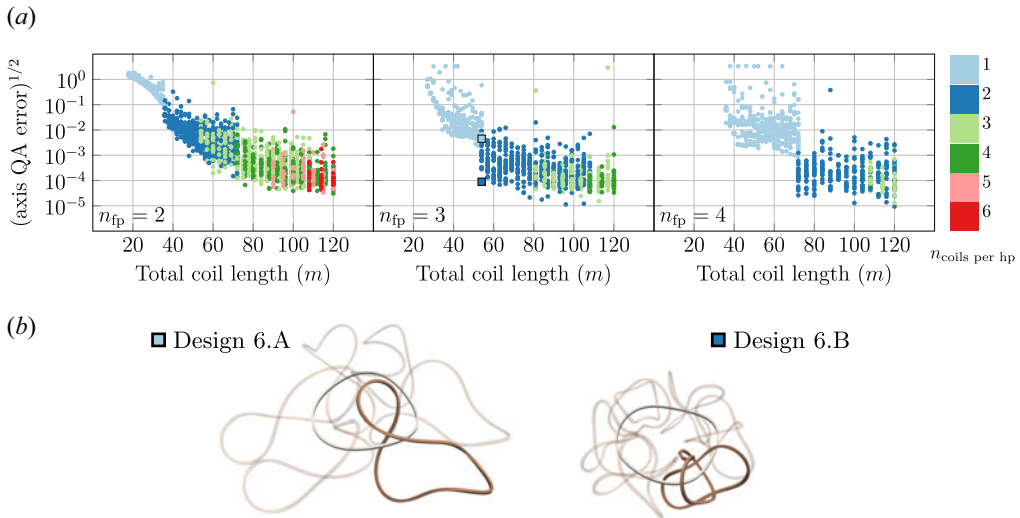


FIGURE 6. The near-axis QA error with respect to total coil length used when $\iota = 0.5$, each point on the scatter plot corresponds to a different device and each panel corresponds to a different number of field periods, and the colour refers to the number of coils per half-period used. Below, we plot two configurations (Designs 6.A and 6.B) that use the same total coil length but different numbers of coils, 1 and 2 coils per half-period, respectively. The axis QA error refers to the terms in (2.2).

we will observe in the next section, this observation does not bear out when optimizing for quasisymmetry on a volume.

The generated coil sets can produce nested flux surfaces, but sometimes only over a small volume in the neighbourhood of the magnetic axis, making it a good set of initial conditions for the next phase of coil optimization.

4. Phases II and III: stellarators optimized for quasisymmetry on a volume

Taking coil sets optimized for near-axis quasisymmetry obtained during phase I, we now optimize for quasisymmetry on a volume with varying aspect ratios during phases II and III. The objective that we minimize is the sum of the average (normalized) QA error and the Boozer residuals on N_s surfaces parametrized in Boozer coordinates φ, θ

$$\begin{aligned} & \hat{f}_{\text{surface}}(\mathbf{c}, \mathbf{I}, \mathbf{s}_1, \iota_1, G_1, \dots, \mathbf{s}_{N_s}, \iota_{N_s}, G_{N_s}) \\ & := \frac{1}{N_s} \sum_{k=1}^{N_s} \left\{ \frac{\int_0^1 \int_0^{1/n_{fp}} B_{\text{non-QA}}(\mathbf{c}, \mathbf{I}, \mathbf{s}_k, \varphi, \theta)^2 \|\mathbf{n}(\mathbf{s}_k, \varphi, \theta)\| \, d\varphi \, d\theta}{\int_0^1 \int_0^{1/n_{fp}} B_{\text{QA}}(\mathbf{c}, \mathbf{I}, \mathbf{s}_k, \theta)^2 \|\mathbf{n}(\mathbf{s}_k, \varphi, \theta)\| \, d\varphi \, d\theta} \right. \\ & \quad \left. + \frac{1}{2} w_r \int_0^1 \int_0^{1/n_{fp}} \|\mathbf{r}(\mathbf{s}_k, \iota_k, G_k, \mathbf{c}, \mathbf{I}, \varphi, \theta)\|^2 \, d\varphi \, d\theta \right\}, \end{aligned} \tag{4.1}$$

where Cartesian coordinates of points on the k th surface $\Sigma(\mathbf{s}_k, \varphi, \theta) \in \mathbb{R}^3$ are expressed using the Fourier representation in Giuliani *et al.* (2022b). A surface has n_s degrees of freedom associated with it, stored in $\mathbf{s}_k \in \mathbb{R}^{n_s}$. In the objective function above, the surface normal is

$$\mathbf{n}(\mathbf{s}_k, \varphi, \theta) = \frac{\partial \Sigma}{\partial \varphi}(\mathbf{s}_k, \varphi, \theta) \times \frac{\partial \Sigma}{\partial \theta}(\mathbf{s}_k, \varphi, \theta), \tag{4.2}$$

and $w_r > 0$ is a weighting parameter in front of a partial differential equation (PDE) residual, described below. The electromagnetic coils also have an associated current stored in the vector $\mathbf{I} \in \mathbb{R}^{N_c-1}$, where the dimensionality is one less than the number of coils because the first coil's current is fixed, preventing the currents from approaching zero. Analogously to the previous phase, the vector of coil degrees of freedom are stored in $\mathbf{c} \in \mathbb{R}^{3N_c(2N_{f,c}+1)}$. The x, y, z positions of the N_c base modular coils are each represented using a Fourier series with $N_{f,c}$ modes. The Boozer residual \mathbf{r} , and scalars G_k and ι_k are described in more detail below.

The terms in the QA error ratio of (4.1), introduced in Giuliani *et al.* (2022b), are given by

$$B_{QA}(\mathbf{c}, \mathbf{I}, \mathbf{s}_k, \theta) = \frac{\int_0^{1/n_{fp}} B(\boldsymbol{\Sigma}(\mathbf{s}_k, \varphi, \theta), \mathbf{c}, \mathbf{I}) \|\mathbf{n}(\mathbf{s}_k, \varphi, \theta)\| \, d\varphi}{\int_0^{1/n_{fp}} \|\mathbf{n}(\mathbf{s}_k, \varphi, \theta)\| \, d\varphi}, \tag{4.3}$$

with

$$B_{\text{non-QA}}(\mathbf{c}, \mathbf{I}, \mathbf{s}_k, \varphi, \theta) = B_{QA}(\mathbf{c}, \mathbf{I}, \mathbf{s}_k, \theta) - B(\mathbf{c}, \mathbf{I}, \mathbf{s}_k, \varphi, \theta), \tag{4.4}$$

where $B(\boldsymbol{\Sigma}(\mathbf{s}_k, \varphi, \theta), \mathbf{c}, \mathbf{I})$ is the field strength evaluated on the k th magnetic surface.

The goal is to design a set of coils that solves the following optimization problem:

$$\min_{\mathbf{c}, \mathbf{I}, \mathbf{s}_1, \iota_1, G_1, \dots, \mathbf{s}_{N_s}, \iota_{N_s}, G_{N_s}} \hat{f}_{\text{surface}}(\mathbf{c}, \mathbf{I}, \mathbf{s}_1, \iota_1, G_1, \dots, \mathbf{s}_{N_s}, \iota_{N_s}, G_{N_s}) \tag{4.5a}$$

$$\text{subject to } \mathbf{g}_k(\mathbf{s}_k, \iota_k, G_k, \mathbf{c}, \mathbf{I}) = 0, \quad \text{for } k = 1, \dots, N_s, \tag{4.5b}$$

$$\iota_{\text{target}} - \frac{1}{N_s} \sum_{k=1}^{N_s} \iota_k = 0, \tag{4.5c}$$

$$R_{\text{major}}(\mathbf{s}_{N_s}) = 1, \tag{4.5d}$$

$$\sum_{i=1}^{N_c} L_i \leq L_{\text{max}}, \tag{4.5e}$$

$$\kappa_i \leq \kappa_{\text{max}}, \quad i = 1, \dots, N_c, \tag{4.5f}$$

$$\frac{1}{L_c^{(i)}} \int_{\Gamma^{(i)}} \kappa_i^2 \, dl \leq \kappa_{\text{msc}}, \quad i = 1, \dots, N_c, \tag{4.5g}$$

$$\|\boldsymbol{\Gamma}^{(i)} - \boldsymbol{\Gamma}^{(j)}\| \geq d_{\text{min}}, \quad \text{for } i \neq j, \tag{4.5h}$$

$$\|\boldsymbol{\Gamma}^{(i)}\| - L^{(i)} = 0, \quad \text{for } i = 1, \dots, N_c. \tag{4.5i}$$

The constraint \mathbf{g}_k is a system of equations that relates the coil and surface degrees of freedom. This constraint is a spatial discretization that approximates solutions to the system of partial differential equations

$$\left. \begin{aligned} \mathbf{r}(\mathbf{s}_k, G_k, \iota_k, \mathbf{c}, \mathbf{I}, \varphi, \theta) &= 0, \\ V(\mathbf{s}_k) - V_{\text{target},k} &= 0, \end{aligned} \right\} \tag{4.6}$$

where the residual $\mathbf{r}(s_k, G_k, \iota_k, \mathbf{c}, \mathbf{I}, \varphi, \theta) \in \mathbb{R}^3$ is

$$\begin{aligned} \mathbf{r}(s_k, G_k, \iota_k, \mathbf{c}, \mathbf{I}, \varphi, \theta) &:= G_k \frac{\mathbf{B}(\boldsymbol{\Sigma}(s_k, \varphi, \theta), \mathbf{c}, \mathbf{I})}{\|\mathbf{B}(\boldsymbol{\Sigma}(s_k, \varphi, \theta), \mathbf{c}, \mathbf{I})\|} \\ &- \|\mathbf{B}(\boldsymbol{\Sigma}(s_k, \varphi, \theta), \mathbf{c}, \mathbf{I})\| \left(\frac{\partial \boldsymbol{\Sigma}(s_k, \varphi, \theta)}{\partial \varphi} + \iota_k \frac{\partial \boldsymbol{\Sigma}(s_k, \varphi, \theta)}{\partial \theta} \right), \end{aligned} \tag{4.7}$$

and $V(s_k)$ is the volume enclosed by the surface $\boldsymbol{\Sigma}(s_k, \varphi, \theta)$ and $V_{\text{target},k}$ is the user-specified target volume of the surface. The inputs to the PDE are the surface label, $V_{\text{target},k}$, and the electromagnetic coils \mathbf{c}, \mathbf{I} . The outputs are the surface degrees of freedom s_k, G_k and rotational transform ι_k . Note that, for a given surface solve, the geometry of the coils \mathbf{c} and their currents \mathbf{I} are fixed.

We have devised two approaches to solve (4.6) using a collocation method. In both approaches, constraint (4.5b) allows us to compute the surface degrees of freedom s_k, ι_k, G_k in terms of the electromagnetic coils \mathbf{c}, \mathbf{I} .

The approach used during phase II, called BoozerLS surfaces, is useful for regimes where nested flux surfaces do not exist, and solutions to (4.6) do not exist. In this case, we search for solutions that satisfy the PDE in a least squares sense. This results in a bilevel optimization problem, comprising an outer optimization problem over the coil geometries and currents to optimize the device for quasisymmetry. The inner optimization problem is over the surface degrees of freedom for fixed coil geometry \mathbf{c} and currents \mathbf{I} , and computes surfaces that solve (4.6) in a least squares sense

$$\begin{aligned} s_k(\mathbf{c}, \mathbf{I}), G_k(\mathbf{c}, \mathbf{I}), \iota_k(\mathbf{c}, \mathbf{I}) &= \operatorname{argmin}_{\tilde{s}_k, \tilde{G}_k, \tilde{\iota}_k} \int_0^1 \int_0^{1/n_{\text{fp}}} \|\mathbf{r}(\tilde{s}_k, \tilde{G}_k, \tilde{\iota}_k, \varphi, \theta, \mathbf{c}, \mathbf{I})\|^2 d\varphi d\theta \\ &+ \frac{1}{2} w_{v,k} (V(\tilde{s}_k) - V_{\text{target},k})^2, \end{aligned} \tag{4.8}$$

where the first term in the objective above measures how accurately the PDE is solved and the second term premultiplied by the weight $w_{v,k}$ ensures that we solve for a magnetic surface with target volume $V_{\text{target},k}$. In this case, (4.5b) is the optimality condition of the inner least squares optimization problem (4.8) solved by BFGS. The approach makes the surface solve more robust because we do not require the residual be zero at a set of collocation points and we can use robust line search based algorithms to minimize the PDE residual. That being said, solving this inner optimization problem to sufficient accuracy can be expensive, which is the price paid for robustness.

The approach used during phase III, called BoozerExact surfaces, is useful for regimes when nested flux surfaces do exist and we would like to polish the quality of QA. In this case, (4.5b) is the vector of residual (4.7) evaluations on a tensor product grid of collocation points in (φ, θ) , which we require to be zero. This is a stronger requirement, and makes the numerical method more brittle: when nested magnetic surfaces do not exist, Newton’s method may fail. At the expense of robustness, solving the constraint here is much less computationally expensive.

Using the implicit function theorem, we minimize the reduced objective

$$f_{\text{surface}}(\mathbf{c}, \mathbf{I}) = \hat{f}_{\text{surface}}(\mathbf{c}, \mathbf{I}, s_1(\mathbf{c}, \mathbf{I}), \iota_1(\mathbf{c}, \mathbf{I}), G_1(\mathbf{c}, \mathbf{I}), \dots, s_{N_s}(\mathbf{c}, \mathbf{I}), \iota_{N_s}(\mathbf{c}, \mathbf{I}), G_{N_s}(\mathbf{c}, \mathbf{I})) \tag{4.9}$$

by eliminating s_k, ι_k, G_k via the constraint (4.5b). Constraint (4.5c) ensures that the mean rotational transform on the volume is the target value ι_{target} . Constraint (4.5d) ensures that

the sum of coil lengths per half-period is less than $L_{\max} = n_{\text{coils per hp}} L_{\text{target}}$, where L_{target} is the coil length used in the near-axis optimization problem (2.1). This is notably different to previously, where each coil was constrained to have the same length L_{target} . Constraint (4.5e) ensures that the major radius of the outermost surface is 1. Constraints (4.5f) and (4.5g) ensure that the maximum coil curvature and mean squared curvature do not respectively exceed κ_{\max} and κ_{msc} . Constraint (4.5g) ensures that the minimum intercoil distance does not decrease below d_{\min} . Finally, constraint (4.5i) ensures that the coils have uniform incremental arclength. Constraints (4.5c)–(4.5i) are enforced using a penalty method, and are satisfied to 0.1% precision. We always use the BFGS quasi-Newton method in phases II and III of the workflow.

Before we can begin solving these optimization problems, we require N_s magnetic surfaces in the volume of the device parametrized in Boozer angles. The initial surfaces are obtained by starting with the optimized magnetic axis obtained during phase I. Then, a first-order near-axis expansion is used to find the surface geometry in the neighbourhood of the axis with minor radius $r = 0.05$ m. That is, the surface degrees of freedom s_1 are found such that we have

$$\Sigma_1(s_1, \varphi, \theta) = \Gamma^{\text{axis}}(\varphi) + rX_1(\varphi, \theta)\mathbf{n}(\varphi) + rY_1(\varphi, \theta)\mathbf{b}(\varphi), \quad (4.10)$$

where \mathbf{n} , \mathbf{b} are the normal and binormal vectors associated with the Frenet frame of the magnetic axis, X_1 , Y_1 are defined in Landreman, Sengupta & Plunk (2019) and Landreman (2023) and r is a radial distance from the axis. The remaining $N_s - 1$ surfaces with lower aspect ratio are computed using a continuation procedure in the surface label, e.g. volume enclosed by the surface. During phase III, we already have BoozerLS surfaces available from phase II and those are used as initial guesses for the BoozerExact surfaces. The BoozerLS surfaces are typically very good initial guesses for the BoozerExact solve.

4.1. Illustration of phases II and III

In this section, we give more details on the final two phases of the coil design algorithm to illustrate how it works. The configurations obtained from the near-axis formulation (phase I) only attempt to find coils that target quasisymmetry on the magnetic axis. The formulation does not control anything about the magnetic field away from the axis and, as a result, magnetic surfaces may not exist on a large volume (Lee *et al.* 2022). The BoozerLS surfaces (phase II) are robust and capable of healing generalized chaos (Giuliani *et al.* 2023), although it can be computationally expensive. Finally, the BoozerExact (phase III) surfaces polish the coil set for precise QA and are computationally much cheaper.

The algorithm for phase II is given in Algorithm 1. The first step of the algorithm is to use the magnetic axis obtained from phase I to compute an approximate magnetic surface from (4.10). Then, the optimization problem (4.5) is solved, but we do not attempt to fully converge the coil sets and only complete a total of 300 iterations of BFGS. Each time BFGS is restarted, an attempt is made to increase the number of Fourier modes used to represent the surface. Starting from $m_{\text{pol}}, n_{\text{tor}} = 2$, we try to compute a BoozerLS surface. If the solve succeeds, we use the surface as an initial guess and increase $m_{\text{pol}}, n_{\text{tor}}$ by 1. If the surface solve fails after increasing the surface degree, then the degree is reverted to the previous one. If the solve succeeds, the continuation in degree continues until $m_{\text{pol}}, n_{\text{tor}} = 4$. Progressively increasing the number of Fourier modes makes the procedure more robust, as we find self-intersecting surfaces are less likely to occur.

The weight w_r is chosen such that the Boozer residual term is an order of magnitude larger than the non-quasisymmetry penalty. This term favours nested flux surfaces and improves the robustness of the optimization algorithm. However, it also competes with




FIGURE 7. Poincaré plots at $\phi = 0$ after phases I and II for aspect ratios $AR = 20, 10$. The point where the magnetic axis intersects the $\phi = 0$ plane corresponds to the red dot in (a). Cross-sections of the surfaces on which QA is optimized are red in (b,c).

the QA error and so the quality of quasisymmetry obtained at this stage can be limited. Since the stellarators obtained here are not fully converged and strongly depend on the value of w , we do not analyse the physics properties of the stellarators here yet. In figure 7, we provide Poincaré plots of a device just after the near-axis optimization, and then as BoozerLS surfaces are added to optimize for nested flux surfaces. The first panel shows that the near-axis optimization algorithm does not guarantee nested surfaces far away from the magnetic axis. The final two panels illustrate that as the surfaces are added, the volume with nested flux surfaces increases.

Algorithm 1 Phase II: BoozerLS

```

compute an initial BoozerLS surface with aspect ratio 20 from near-axis formula (4.10).
for aspect ratio (AR) in {20, 10, 6.66, 5, 4, 3.33, 2.85} do
  for  $i$  in 1, 2, 3 do
    attempt to increase the degree of the BoozerLS surfaces up to  $m_{\text{pol}}, n_{\text{tor}} \leq 4$ 
    solve (4.5) on surfaces with a fixed budget of iterations.    ▷ Island and chaos
  healing
  end for
  add another BoozerLS surface with the next smallest AR.
end for

```

Algorithm 2 Phase III: BoozerExact

```

for  $i$  in 1, 2, 3, 4 do
  keep rational surfaces as BoozerLS, otherwise convert to BoozerExact.
  attempt to increase the degree of the BoozerLS surfaces up to  $m_{\text{pol}}, n_{\text{tor}} \leq 4$ 
  attempt to increase the degree of the BoozerExact surfaces up to  $m_{\text{pol}}, n_{\text{tor}} \leq 10$ 
  solve (4.5) on surfaces with a fixed budget of iterations.    ▷ Polish
end for

```

After phase II terminates, we proceed to phase III, where we optimize for precise quasisymmetry. It can be shown that the island width scales like $\sqrt{|B_{n,m}|/m'}$ (Hudson, Monticello & Reiman 2001), where l' is the magnetic shear. Thus, for large enough m , the island width should be small. Informed by this estimate, we retain the BoozerLS formulation when the rotational transform on the surface is within 1% of a rotational transform of the form $\iota = n_{\text{ip}}/m, 2n_{\text{ip}}/m$ for $m = 1, \dots, 15$. During this final phase, the weight w is set to zero on BoozerExact surfaces, but it maintained for BoozerLS surfaces in the neighbourhood of low-order rationals. At the start of every BFGS run, we attempt to increase the degree of the surface representation as we did during phase II, up to $m_{\text{pol}}, n_{\text{tor}} = 10$ for BoozerExact surfaces, and 4 for the remaining BoozerLS surfaces. When all surfaces are of type BoozerExact, we cap the total number of BFGS iterations to 20 000. When there is a surface of type BoozerLS, we cap the total number of BFGS iterations 4000. The algorithm for phase III is given in Algorithm 2.

During these last two phases of coil optimization, we also increase the number of Fourier harmonics in the coils to $N_{f,c} = 16$, because using too few Fourier modes in the coils can limit the attainable quality of quasisymmetry.

4.2. Comparison of devices from TuRBO and naive globalization after phase III

After the full workflow is completed (globalization, then phases I, II, then III), we compare the devices found depending on whether the TuRBO or naive globalization approach was applied just before phase I. In figure 8, we plot volume quasisymmetry error for configurations found for both globalization techniques. It appears that the highly favourable devices discovered by TuRBO do not necessarily translate to devices that have comparatively favourable volume quasisymmetry. We also find the TuRBO devices resulted in many more configurations for shorter coils than the ones found by the naive approach, although there is a larger spread of device performance. The Design 5.B does not persist upon optimizing for volume quasisymmetry, which we suspect is because the coils are too close to the magnetic axis (16.5 cm). As a result, when the surface generated by the near-axis formula (4.10) is used at the start of phase II, the BoozerLS surface solve in (4.5b) does not converge and the optimization cannot proceed. The devices discovered by TuRBO also do not appear to outperform those found by the naive approach. Nevertheless, we emphasize that this may only be because the globalization is performed on the near-axis problem just before phase I and not directly to the volume QA optimization, just before phase II. If globalization is applied directly to the BoozerLS phase of the workflow, it is possible that new and interesting devices could still be discovered. Finally, the magnitude and form of the perturbation ϵ used in the naive approach was determined after somewhat lengthy tuning (§ 2.2). There are much fewer *ad hoc* choices in the bounding box approach used by TuRBO.

4.3. Computational details

The run time of each phase can vary depending on the design targets, as the difficulty of the optimization problems strongly depends on the total coil length used by the device. In previous work, we used various levels of parallelism simultaneously such as MPI, OpenMP and SIMD. In this work, we take the approach of only allocating one core per stellarator and only use SIMD parallelism locally to a core when possible. The full workflow for a single device takes of the order of a day or two on a single core. However, this duration can vary substantially depending on how accurately each individual optimization problem is solved or if more cores are used per problem.

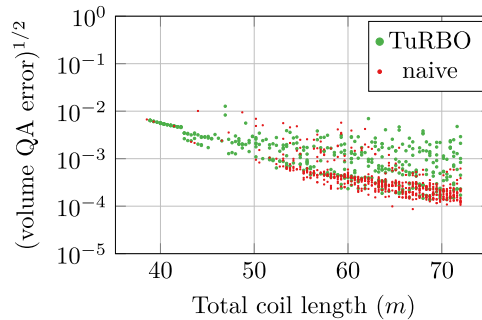


FIGURE 8. Taking the stellarators from figure 5, and using them as initial guesses for the volume QA optimization phase (§ 4), we obtain the stellarators plotted above. So that the data sets can be more easily distinguished, we use smaller marker sizes for devices corresponding to the naive algorithm.

5. Quasi-symmetric stellarator repository (QUASR)

The algorithms presented in the previous sections have culminated in the comprehensive database that we detail in this section. Before discussing the stellarators in the database, we would like to repeat some of the database's limitations:

- (i) The database only contains curl-free, stellarator symmetric magnetic fields with optimized QA. Note that our algorithms are generic, however, and may also be applied to other flavours of quasisymmetry as well.
- (ii) Our goal was both to produce a large data set of stellarators and to visualize the trade-off of target physics characteristics. Given that this approach might find stellarators that do not lie on the Pareto front, it may not be the most computationally efficient. Approaches based on continuation might reduce this computational cost (Bindel *et al.* 2023), at the expense of some parallelism.
- (iii) The algorithm might have discovered the same device multiple times. Three possible ways that this can occur are as follows. Given a local minimum of (2.1) or (4.5), a visually distinct local minimum can be found by reflecting the device about the XY plane, i.e. applying the transformation $-Z \rightarrow Z$. Another visually distinct local minimum can also be found by rotating the device by a half-period. We have also noticed that local minima can lie in tricky valleys of the objective. In particular, initially the gradient-based optimizer makes a lot of progress, and does so quickly. But after a few thousand iterations, progress slows, especially for devices that use longer coil lengths. Therefore, it might also be that visually distinct local minima will merge with one another after more iterations, but with only marginal reduction of the objective. This is the price paid for an extensive scan.

Due to the computational expense of generating these devices, we also include in the database the ones discovered after executing previous versions of the workflow described above. In the following sections, we compare devices in the database with ones discussed previously in the literature, highlight a device with rotational transform profiles that pass through low-order rationals and perform a couple of trade-off analyses of the devices. We look at how accurately quasisymmetry can be attained when the total coil length, number of coils per half-period, number of field periods, device aspect ratio and target mean rotational transform are varied. In addition, we examine the relationship between

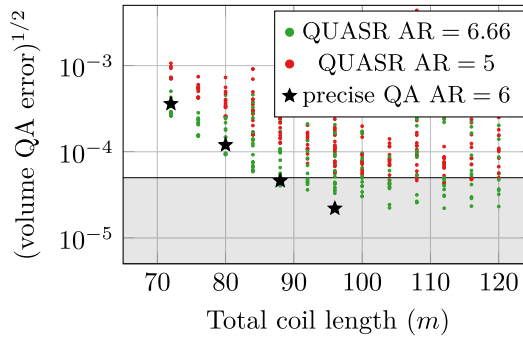


FIGURE 9. Comparison of devices from QUASR with aspect ratio 6.66 and 5, target mean rotational transform $\iota = 0.4$ and the precise QA coil sets of Giuliani *et al.* (2022b) with aspect ratio 6 and mean rotational transform 0.42. All devices have $n_{\text{fp}} = 2$ and $n_{\text{coils per hp}} = 4$ here. The grey region corresponds to the Earth's background magnetic field of $50 \mu\text{T}$, and volume QA error refers to the mean non-QA ratio in (4.1).

elongation and quality of quasisymmetry. Due to the multiple possible analyses that might be performed, we only scratch the surface here.

5.1. Comparison with previous devices

As a first examination, we select the devices in our database that have the closest design targets to previous configurations computed in Giuliani *et al.* (2022b), where devices have aspect ratio 6, $\iota = 0.42$, $n_{\text{coils per hp}} = 4$, total coil length 72 m, 80 m, 88 m, 96 m and $n_{\text{fp}} = 2$. The devices in QUASR that are closest to these devices have the same number of coils, aspect ratio 6.66 or 5 and target mean rotational transform 0.4. In figure 9, these devices are compared, where we observe that the QUASR devices perform comparably, although notably the precise QA device with length 96 m is better. This might be because we only allow a total of 20 000 iterations of BFGS during phase III, while in generating the precise QA coil sets, more than double the number of iterations were allowed.

5.2. A stellarator with a rotational transform profile passing through a low-order rational

Consider a device with two field periods, $n_{\text{coils per hp}} = 3$, $\iota_{\text{target}} = 0.5$ and aspect ratio 4. Since the target mean rotational transform is a low-order rational, we would expect there to be a strong likelihood of an island chain somewhere in the toroidal volume. However, the island chains are small, as shown in figure 10.

5.3. Impact of design targets on attainable quasisymmetry

In figure 11, we plot the volume QA error with respect to coil length for various values of $n_{\text{fp}} = 2, 3, 4$, when optimizing for QA on a single surface of aspect ratio 20. Since this is such a high aspect ratio surface, this set-up is a close approximation to the near-axis problem in § 2, although the objective is asking for more out of the magnetic field. A notable difference with the near-axis results is that increasing the number of field periods does not appear to improve the attainable on-axis quasisymmetry, as is observed in figure 6. One possible reason for this is that the near-axis design problem (2.1) only optimizes for quasisymmetry to first order and introducing second-order near-axis penalties might explain this discrepancy. There do appear to be some highly performant devices that are outliers, so it might also be that our algorithm is just not discovering those

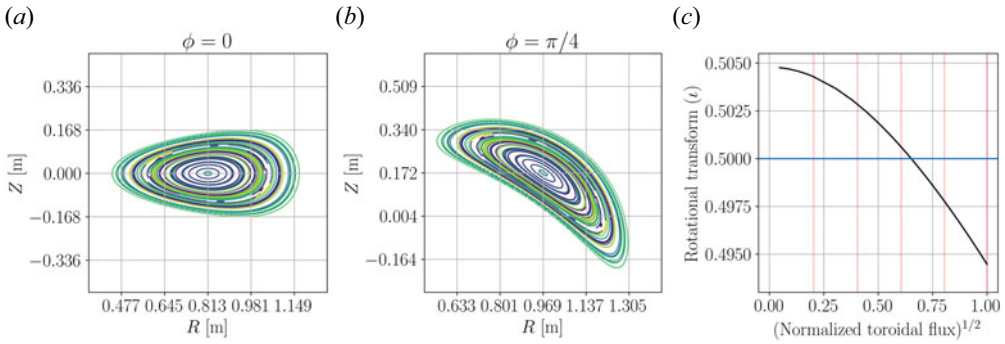


FIGURE 10. Poincaré plots and rotational transform profiles of a device with aspect ratio 4, mean rotational transform 0.5 and $n_{fp} = 2$. We visually could not identify trajectories associated with islands. The horizontal blue lines on the rotational transform profiles indicate the low-order rational $\iota = 1/2$. The vertical red lines on the rotational transform profiles indicate the normalized toroidal flux label on which quasisymmetry was optimized.




FIGURE 11. Devices on which quasisymmetry is optimized on a surface of aspect ratio 20, with a target rotational transform $\iota = 0.5$. The grey region corresponds to the Earth’s background magnetic field of $50 \mu\text{T}$, and volume QA error refers to the mean non-QA ratio in (4.1).

solution branches. There is also a clear preference for $n_{fp} = 2$ stellarators as the algorithm has difficulty finding devices with precise quasisymmetry for $n_{fp} = 3, 4$. Devices with $n_{fp} = 1, 5$ are not shown in the figure, but in both cases, the algorithm had varying difficulty finding devices with precise QA.

We also examine the trade-off between the quality of quasisymmetry, total coil length and device aspect ratio in figure 12. Increasing the target mean rotational transform does not appear to greatly affect the quality of attainable quasisymmetry, while as expected decreasing the device aspect ratio appears to have a much larger negative impact.

5.4. Maximum elongation

Finally, we study the values of elongation appearing in the database when $\iota_{\text{target}} = 0.6$ and an aspect ratio of 10. High elongation increases the ratio of surface area to volume of the magnetic surfaces, moreover, it can increase bunching of the flux surfaces and this can have negative implications for stability (Goodman *et al.* 2023). For each stellarator in our database, we compute the maximum elongation in the device by computing a surface with aspect ratio 80 in the neighbourhood of the magnetic axis. Then, we compute $N = 10$ cross-sections of that surface at cylindrical angles $\phi = \pi(i + 1/2)/n_{fp}/N$ for $i = 0, \dots, N - 1$. For each cross-section, we fit an ellipse in a least squares sense (Halir & Flusser 1998) and compute the ratio of the ellipse’s major to minor axes. The value reported here is the maximum ratio observed at these cross-sections.

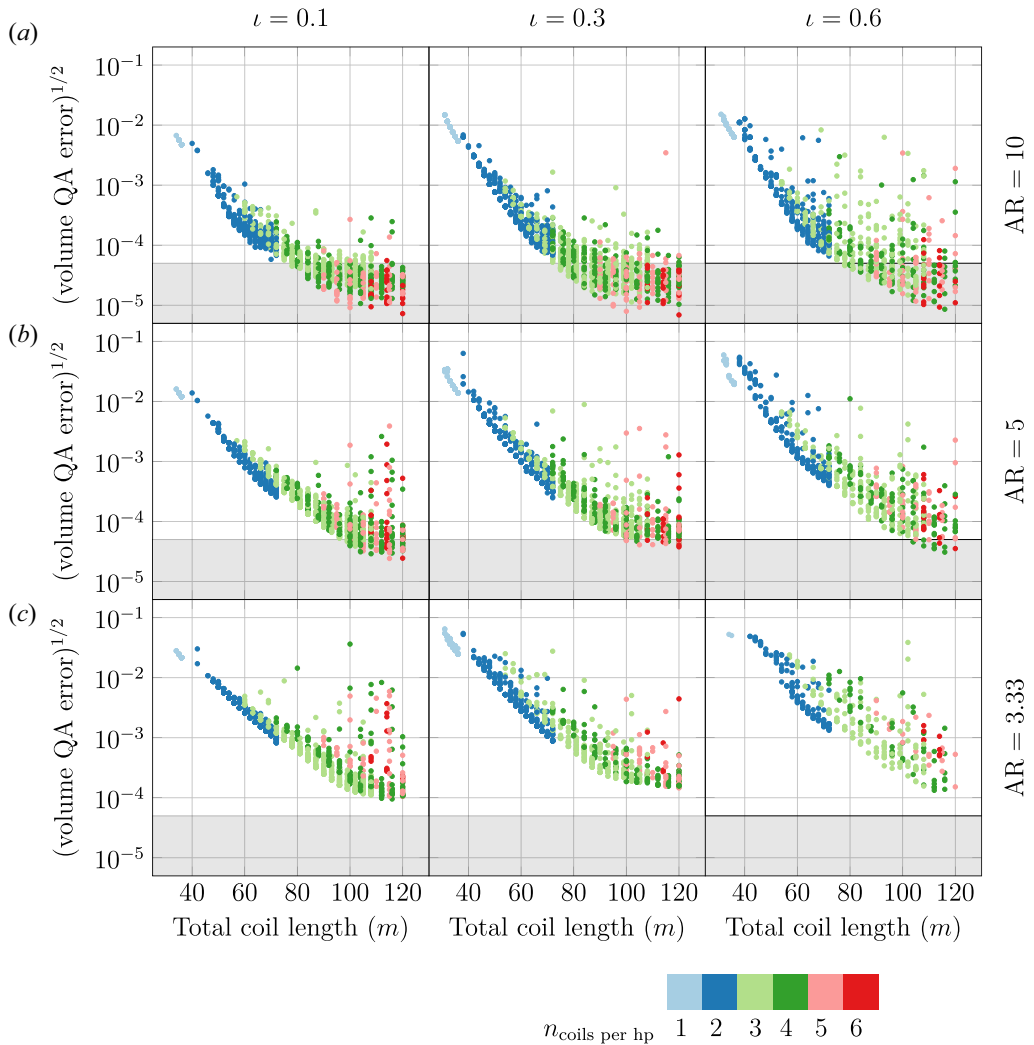


FIGURE 12. Trade-off between quality of quasisymmetry and total coil length for various device aspect ratios, and target mean rotational transform when $n_{fp} = 2$. The grey region corresponds to the Earth's background magnetic field of $50 \mu\text{T}$, and volume QA error refers to the mean non-QA ratio in (4.1).

We did not target any particular value of elongation but nevertheless there do appear to be favoured values (figure 13). For $n_{fp} = 1$ and 2, the lowest values of quasisymmetry error occur for an elongation around 7, while for $n_{fp} = 3$ an elongation of 4 is favoured. It is unclear whether this picture would change if a stage one optimization for QA were done instead, without coils. In Goodman *et al.* (2023), various QI stellarator designs were proposed, where a maximum elongation below approximately 6 was targeted. The stellarators in QUASR here illustrate that a good approximation of QA can also be found when requiring a maximum elongation below 6 too. We also observe that as expected, the lower quasisymmetry errors occur for longer coil lengths.




FIGURE 13. The volume quasisymmetry error with respect to the maximum elongation on a very high aspect ratio surface in the neighbourhood of the magnetic axis, for various values of n_{fp} . For all configurations plotted, $t_{target} = 0.6$, where the aspect ratio of the outermost surface is 10. The grey region corresponds to the Earth's background magnetic field of $50 \mu\text{T}$, and volume QA error refers to the mean non-QA ratio in (4.1).

6. Conclusions

We have proposed a direct stellarator coil design algorithm that is globalized using the TuRBO optimization algorithm, where box constraints on anchor points of the coils are applied. The algorithm combines three direct coil optimization algorithms, and has allowed the construction of a large database of approximately 200 000 vacuum-field stellarators for various design targets, e.g. aspect ratio, n_{fp} , n_{coils} per hp, rotational transform and total coil length. Using the database, we have examined the trade-off between accuracy of QA, total coil length, rotational transform and aspect ratio of the device.

Since the techniques in this work are quite general, there are many other directions that we would like to explore. Applying all of these approaches to other flavours of quasisymmetry, such as quasi-helical symmetry, is the next logical step. Adding windowpane coils and other coil geometries, e.g. helical coils, to the database would enrich the database. There is still room to improve the globalization algorithms and one possibility is to directly globalize the BoozerLS phase of the optimization.

Finally, we have only scratched the surface of possible physics analyses of the data set. Since it is publicly available, we hope that the stellarator community might explore it further.

Acknowledgements

The author would like to thank the Flatiron Institute's Scientific Computing Core for their support, G. Stadler, G.P. Langlois and R. Jorge for the helpful discussions, and M. Padidar for the tips for getting set up with TuRBO. We also would like to thank R. Jorge for linking PyQSC to SIMSOPT for obtaining initial surface guesses at the start of phase II in the workflow. Finally, the author would like to thank J. Soules for writing the web application to navigate the database.

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

Declaration of interests

The author reports no conflict of interest.

Data availability statement

- (i) Scripts that wrap phase I with the TuRBO globalization are available at <https://github.com/andrewgiuliani/Global-Direct-Coil-Optimization-I>.
- (ii) Scripts for phases II and III are available in SIMSOPT (Landreman *et al.* 2021) and <https://github.com/andrewgiuliani/Global-Direct-Coil-Optimization-II>.

- (iii) The entire set of stellarators in VMEC and SIMSOPT formats is archived on Zenodo at <https://doi.org/10.5281/zenodo.10050656>.
- (iv) Jeff Soules has written a web application for navigating the data set, hosted at <https://quasr.flatironinstitute.org/>.

REFERENCES

- BERTOLAZZI, E., GHILONI, R. & SPECOGNA, R. 2019 Efficient computation of linking number with certification, [arXiv:1912.13121](https://arxiv.org/abs/1912.13121).
- BINDEL, D., LANDREMAN, M. & PADIDAR, M. 2023 Understanding trade-offs in stellarator design with multi-objective optimization. *J Plasma Phys* **89** (5), 905890503.
- BOOZER, A.H. 2019 Curl-free magnetic fields for stellarator optimization. *Phys. Plasmas* **26** (10), 102504.
- ERIKSSON, D., PEARCE, M., GARDNER, J., TURNER, R.D. & POLOCZEK, M. 2019 Scalable global optimization via local Bayesian optimization. In *Advances in Neural Information Processing Systems* (eds. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox & R. Garnett), vol. 32. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2019/file/6c990b7aca7bc7058f5e98ea909e924b-Paper.pdf
- GIULIANI, A., WECHSUNG, F., CERFON, A., LANDREMAN, M. & STADLER, G. 2023 Direct stellarator coil optimization for nested magnetic surfaces with precise quasi-symmetry. *Phys. Plasmas* **30** (4), 042511.
- GIULIANI, A., WECHSUNG, F., CERFON, A., STADLER, G. & LANDREMAN, M. 2022a Single-stage gradient-based stellarator coil design: optimization for near-axis quasi-symmetry. *J. Comput. Phys.* **459**, 111147.
- GIULIANI, A., WECHSUNG, F., STADLER, G., CERFON, A. & LANDREMAN, M. 2022b Direct computation of magnetic surfaces in Boozer coordinates and coil optimization for quasisymmetry. *J. Plasma Phys.* **88** (4), 905880401.
- GLAS, S., PADIDAR, M., KELLISON, A. & BINDEL, D. 2022 Global stochastic optimization of stellarator coil configurations. *J. Plasma Phys.* **88** (2), 905880208.
- GOODMAN, A.G., CAMACHO MATA, K., HENNEBERG, S.A., JORGE, R., LANDREMAN, M., PLUNK, G.G., SMITH, H.M., MACKENBACH, R.J.J., BEIDLER, C.D. & HELANDER P. 2023 Constructing precisely quasi-isodynamic magnetic fields. *J Plasma Phys* **89** (5), 905890504.
- HALIR, R. & FLUSSER, J. 1998 Numerically stable direct least squares fitting of ellipses. In *Proc. 6th International Conference in Central Europe on Computer Graphics and Visualization*. WSCG, vol. 98, pp. 125–132.
- HUDSON, S.R., MONTICELLO, D.A. & REIMAN, A.H. 2001 Reduction of islands in full-pressure stellarator equilibria. *Phys. Plasmas* **8** (7), 3377–3381.
- LANDREMAN, M. 2023 PyQSC version 0.1.2. <https://github.com/landreman/pyQSC>.
- LANDREMAN, M., MEDASANI, B., WECHSUNG, F., GIULIANI, A., JORGE, R. & ZHU, C. 2021 SIMSOPT: a flexible framework for stellarator optimization. *J. Open Source Softw.* **6** (65), 3525.
- LANDREMAN, M. & SENGUPTA, W. 2018 Direct construction of optimized stellarator shapes. Part 1. Theory in cylindrical coordinates. *J. Plasma Phys.* **84** (6), 905840616.
- LANDREMAN, M., SENGUPTA, W. & PLUNK, G.G. 2019 Direct construction of optimized stellarator shapes. Part 2. Numerical quasisymmetric solutions. *J. Plasma Phys.* **85** (1), 905850103.
- LEE, B.F., PAUL, E.J., STADLER, G. & LANDREMAN, M. 2022 Stellarator coil optimization supporting multiple magnetic configurations. *Nucl. Fusion* **63** (1), 014002.
- LOBSIEN, J.-F., DREVLAK, M., KRUGER, T., LAZERSON, S., ZHU, C. & PEDERSEN, T.S. 2020 Improved performance of stellarator coil design optimization. *J. Plasma Phys.* **86** (2), 815860202.
- LOBSIEN, J.-F., DREVLAK, M., PEDERSEN, T.S., & W7-X Team. 2018 Stellarator coil optimization towards higher engineering tolerances. *Nucl. Fusion* **58** (10), 106013.
- MINER, W.H., VALANJU, P.M., HIRSHMAN, S.P., BROOKS, A. & POMPHREY, N. 2001 Use of a genetic algorithm for compact stellarator coil design. *Nucl. Fusion* **41** (9), 1185.
- PADIDAR, M., ZHU, X., HUANG, L., GARDNER, J. & BINDEL, D. 2021 Scaling Gaussian processes with derivative information using variational inference. *Adv. Neural Inf. Process. Syst.* **34**, 6442–6453.

- RODRIGUEZ, E., SENGUPTA, W. & BHATTACHARJEE, A. 2022 Phases and phase-transitions in quasisymmetric configuration space. *Plasma Phys. Control. Fusion* **64** (10), 105006.
- WECHSUNG, F., GIULIANI, A., LANDREMAN, M., CERFON, A. & STADLER, G. 2022 Single-stage gradient-based stellarator coil design: stochastic optimization. *Nucl. Fusion* **62** (7), 076034.
- ZHU, C., HUDSON, S.R., SONG, Y. & WAN, Y. 2017 New method to design stellarator coils without the winding surface. *Nucl. Fusion* **58** (1), 016008.