

# Understanding trade-offs in stellarator design with multi-objective optimization

David Bindel<sup>1</sup>, Matt Landreman<sup>2</sup> and Misha Padidar<sup>3,†</sup>

<sup>1</sup>Department of Computer Science, Cornell University, Ithaca, NY 14853, USA

<sup>2</sup>Institute for Research in Electronics and Applied Physics, University of Maryland, College Park, MD 20742, USA

<sup>3</sup>Center for Applied Mathematics, Cornell University, Ithaca, NY 14853, USA

(Received 17 April 2023; revised 17 July 2023; accepted 18 July 2023)

In designing stellarators, any design decision ultimately comes with a trade-off. Improvements in particle confinement, for instance, may increase the burden on engineers to build more complex coils, and the tightening of financial constraints may simplify the design and worsen some aspects of transport. Understanding trade-offs in stellarator designs is critical in designing high performance devices that satisfy the multitude of physical, engineering and financial criteria. In this study, we show how multi-objective optimization (MOO) can be used to investigate trade-offs and develop insight into the role of design parameters. We discuss the basics of MOO, as well as practical solution methods for solving MOO problems. We apply these methods to bring insight into the selection of two common design parameters: the aspect ratio of an ideal magnetohydrodynamic equilibrium and the total length of the electromagnetic coils.

**Keywords:** fusion plasma, plasma devices

---

## 1. Problem description

Design criteria for stellarators stem from many physical phenomena as well as many disciplines. For instance, designs must consider fusion performance metrics from multiple models and scales, engineering specifications that determine the realizability of the design, and financial timelines that set the pace of research and construction. The collection of these criteria give rise to highly constrained design optimization problems with many, potentially competing, objectives. Any design decision ultimately comes with a trade-off. For instance, an improvement in particle confinement may increase the burden on engineers to build more complex coils, and tightening of financial constraints may simplify the design and worsen some aspects of transport. When making design decisions, such as how to weight or balance objectives in an optimization or how tightly to enforce a constraint, it is critical to understand how each choice affects other aspects of the design.

† Email address for correspondence: [map454@cornell.edu](mailto:map454@cornell.edu)

To understand trade-offs that appear in stellarator design, we turn to the use of multi-objective optimization (MOO) (Ehrgott 2005; Miettinen 2012; Emmerich & Deutz 2018). MOO treats problems of the form

$$\min_{x \in \Omega} F(x) := (F_1(x), \dots, F_{n_F}(x)), \quad (1.1)$$

where  $x \in \mathbb{R}^{n_x}$  is a vector of design variables,  $\Omega \subseteq \mathbb{R}^{n_x}$  is a compact set and  $F_1, \dots, F_{n_F}$  are, typically differentiable, objectives. MOO explores the ‘Pareto optimal’ space of solutions – those which are neither better nor worse than one another and provide a distinct trade-off in the value of objectives – as well as develops an understanding of how a reduction in one objective may require worsening of another.

In this study, we introduce multi-objective optimization in the context of stellarator design. We discuss the basics of MOO, as well as a practical solution method for solving MOO problems: the  $\epsilon$ -constraint method. We also present a continuation method which leverages a local expansion of the Pareto front to explore it efficiently. We apply these MOO methods to bring insight into the selection of two common design parameters: the aspect ratio of an ideal magnetohydrodynamic (MHD) equilibrium and the total length of the electromagnetic coils. The aspect ratio has long been considered to have a trade-off with the degree to which a configuration is quasi-symmetric. A conjecture by Garren & Boozer (1991) states that quasi-symmetry can only be achieved through second order in the inverse aspect ratio, and hence that decreasing aspect ratio would result in a worsening of the degree of quasi-symmetry. Our numerical experiments suggest that this is indeed the case, but the trade-off is modest. We also explore the relationship between the total allowable length of the coil pack and the ability for the coils to reproduce a target magnetic field. There is a natural trade-off here since longer more complex coils can reproduce more intricate fields than the same number of shorter coils (Landreman 2017; Wechsung *et al.* 2022b).

Upon computing the Pareto front, a stellarator designer can visualize the trade-off between multiple objectives and decide how to balance this trade-off in their design. Since the Pareto front provides a set of solutions rather than a single solution, an engineer or physicist can choose between multiple Pareto optimal designs by evaluating them under additional design criteria.

This paper is structured as follows: in §2, we review MOO, as well as two methods for solving MOO problems. Subsequently, in §3, we apply the optimization methods to explore two trade-off problems in stellarator design. Finally, in §4, we look beyond our case studies to the more general use of MOO for stellarator design and discuss future directions.

We use bold characters, such as  $x$ , to denote vectors. We denote a vector  $x$  with entry  $j$  removed as  $\bar{x}^j = (x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_{n_x})$ . We compare vectors using vector inequalities: if  $x \leq y$ , then  $x_j \leq y_j$  for  $j = 1, \dots, n_x$ .

## 2. Multi-objective optimization

Multi-objective optimization problems typically do not have a single solution, but rather an entire set of solutions. These are points which balance a trade-off between objectives: improvement in one objective implies a worsening of the other. The set of solutions to a MOO problem is called the Pareto front. In this section, we briefly formalize this notion, before introducing practical solution techniques for MOO problems.

The idea of a trade-off can be formalized through the notion of non-dominance. A point  $x \in \Omega$  dominates  $y \in \Omega$ , with respect to the objectives  $F$ , if  $F(x) \leq F(y)$  and  $F(x) \neq F(y)$ . Informally,  $x$  is at least as good as  $y$  in all the objectives. If  $x$  is not

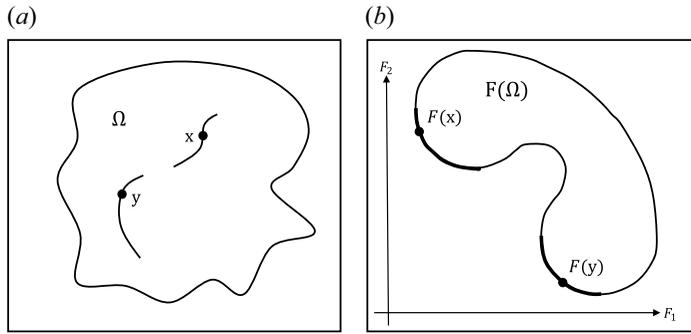


FIGURE 1. (a) Visualization of two points on the efficient set (black curves). (b) Pareto front is indicated as the thicker black lines on the edge of the image of  $\Omega$ ,  $F(\Omega)$ . The Pareto front is disconnected into two components, each of which is a mapping under  $F$  of one of the efficient curves in panel (a). Each efficient curve is called a ‘branch’ of solutions.

dominated by any other point in  $\Omega$ , then  $x$  is called efficient, or Pareto optimal, and  $F(x)$  is a non-dominated point. Plainly,  $x$  is efficient if there is no other point that is at least as good in all objectives and better in at least one objective. For example, if  $x$  is efficient, then another point  $y$  could still satisfy  $F_1(y) < F_1(x)$ , but then it would have to hold that  $F_j(y) > F_j(x)$  for some  $j \neq 1$ . The set of all efficient points forms the efficient set and the set of all non-dominated points forms the non-dominated set, which is more commonly known as the Pareto front. Finding the Pareto front is the goal of multi-objective optimization. A visualization of an efficient set and Pareto front are shown in figure 1. Intuitively, efficient points are those which cannot be compared by their objective values alone. For instance,  $x$  may have less complex coil shapes than  $y$ , but may not reproduce a target magnetic field as well. Based off of these facts alone, neither  $x$  nor  $y$  is a dominant configuration and additional information would be needed to determine which configuration is preferable.

In the absence of convex objectives, finding globally Pareto efficient solutions is difficult. Instead, we settle for searching for weakly, locally efficient points, which is all most MOO algorithms can guarantee in this setting, given a finite sample size. Here,  $x \in \Omega$  is weakly efficient (Ehrgott 2005) if there is no  $y \in \Omega$  such that  $y$  strictly dominates  $x$ ,  $F(y) < F(x)$ . The set of weakly efficient points contains all efficient points but can also contain ‘extensions’ from the edge of the Pareto front, see figure 2. A point  $x \in \Omega$  is locally efficient (Emmerich & Deutz 2018) if there exists a non-empty open ball containing  $x$  such that  $F(x)$  is non-dominated in the intersection of the ball and  $\Omega$ . A local Pareto front can be found by solving (1.1) with local optimization methods, rather than globally solving (1.1) to compute the true Pareto front. In figure 2, the dashed line indicates a local efficient set, which is Pareto optimal within the shaded region. We find that finding weakly, locally efficient solutions is enough to provide insight into the trade-offs of interest.

MOO problems are most commonly solved by either directly applying MOO algorithms (Deb *et al.* 2002; Knowles 2006; Wagner *et al.* 2010; Daulton *et al.* 2022; Chang & Wild 2023) or by reformulating the MOO problem as a series of scalar optimization problems that can be solved with scalar optimization methods. Algorithmic approaches attempt to find and explore the efficient set by taking steps to jointly minimize a combination of the objectives. These methods are a practical option for problems with many objectives, since the algorithm can weigh how to efficiently explore the high dimensional Pareto front. Scalarization methods, however, allow for more user involvement and incorporation of

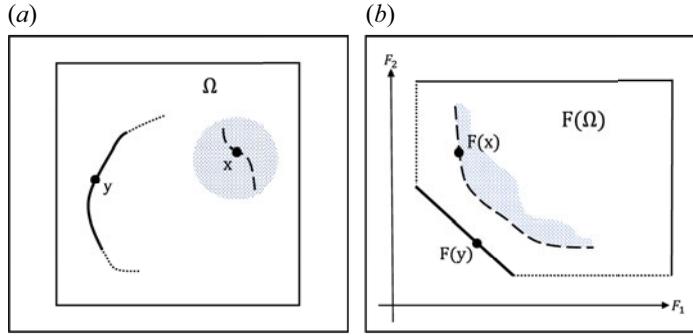


FIGURE 2. (a) The efficient set (thick black line), weakly efficient set (dotted lines) and locally efficient set (dashed line) in the domain  $\Omega$ . The locally efficient points are Pareto optimal within the shaded region around the curve. (b) The Pareto front is indicated as the thicker black line on the edge of the image of  $\Omega$ ,  $F(\Omega)$ . The weak Pareto front (dotted lines) extends from the Pareto front and is not strictly dominated by any other point. A local Pareto front is shown as the dashed line, which is the Pareto front of the shaded region.

problem specific information, such as derivatives, which many MOO algorithms do not use. Scalarization methods can be formulated and solved efficiently because they rely on scalar optimization methods, and allow the user to control the exploration of the efficient set. A particularly popular scalarization approach is the ‘linearization method’ (Ehrgott 2005). The linearization method finds efficient points by solving the linearized problem

$$\min_{x \in \Omega} \sum_{i=1}^{n_F} \alpha_i F_i(x), \tag{2.1}$$

where the user-selected weight vector,  $\alpha$ , must be non-negative and sum to one. The linearization method, however, has two major drawbacks (Das & Dennis 1996): (1) if the Pareto front is non-convex, then no weight vector  $\alpha$  exists such that the solution to the linearized problem, (2.1), lies on the non-convex portion of the Pareto front; (2) a uniformly spaced selection of weight vectors does not uniformly explore the Pareto front. For these reasons, it is worthwhile to look beyond the linearization method when solving non-convex MOO problems.

In the remainder of this section, we discuss a scalarization method, the  $\epsilon$ -constraint method, which overcomes the drawbacks of the linearization method and is particularly useful for trade-offs in stellarator design. We also discuss a continuation method for locally exploring the Pareto front efficiently.

### 2.1. The $\epsilon$ -constraint method

Multi-objective optimization problems can be transformed into a series of scalar optimization problems via scalarization methods. One such scalarization method, the  $\epsilon$ -constraint method (Ehrgott 2005), finds (weakly) Pareto optimal points by minimizing a single objective subject to upper bound inequality constraints on all other objectives, with an upper bound parameter  $\epsilon \in \mathbb{R}^{n_F-1}$ . To find points on the Pareto front, the  $\epsilon$ -constraint method solves problems of the form

$$\left. \begin{aligned} & \min_{x \in \Omega} F_j(x) \\ \text{such that } & F_1(x), \dots, F_{j-1}(x), F_{j+1}(x), \dots, F_m(x) \leq \epsilon \end{aligned} \right\}, \tag{2.2}$$

where the user selects the index  $j$  and vector  $\epsilon$ . The  $\epsilon$ -constraint method theoretically guarantees that it can be used to find any point on the Pareto front: if  $\mathbf{x}$  is a (local) solution to (2.2), then it is weakly (locally) efficient. Furthermore,  $\mathbf{x}$  is efficient if and only if it is a solution to (2.2) for all choices of  $j \in \{1, \dots, n_F\}$  (Ehrgott 2005).

The  $\epsilon$ -constraint method is particularly relevant to stellarator optimization since we often have bounds on the range of the objectives, which allows us to easily select the upper bound parameter,  $\epsilon$ . For example, devices often have an aspect ratio between 3 and 10, a coil must be at least as long as the minor circumference of the plasma, and it should not be longer than a few times that. By selecting  $\epsilon$ , targeted areas of the Pareto front can be searched. This is more interpretable than using other scalarization techniques, such as a weighted sum of objectives, which requires selecting abstract weights that do not have any clear relation to the Pareto front.

## 2.2. Continuation methods

Continuation methods (Hillermeier *et al.* 2001; Schtze, Dell'Aere & Dellnitz 2005; Gkaragkounis *et al.* 2018; Peitz, Ober-Blöbaum & Dellnitz 2019; Vasilopoulos *et al.* 2021) in MOO explore the Pareto front locally around a given efficient point. By using optimality conditions and the implicit function theorem (Krantz & Parks 2002), continuation methods build local models of the Pareto front, which they use to estimate nearby efficient points. Local exploration of the Pareto front with continuation methods comes at a relatively low computational cost when compared to restarting a scalarization solver, such as the  $\epsilon$ -constraint method, from scratch. However, when combined, continuation methods and scalarization solvers make a powerful pair: the scalarization methods find efficient points over distinct parts of the Pareto front and the low-cost continuation method is used to fill in the space between points. In this section, we discuss a predictor–corrector continuation method based off of the  $\epsilon$ -constraint method, which first predicts an estimate of an efficient point using a Taylor expansion and subsequently uses the  $\epsilon$ -constraint method to correct the predictions, see Algorithm 1. While other predictor–corrector continuation schemes (Schütze *et al.* 2005; Peitz *et al.* 2019; Vasilopoulos *et al.* 2021) write the expansion in terms of weights, or target points, we expand the Pareto front in terms of  $\epsilon$  because it is easily interpreted and allows for the  $\epsilon$ -constraint method to be used in the corrector step.

The key idea of the predictor step is to notice that the  $\epsilon$ -constraint method allows us to parametrically describe weakly efficient points in terms of  $\epsilon$ , i.e. we can write weakly efficient points as  $\mathbf{x}(\epsilon)$ . In fact, as we will show, under reasonable conditions,  $\mathbf{x}(\epsilon)$  is a continuously differentiable map. The predictor step leverages this parametric representation to approximate the efficient set by a Taylor expansion. Given an efficient point  $\mathbf{x}$ , which is a solution to (2.2) with parameter  $\epsilon$ , the predictor step locally approximates the efficient set by

$$\mathbf{x}(\epsilon + \Delta\epsilon) \approx \mathbf{x}(\epsilon) + \nabla_{\epsilon}\mathbf{x}(\epsilon)^T \Delta\epsilon, \quad (2.3)$$

where  $\Delta\epsilon$  is a small change in  $\epsilon$ , and  $\nabla_{\epsilon}\mathbf{x}$  is the Jacobian of  $\mathbf{x}(\epsilon)$  with respect to  $\epsilon$ . Though less accurate, a zeroth-order Taylor expansion could be used in place of the first-order expansion. A zeroth-order expansion simplifies the predictor–corrector method, Algorithm 1, by using the last known efficient point  $\mathbf{x}(\epsilon)$  as a prediction of a nearby efficient point  $\mathbf{x}(\epsilon + \Delta\epsilon)$ , and is equivalent to warm starting the  $\epsilon$ -constraint method with  $\mathbf{x}(\epsilon)$ . The zeroth-order predictor is useful when the Jacobian  $\nabla_{\epsilon}\mathbf{x}$  is intractable to approximate or the cost of running the  $\epsilon$ -constraint method is low. However, if  $\nabla_{\epsilon}\mathbf{x}$  can be approximated, then the first-order expansion will give more accurate predictions of  $\mathbf{x}(\epsilon + \Delta\epsilon)$  than the zeroth-order expansion, and should be used. A visualization of the

---

**Algorithm 1:**  $\epsilon$ -Constraint predictor–corrector method

---

**Input:** Efficient point  $x_0$ , small value  $\Delta\epsilon \in \mathbb{R}^{n_F-1}$ , index  $j \in \{1, \dots, n_F\}$ , backtracking parameter  $\gamma \in (0, 1)$ , small tolerance  $\tau_\Delta > 0$ .

**Result:** Approximate efficient points  $\{x_k\}$ .

- 1 **for**  $k = 0, 1, \dots$  **do**
- 2 Set  $\epsilon_k = \bar{F}^j(x_k)$ ;
- 3 Compute the Lagrange multiplier  $\lambda_k$  from (2.8) using  $\epsilon = \epsilon_k$ ;
- 4 Compute the Hessian approximations for  $F$  at  $x_k$ ;
- 5 Build the Jacobian  $D(x_k, \lambda_k)$  from the Hessians approximations and gradients;
- 6 Solve for  $\nabla_\epsilon x$  using (2.7);
- 7 Predict: Approximate  $x(\epsilon_k + \Delta\epsilon)$  using (2.3);
- 8 **while**  $F$  cannot be evaluated at the predicted point **do**
- 9 Set  $\Delta\epsilon \rightarrow \gamma \Delta\epsilon$ ;
- 10 **if**  $\Delta\epsilon < \tau_\Delta$  **then**
- 11 | Terminate Algorithm.
- 12 **else**
- 13 | Predict: Approximate  $x(\epsilon_k + \Delta\epsilon)$  using (2.3)
- 14 **Correct:** Compute  $x_{k+1}$  with the  $\epsilon$ -constraint method by solving (2.2) using  $\epsilon = \epsilon_k + \Delta\epsilon$ , and starting the optimization from the predicted point;

---

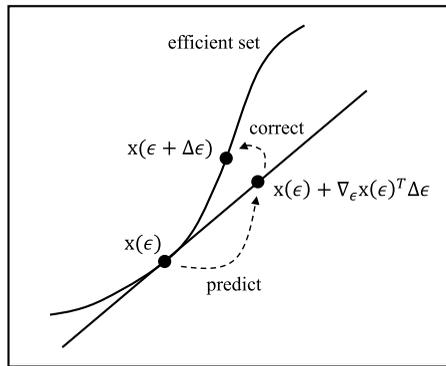


FIGURE 3. Visualization of the predictor–corrector continuation method.

first-order predictor corrector method is shown in figure 3. In the remainder of this section, we will discuss conditions which determine when the expansion, (2.3), exists and a method of computing the Jacobian  $\nabla_\epsilon x$ .

The existence of the expansion, (2.3), and the computation of the Jacobian  $\nabla_\epsilon x$  rely on applying the implicit function theorem (Krantz & Parks 2002) to the first order necessary conditions for optimality of (2.2), the Karush–Kuhn–Tucker (KKT) conditions (Nocedal & Wright 1999). The KKT conditions are necessary conditions for a point  $x$  to be an optimal solution to an inequality constrained optimization problem, such as the  $\epsilon$ -constraint problem (2.2). In the case of the  $\epsilon$ -constraint method, points satisfying the KKT conditions for (2.2) are also weakly locally efficient.

Suppose that  $x$  is an optimal solution to (2.2), and  $\mathcal{A} \subseteq \{1, \dots, j - 1, j + 1, \dots, n_F\}$  is the set of  $m$  active constraints, i.e. the set of constraints  $i$  for which  $F_i(x) = \epsilon_i$ . The

KKT conditions with the linear independence constraint qualification (LICQ) (Nocedal & Wright 1999) state that there exist  $\lambda \in \mathbb{R}^m$  such that  $\lambda \geq 0$  and

$$\nabla F_j(\mathbf{x}) + \sum_{i \in \mathcal{A}} \lambda_i \nabla F_i(\mathbf{x}) = 0, \tag{2.4}$$

$$\lambda_i (F_i(\mathbf{x}) - \epsilon_i) = 0 \quad \forall i \in \mathcal{A}, \tag{2.5}$$

$$F_i(\mathbf{x}) \leq \epsilon \quad \forall i \neq j. \tag{2.6}$$

Notice that (2.4) and (2.5) show that if a constraint is active and  $\lambda_i > 0$ , then the solution of the  $\epsilon$ -constraint method is dependent on  $\epsilon_i$  in the sense that a small change in  $\epsilon_i$  could change the solution to the system of equations. In this case, it is reasonable to write  $\mathbf{x}$  as a function of  $\epsilon_{\mathcal{A}} = (\epsilon_i)_{i \in \mathcal{A}}$  to stress the dependence of solutions on the parameters,  $\mathbf{x}(\epsilon_{\mathcal{A}})$ . The implicit function theorem gives us the mathematical leeway to do so.

Informally, the implicit function theorem states that if there exists a point  $\mathbf{x}^*$ , with Lagrange multipliers  $\lambda^*$  and constraint parameters  $\epsilon_{\mathcal{A}}^*$ , which satisfy the KKT conditions, and if the Jacobian,  $D(\mathbf{x}^*, \lambda^*)$ , of the left-hand side of (2.4) and (2.5) with respect to  $(\mathbf{x}, \lambda)$  is invertible, then  $\mathbf{x}(\epsilon_{\mathcal{A}})$ ,  $\lambda(\epsilon_{\mathcal{A}})$  are continuously differentiable functions of  $\epsilon_{\mathcal{A}}$  around  $\epsilon_{\mathcal{A}}^*$ . Taking the form of a continuously differentiable function, the space of solutions to the  $\epsilon$ -constraint problem (2.2) can be approximated with the first-order Taylor expansion, (2.3). The Jacobian  $\nabla_{\epsilon} \mathbf{x}$  can be computed by solving the system

$$D(\mathbf{x}^*, \lambda^*) \begin{bmatrix} \nabla_{\epsilon} \mathbf{x} \\ \nabla_{\epsilon} \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \lambda^* \end{bmatrix}. \tag{2.7}$$

The system (2.7) is derived by treating  $\mathbf{x}$  and  $\lambda$  as a function of  $\epsilon$  and differentiating (2.4) and (2.5) with respect to  $\epsilon$ . Importantly,  $D$  depends on second derivatives of the objective functions. If  $D$  is invertible, then the implicit function theorem guarantees existence of the expansion (2.3), and a unique solution exists to (2.7) (Krantz & Parks 2002).

Invertibility of the Jacobian matrix  $D$  occurs naturally when  $\mathbf{x}, \lambda$  satisfy the KKT conditions, and  $\mathbf{x}$  is a strict local minimum:  $\nabla^2 F_j(\mathbf{x}) + \sum_{i \neq j} \lambda_i \nabla^2 F_i(\mathbf{x})$  is positive definite over the orthogonal subspace to  $\{\nabla F_i\}$  for all  $i \in \mathcal{A}$  such that  $\lambda_i > 0$ . However, the expansion does not exist if there are no active constraints: if all  $F_i(\mathbf{x}) < \epsilon_i$ , then a small variation in any  $\epsilon_i$  should not affect the solution. For example, in figure 1, the predictor–corrector method cannot be used beyond the tips of either branch of the efficient set since the KKT conditions cease to hold. Furthermore, if some constraint  $i$  is active but the associated Lagrange multiplier  $\lambda_i = 0$ , then a small increase in  $\epsilon_i$  will not change the solution to (2.2), but a small decrease would. In this case,  $\lambda_i$  is only sub-differentiable (Clarke 1990) with respect to  $\epsilon_i$ . While an expansion can be derived in this case, we ignore it for simplicity, as this case only appears at the edges of the efficient set or at saddle points of  $F_j$ . We are primarily concerned with the case when all active constraints have strictly positive Lagrange multipliers, since in this case, there is a trade-off to be made between objectives: reduction in  $F_j$  comes at the cost of increase in  $\epsilon_i$ .

Now that we have established when a local expansion of the efficient set exists, and how to compute it, we can discuss the predictor–corrector method. The predictor step of the predictor–corrector method takes as input an efficient point  $\mathbf{x}$  and a step size  $\Delta \epsilon$ . Subsequently, we compute the Lagrange multiplier  $\lambda$ . If some components of  $\lambda$  are strictly positive, then we compute  $D(\mathbf{x}, \lambda)$ , solve (2.7) for  $\nabla_{\epsilon} \mathbf{x}$ , then predict a weakly efficient point via the expansion, (2.3).

When a simulation is used to compute  $F$ , it may not be possible to evaluate  $F$  at the predicted point since the simulation can ‘fail’. For example, the Variational Moments

Equilibrium Code ( $\nabla$ MEC) may fail to converge if the predicted point represents an ideal MHD equilibrium with self-intersecting flux surfaces. To proceed with the corrector step, the predicted point must evaluate under  $F$ . To ensure that such a point can be found, the predicted point can be modified using a backtracking line search (Nocedal & Wright 1999). The backtracking procedure iteratively reduces  $\Delta\epsilon$  to  $\gamma\Delta\epsilon$  for  $\gamma \in (0, 1)$  and predicts the value of  $\mathbf{x}(\epsilon + \Delta\epsilon)$ , until  $F$  can be evaluated at the predicted point  $\mathbf{x}(\epsilon + \Delta\epsilon)$  or  $\Delta\epsilon$  is smaller than some tolerance  $\tau_\Delta$ .

For small enough  $\Delta\epsilon$ , the prediction will typically be a very good estimate of an efficient point. Nonetheless, it is helpful to finely resolve the point, particularly as this allows us to compute the subsequent predictor step accurately. The corrector step corrects the prediction by simply solving the  $\epsilon$ -constraint problem, (2.2), with new constraint parameters  $\epsilon + \Delta\epsilon$ . This method is detailed in Algorithm 1.

### 2.2.1. Computation of the Lagrange multiplier

In practice, we rarely have a point  $\mathbf{x}$  that exactly satisfies any constraint with equality. Typically, an inequality constraint may approximately be satisfied with equality, i.e.  $\bar{F}^j(\mathbf{x}) = \epsilon - \nu$  for a small value of  $\nu > 0$ . In this case, theoretically the constraint is not active and so the Lagrange multipliers are equal to zero. However, the numerics suggest that the constraint is active and the Lagrange multiplier should not be zero. In practice, we determine the active set  $\mathcal{A}$  as the set of all constraints with  $\nu_i < \tau$  for some small tolerance  $\tau$ . The Lagrange multiplier can then be computed by solving the KKT conditions for a perturbed version of (2.2) where the constraint right hand side,  $\epsilon_i$ , is shifted to  $\epsilon_i - \nu$  for all constraints in  $\mathcal{A}$ . In this way, the constraints hold with exact equality, since  $\bar{F}^j(\mathbf{x}) = \epsilon - \nu$ , and we are justified in computing the potentially non-zero Lagrange multipliers. The Lagrange multipliers are then estimated as the solution to

$$\min_{\lambda \geq 0} \left\| \nabla F_j(\mathbf{x}) + \sum_{i \in \mathcal{A}} \lambda_i \nabla F_i(\mathbf{x}) \right\|^2. \quad (2.8)$$

Importantly, this procedure affects the prediction step in that the expansion is taken around  $\epsilon + \nu$ . However, this is accounted for in step 1 of Algorithm 1.

## 3. Numerical experiments

In this section, we use MOO to bring insight into the selection of two common stellarator design parameters: the aspect ratio of an ideal MHD equilibrium and the total length of the electromagnetic coils.

Our first experiment attempts to determine if there is a trade-off between achieving precise quasi-symmetry in a ‘stage-one’ stellarator design (an ideal MHD equilibrium) and having low aspect ratio. The Garren–Boozer conjecture (Garren & Boozer 1991) suggests that exact quasi-symmetry is only possible at high aspect ratio. In addition, precise quasi-symmetry has been achieved throughout a volume in high aspect ratio stellarators (Giuliani *et al.* 2022b; Landreman 2022; Landreman & Paul 2022; Wechsung *et al.* 2022b). However, it is unclear at what rate quasi-symmetry decays as the aspect ratio is increased or if this trade-off applies to precise quasi-symmetry at all. Our first experiment answers the following question: ‘To what extent does the aspect ratio limit the degree to which quasi-symmetry can be achieved throughout a volume?’

Our second experiment considers a trade-off in the ‘stage-two’ design problem, where coils are optimized to fit a target magnetic field. A target magnetic field can be recreated arbitrarily well by coils which have no constraint on their length. However, when restricted

to have a reasonably short length for engineering purposes, coils may not be able to reproduce a target magnetic field. In this problem, we aim to understand how reduction in the allowable coil length worsens the reproduction of the target magnetic field.

3.1. *Problem 1: the aspect ratio and quasi-symmetry trade-off*

We seek a plasma boundary shape, parametrized by  $\mathbf{x}$ , of an ideal MHD equilibrium for a quasi-helical (QH) stellarator configuration that has minimal aspect ratio  $A(\mathbf{x})$  and deviation from quasi-symmetry  $Q_{M,N}(\mathbf{x})$ ,

$$\min_{A(\mathbf{x}) \in [A_l, A_u]} (A(\mathbf{x}), Q_{M,N}(\mathbf{x})). \tag{3.1}$$

Including bound constraints on the aspect ratio  $A_l = 3, A_u = 10$ , restricts our decision space to a realistic range of configurations. For convenience, we collect the  $n_F = 2$  objectives into the vector  $\mathbf{F}(\mathbf{x}) := (A(\mathbf{x}), Q_{M,N}(\mathbf{x}))$ .

The violation of quasi-symmetry is defined as the quasi-symmetry-ratio-residual objective employed by Landreman & Paul (2022). The objective measures the departure from quasi-symmetry throughout the plasma volume as the sum of flux surface averages across surfaces  $(s_1, \dots, s_{n_s})$ ,

$$Q_{M,N}(\mathbf{x}) = \sum_{s_j} \left\langle \left( \frac{1}{B^3} [(N - lM)\mathbf{B} \times \nabla \mathbf{B} \cdot \nabla \psi - (MG - NI)\mathbf{B} \cdot \nabla \mathbf{B}] \right)^2 \right\rangle. \tag{3.2}$$

The helicity parameters,  $M, N$ , determine the type of quasi-symmetry that  $Q_{M,N}$  measures:  $M = 1, N = 0$  being quasi-axisymmetry, and  $M = 1, N = \pm kn_{fp}$  with non-zero integers  $k$  being quasi-helical symmetry (Landreman & Paul 2022). For our experiments, the helicity parameters were set to  $M = 1, N = -n_{fp}$  for the  $n_{fp} = 4$  field period magnetic field. Discretization of the flux surface average over the poloidal and toroidal angles results in an objective with sum-of-squares structure,  $Q_{M,N}(\mathbf{x}) = \sum_{i=1}^{n_O} Q_{M,N,i}^2(\mathbf{x})$ , where  $Q_{M,N,i}$  measures the violation of quasi-symmetry at a point throughout the volume. For a given plasma boundary,  $\mathbf{x}$ , we compute the aspect ratio using the definition in the VMEC code,

$$A(\mathbf{x}) = R_{\text{major}}(\mathbf{x})/R_{\text{minor}}(\mathbf{x}) \quad R_{\text{minor}} = \sqrt{\bar{C}_A(\mathbf{x})/\pi} \quad R_{\text{major}(\mathbf{x})} = \frac{V(\mathbf{x})}{2\pi^2 R_{\text{minor}}^2(\mathbf{x})}, \tag{3.3a-c}$$

where  $\bar{C}_A$  is the average cross-sectional area of the surface and  $V$  is the volume enclosed by the surface.

The decision variables,  $\mathbf{x}$ , are Fourier amplitudes that describe the shape of the plasma boundary. The plasma boundary is represented in terms of the standard cylindrical coordinates  $(R, \phi, Z)$  where  $R, Z$  are parametrized as a Fourier series in the poloidal and toroidal angles  $\theta$  and  $\phi$ ,

$$\left. \begin{aligned} R(\theta, \phi) &= \sum_{n=0}^{n_{sm}} R_{0,n} \cos(-n_{fp}n\phi) + \sum_{m=1}^{n_{sm}} \sum_{n=-n_{sm}}^{n_{sm}} R_{m,n} \cos(m\theta - n_{fp}n\phi), \\ Z(\theta, \phi) &= \sum_{n=1}^{n_{sm}} Z_{0,n} \sin(-n_{fp}n\phi) + \sum_{m=1}^{n_{sm}} \sum_{n=-n_{sm}}^{n_{sm}} Z_{m,n} \sin(m\theta - n_{fp}n\phi). \end{aligned} \right\} \tag{3.4}$$

The number of modes describing the surface  $n_{sm} = \{0, 1, 2, \dots\}$  can be increased to achieve more intricate boundary representations. Field period symmetry with  $n_{fp}$  periods

and stellarator symmetry have been assumed. The major radius  $R_{0,0}$  is held fixed throughout the optimization to fix the scale of the design. The Fourier amplitudes are collected into the decision variable via  $\mathbf{x} = (R_{0,1}, \dots, Z_{0,0}, \dots)$ . The total number of decision variables satisfies  $n_x = 4n_{\text{sm}}^2 + 4n_{\text{sm}}$ .

Numerical experiments were performed using SIMSOPT (Landreman *et al.* 2021) to handle variables, compute objectives and interface with VMEC which computed the ideal MHD equilibria from the plasma boundary representation. The  $\epsilon$ -constraint method was used to find points along the Pareto front where quasi-symmetry was set to be the target function for minimization while the aspect ratio was constrained by  $\epsilon$ . The bound constraints on the aspect ratio made specifying  $\epsilon$  straightforward:  $\epsilon$  was set to linearly spaced values between  $A_l$  and  $A_u$ . The  $\epsilon$ -constraint problem (2.2) was solved by reformulating the constrained optimization problem with a quadratic penalty method (Nocedal & Wright 1999), see § A. The penalty parameter was increased from an initial value of 1 by a factor of 10 at each iteration, and the quadratic penalty sub-problems were solved by applying a Gauss–Newton optimization routine. The optimization was run until the change in  $Q_{M,N}(\mathbf{x})$  was of the order of machine precision,  $10^{-16}$ , or the norm of the gradient reached  $10^{-10}$ . Upon taking a step with size of the order of  $10^{-10}$ , the optimization algorithm would terminate due to significant computational noise in the output of VMEC. SIMSOPT was used to compute forward difference gradients via MPI-based concurrent function evaluations. To avoid local minima, the number of Fourier modes,  $n_{\text{sm}}$ , was increased iteratively from 1 to 5, reaching  $n_x = 120$  variables. The  $\epsilon$ -constraint problem was solved with the penalty method after each increase of  $n_{\text{sm}}$ .

To reduce the computational overhead of using the  $\epsilon$ -constraint method, we used the predictor–corrector method, introduced in § 2.2, to explore the Pareto front between solutions of the  $\epsilon$ -constraint method. For the high dimensional decision space of interest, the full Hessians  $\nabla^2 Q_{M,N}$  and  $\nabla^2 A$  are too expensive to compute with finite differences. Instead, we approximated the Hessian  $\nabla^2 Q_{M,N}$  with a Gauss–Newton Hessian approximation and use a diagonal second-order central difference approximation to  $\nabla^2 A$ .

Figure 4 shows the Pareto front for the problem (3.1). Figure 5 shows three-dimensional renderings of the three solutions highlighted by the square, star and diamond markers in figure 4, as well as contour plots of the field strength in Boozer coordinates. From figure 4, it is clear that low values of the quasi-symmetry violation,  $Q_{1,-4}(\mathbf{x}) < 10^{-3}$ , can be achieved at all aspect ratios in our range. There is a slight trend indicating that quasi-symmetry may be achieved more precisely at higher aspect ratios. Nonetheless, by viewing the contour plots of the magnetic field strength in Boozer coordinates in figure 5, we see that configurations from all parts of the Pareto front have visibly precise quasi-symmetry.

While it seems that precise quasi-symmetry can be achieved at all aspect ratios considered, it is not clear how well the reduction in quasi-symmetry translates to an improvement in particle confinement. After all, achieving quasi-symmetry is a proxy for the true goal of achieving good confinement. To this end, we computed the fraction of alpha particles lost from 10 distinct Pareto optimal configurations, with aspect ratios 3, 3.5, 4, 4.5, 4.9, 5.6, 6.0, 8.5, 8.7 and 9. To compute the losses, we scaled the Pareto optimal configurations to the ARIES-CS reactor (Najmabadi *et al.* 2008) scale (1.7 m minor radius and  $B_{0,0}(s=0) = 5.7$  T field strength on axis), and traced 5000 particles born on the  $s = 1/4$  flux surface as well 5000 particles born on the  $s = 1/2$  flux surface until a terminal time of 0.1 seconds. The reactor scaling and particle count were also used by Bader *et al.* (2019) and Landreman & Paul (2022). Particles were traced according to the vacuum guiding center equations in Boozer coordinates and were deemed lost if they crossed the  $s = 1$  flux surface. The loss fractions are shown in table 1. Not a single

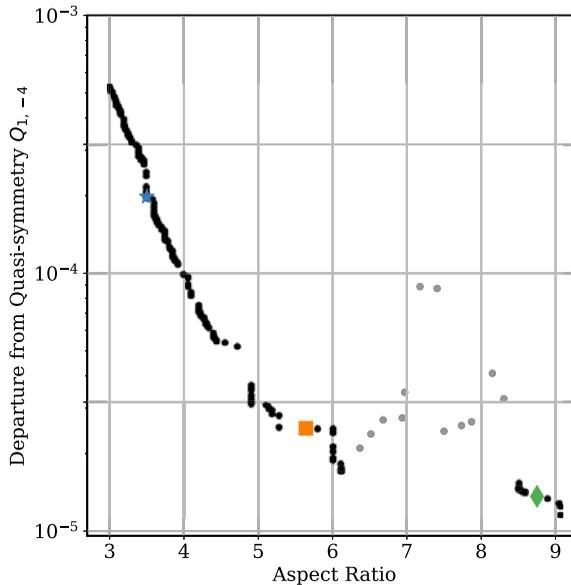


FIGURE 4. Pareto front of the aspect ratio and quasi-symmetry objectives over the domain  $A_l \leq A(\mathbf{x}) \leq A_u$  (black points). The configurations corresponding to the blue star, orange square and green diamond markers are plotted in figure 5. We find that the efficient set undergoes a branch change as the aspect ratio crosses the large gap from  $\approx 6.11$  to  $\approx 8.5$ . The configurations in this range are not Pareto optimal and are shown in grey.

particle born on the  $s = 1/4$  flux surface was lost from nine of the ten configurations, and the remaining configuration, with aspect ratio 5.6, lost only 0.22 % of the alpha particles. A slightly larger fraction of the alpha particles born on the  $s = 1/2$  flux surface were lost, between 0 % and 2.6 % for each configuration. The loss fraction of alpha particles born on the  $s = 1/2$  flux surface increases slightly as the aspect ratio increases up until  $A \approx 6$ , then drops to approximately zero on the right tail of the Pareto front. It seems that there is no clear trend between the aspect ratio of the Pareto optimal designs and the loss fractions for particles born on either of the two surfaces. Thus, while there is a slight trade-off between quasi-symmetry and aspect ratio, the trade-off between confinement and aspect ratio may be more complex.

Spatially, we may expect or hope that the set of desirable configurations forms a compact, connected region. However, we find the contrary – the efficient set is not connected. Qualitatively, the large gap in the Pareto front in figure 4 hints that the efficient set is undergoing a ‘branch change’. Branched curves, or branches, are parametric families of solutions to (2.4)–(2.6) which can be computed by continuation. Two branches can either be connected by a non-smoothness or be entirely disconnected sets, as in the case of figure 1. Branch changes are reflected in the Pareto front via a non-smoothness or gap where the Pareto front transitions from representing points on one branch to representing points on another branch. Quantitatively, we measure the branch change by evaluating the Lagrange multipliers at the left edge of the gap in the Pareto front (roughly  $A = 6.11$ ) in figure 4; a branch change is indicated by  $\lambda = 0$ , since this implies that the aspect ratio can be increased without causing a reduction in the quasi-symmetry. Indeed, we find that the Lagrange multipliers are approximately zero and hence that the efficient set corresponding to the small piece of the Pareto front, with aspect ratio greater than or

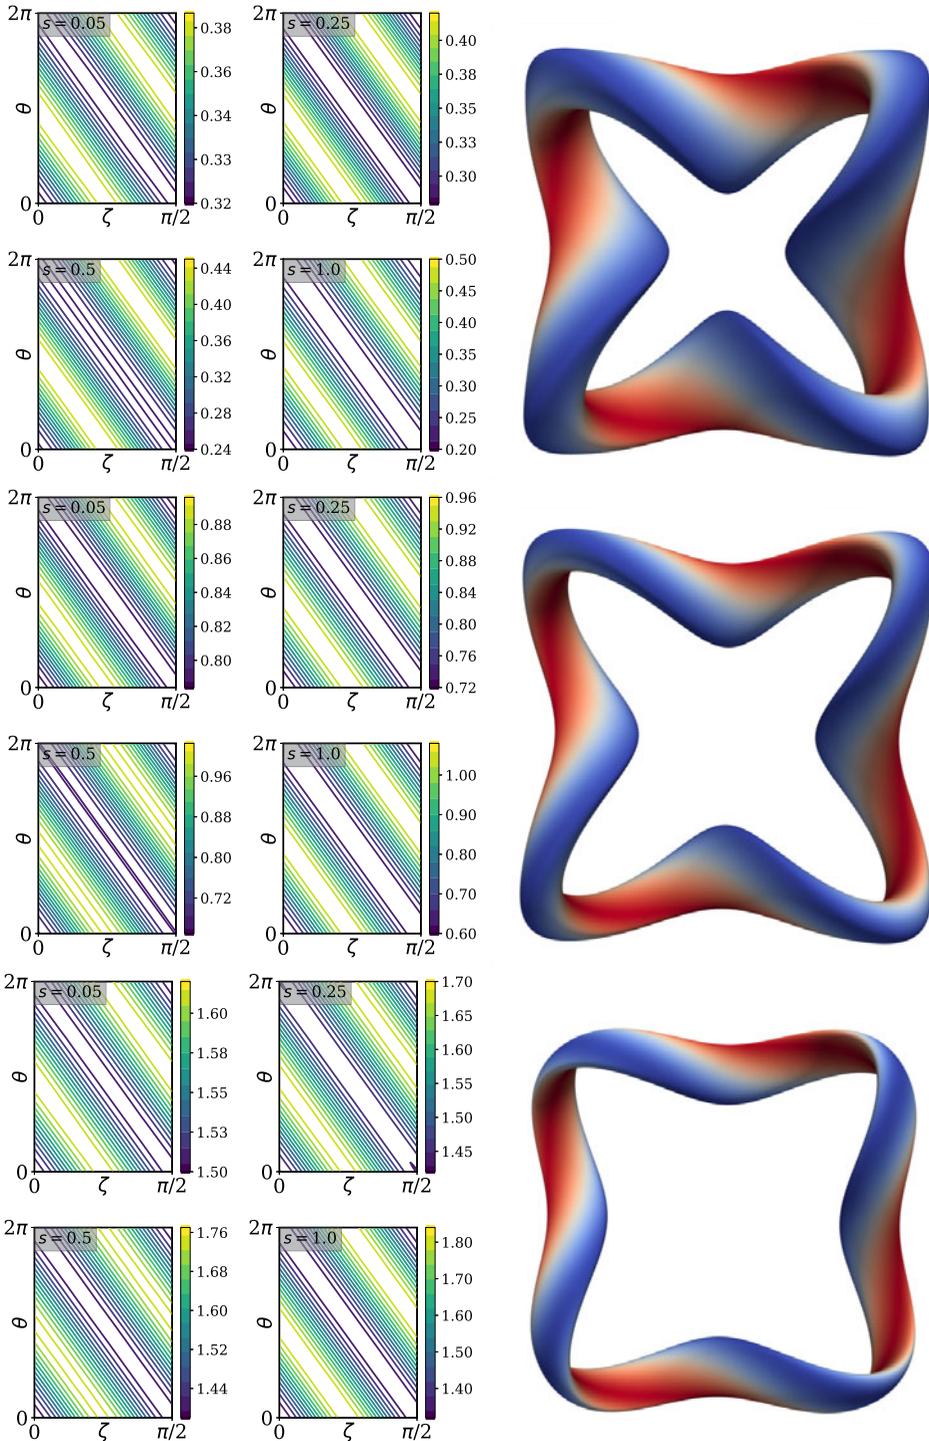


FIGURE 5. (Left columns) Contour plots of the magnetic field strength in Boozer coordinates  $(\theta, \zeta)$  on four flux surfaces  $s = 0, 0.25, 0.5, 1.0$  for three Pareto optimal configurations denoted by a star (aspect 3.5), square (aspect 5.6) and diamond (aspect 8.8) in figure 4. (Right column) Three-dimensional (3-D) renderings of corresponding Pareto optimal designs. The colour of the 3-D configurations indicates the field strength, where red is stronger.

		Aspect Ratio									
		3.0	3.5	4.0	4.5	4.9	5.6	6.0	8.5	8.7	9.0
Surface $s$	1/4	0.0	0.0	0.0	0.0	0.0	0.0022	0.0	0.0	0.0	0.0
	1/2	0.010	0.015	0.018	0.021	0.023	0.026	0.026	0.0002	0.0002	0.0

TABLE 1. Fraction of alpha particles lost for ten Pareto optimal configurations from figure 4. The loss fractions were computed by sampling 5000 particles on the  $s = 1/4$  flux surface or  $s = 1/2$  flux surface and evolving their trajectories by the vacuum guiding centre equations in Boozer coordinates until a terminal time of 0.1 seconds or the particle breached the  $s = 1$  surface and was lost.

equal to 8.5, is disconnected from the remainder of the efficient set. Within the gap in the Pareto front, configurations still can achieve exceptional levels of quasi-symmetry, as shown by the grey points in figure 4. The increase in quasi-symmetry in the gap may not be a physical phenomena, but may instead be due to poor numerical refinement in VMEC, local optimization or the representation of decision variables. It is interesting to see nonetheless that the set of desirable configurations is disparate.

All the Pareto optimal configurations found achieve exceptional levels of quasi-symmetry and particle confinement. From this, it is clear that the aspect ratio is not limiting the device performance. However, engineering criteria on coil shape, fabrication tolerances and placement tolerances may exhibit a greater trade-off with particle confinement, since the magnetic field generated by coils may not align well with the target magnetic field computed by ideal MHD. In the next section, we consider the trade-off between coil length and the ability for coils to reproduce a target magnetic field.

### 3.2. Problem 2: the coil length and quadratic flux trade-off

Given the shape of the last close flux surface,  $\mathcal{S}$ , the ‘stage-two’ stellarator optimization problem seeks to find magnetic coils with magnetic field  $\mathbf{B}$ , such that the field is orthogonal to the surface normal,  $\mathbf{B} \cdot \mathbf{n} = 0$  (Merkel 1987; Zhu *et al.* 2017). Coils with a tight constraint on their length will not be able to form intricate shapes and reduce the normal component of the magnetic field to zero, whereas coils which are longer can form arbitrary shapes that better minimize  $\mathbf{B} \cdot \mathbf{n}$ . Long coils, on the other hand, are undesirable from a financial and engineering standpoint: longer coils are more expensive to build, often have higher curvature and are more difficult to fit into the space around the device. In this experiment, we seek to understand the trade-off between coil length and the ability of coils to reproduce a target magnetic field.

We develop insight into this trade-off by solving the multi-objective problem

$$\min_{x \in \Omega} (J_n, L), \tag{3.5}$$

where  $J_n$  is the quadratic flux,

$$J_n(x) = \frac{1}{2} \int_{\mathcal{S}} (\mathbf{B} \cdot \mathbf{n})^2 ds, \tag{3.6}$$

$L$  is the total length of the  $n_C$  coils considered,

$$L = \sum_{i=1}^{n_C} L_i, \tag{3.7}$$

and the decision variables  $\mathbf{x}$  represent coil shape parameters and the coil currents. The quadratic flux is a standard metric used to find coils in the stage-two problem (Merkel 1987; Landreman 2017; Zhu *et al.* 2017; Singh *et al.* 2020; Kruger *et al.* 2021; Glas *et al.* 2022; Wechsung *et al.* 2022a). To solve the bi-objective coil problem, (3.5), we use the coil optimization framework implemented in SIMSOPT. The modular coils are described by a Fourier representation of the Cartesian coordinates with  $n_{\text{cm}}$  modes. The Fourier representation of the  $x$ -coordinate of the  $i$ th coil is

$$x^i(t, \mathbf{x}) = x_{c,0}^i + \sum_{k=1}^{n_{\text{cm}}} [x_{c,k}^i \cos(kt) + x_{s,k}^i \sin(kt)], \quad (3.8)$$

with analogous forms for the  $y$  and  $z$  coordinates. Each coil is additionally equipped with a current, though the current of the first coil is held fixed at  $I_1 = 10^5$  A throughout the optimization. The number of Fourier amplitudes per coil is  $3(2n_{\text{cm}} + 1)$  making the total number of design variables  $n_x = n_c 3(2n_{\text{cm}} + 1) + n_c - 1$ . The coils considered satisfy field-period symmetry, as well as stellarator symmetry. Thus, the quadratic flux is only computed over a half-field-period, the total coil length is only taken over the coils in a half-field-period and the decision variables  $\mathbf{x}$  only represent the  $n_c$  coils in a half-field-period. The plasma boundary shape  $\mathcal{S}$  used in this experiment is that of the  $n_{\text{fp}} = 2$  field-period quasi-axisymmetric (QA) configuration from Landreman & Paul (2022), which we will refer to as LP-QA. The  $n_{\text{cm}} = 5$  Fourier modes were used to describe the coil shapes.

To avoid ill-posedness of the objectives due to non-uniqueness of the underlying coil parametrization, a penalty on the variation of the arc length was used to encourage the coils to have uniform incremental arc length. The arc length variation objective discretizes each coil into  $n_{\text{av}}$  segments and computes the variance of the set of arc lengths  $\{l_i\}_{i=1}^{n_{\text{av}}}$  for each coil, i.e. for a single coil,  $V = \text{Var}[\{l_i\}_{i=1}^{n_{\text{av}}}]$ . The total arc length variation over the coil set is

$$J_{\text{av}}(\mathbf{x}) = \sum_{i=1}^{n_c} V_i. \quad (3.9)$$

The bi-objective problem was solved by applying the  $\epsilon$ -constraint method, where quadratic flux was the objective and the total coil length was the constraint. The predictor–corrector method was not used to investigate this trade-off, since the computational expense of solving the coil optimization problem with the  $\epsilon$ -constraint method was minimal. However, a warm-starting procedure or ‘zeroth-order’ predictor–corrector method, was used in which solutions of the  $\epsilon$ -constraint method were used as a good initial guess of efficient points with similar values of  $\epsilon$ . The  $\epsilon$ -constraint problems were solved with a quadratic penalty method (Nocedal & Wright 1999), see § A. The penalty parameter was increased from an initial value of 1 by a factor of 10 at each iteration, and the quadratic penalty sub-problems were solved by the L-BFGS-B algorithm (Zhu *et al.* 1997) until the norm of the gradient reached  $10^{-14}$  or 30 000 iterations elapsed. The arc length variation (3.9) was included as a penalty term with fixed penalty parameter  $10^{-4}$ .

Figure 6 shows the Pareto front for the bi-objective coil optimization problem (3.5) when  $n_c = 4$  coils are used per half-field-period. Figure 6 also shows three-dimensional renderings of the two Pareto optimal coils sets, highlighted as the star and diamond in the plot of the Pareto front. Figure 6 shows that increasing the coil length substantially improves the ability for coils to reduce the quadratic flux up until the average coil length over the minor circumference of the plasma boundary is approximately 6, after which we

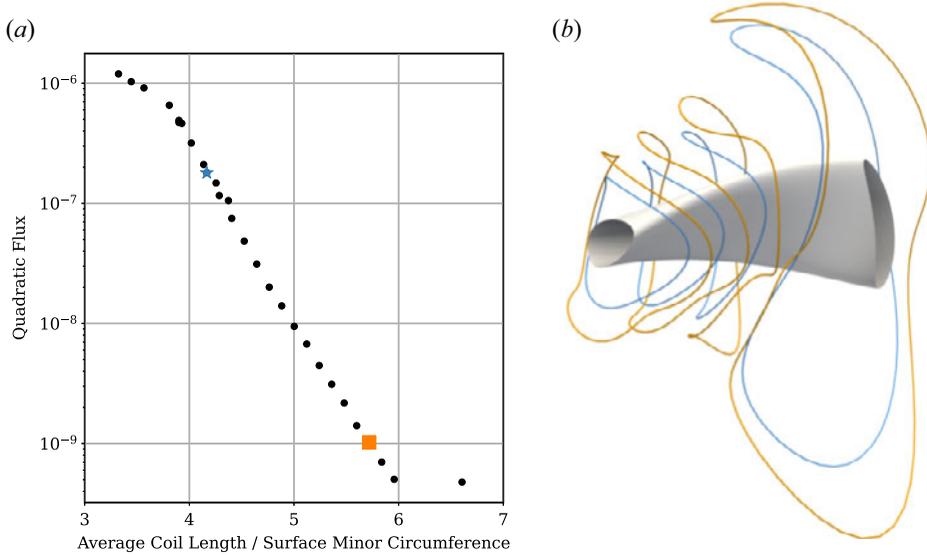


FIGURE 6. (a) Pareto front for coil length versus quadratic flux. The  $x$ -axis is non-dimensionalized by dividing by the minor circumference of the surface ( $2\pi$  times the minor radius). (b) Three-dimensional rendering of the Pareto optimal coil sets, denoted by the blue star (smaller blue coils) and the orange square (larger orange coils).

numerically find no improvement. In addition, we find that when the average coil length over the minor circumference of the plasma boundary exceeds roughly 6, the coils become exceedingly complex: coils curvature becomes large and coils begin to pass under one another, competing for space. At these coil lengths, coil curvature and the coil-to-coil separation should also be constrained.

While figure 6 shows that constraining the total coil length limits the ability for coils to reproduce a magnetic field, it does not show how the limits on the total coil length relate to a loss in quasi-symmetry or loss in confinement. In figure 7, we measure the extent to which the Pareto optimal coil configurations from figure 6 generate quasi-symmetric magnetic fields. To do so, we fit a quadratic flux minimizing (QFM) surface (Dewar, Hudson & Price 1994) to the coil-generated magnetic field  $\mathbf{B}$ , evaluate VMFC using the QFM surface as the boundary shape and compute the quasi-symmetric metric  $Q_{1,0}$ . The QFM surface is constrained to have identical volume to the LP-QA target surface. We find the QFM surface by solving the following problem in SIMSOPT:

$$\min_S \frac{\int_S (\mathbf{B} \cdot \mathbf{n})^2 ds}{\int_S \|\mathbf{B}\|^2 ds}, \tag{3.10}$$

$$\text{such that Volume}(S) = \text{Volume}(\text{LP-QA}). \tag{3.11}$$

Figure 7 shows that the quasi-axisymmetry degrades significantly, by four orders of magnitude, as the coil length constraints are tightened. Unlike the longer coil sets, the shorter coil sets generate significant coil ripple, degrading the quasi-symmetry. While it may not be possible to entirely avoid the loss of quasi-symmetry when coil constraints are tightened, the loss may be diminished by using single-stage optimization approaches to

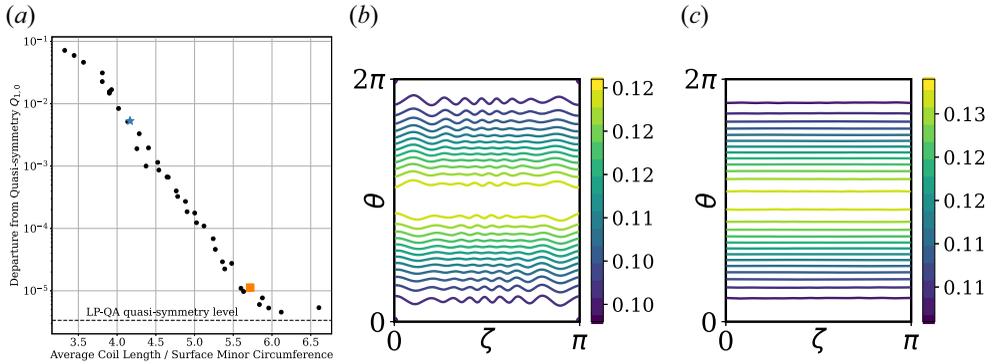


FIGURE 7. (a) Trade-off between quasi-axisymmetry and coil length for the Pareto optimal points shown in figure 6. For each Pareto optimal coil set, the quasi-axisymmetry metric  $Q_{1,0}$  was computed by running VMEC with a boundary surfaced obtained by fitting a quadratic flux minimizing surface to the field generated by the coils. (b,c) Contour plots of the coil generated field strength in the Boozer toroidal and poloidal angles  $\zeta, \theta$  on the  $s = 1$  surface of the LP-QA configuration for the coil sets of the blue star and orange square shown in figure 6.

design the coils. Part of the loss of quasi-symmetry is inherent to the two-stage approach to coil design: the coils are optimized to reproduce a target magnetic field, but not to generate a quasi-symmetric field. Single-stage optimization approaches (Giuliani *et al.* 2022a,b ; Jorge *et al.* 2023), however, directly optimize coils for quasi-symmetry, ensuring that they achieve optimal quasi-symmetry levels over the space of coils satisfying the coil length constraint. Single-stage approaches would hence improve the unfortunate trade-off between coil length and quasi-symmetry.

#### 4. Discussion

Understanding trade-offs in stellarator designs is particularly important in designing high-performance devices that satisfy the multitude of physical, engineering and financial criteria. Throughout this study, we have shown how MOO can be used to investigate trade-offs and develop insight into the role of design parameters. By quantifying the Pareto front, stellarator designers can develop an intuitive understanding of how one design criteria can affect other aspects of the design. Furthermore, the Pareto front provides a host of candidate configurations from which engineers and physicists can choose.

The example problems considered here were biobjective problems, but many stellarator design problems may have three or more competing objectives. MOO methods are useful in this context, however, the dimension of the Pareto front grows with the number of objectives, which makes visualizing the Pareto front difficult and thoroughly exploring the Pareto front often intractable. For this reason, we generally recommend exploring trade-offs between two or three objectives at a time. For three objectives, scalarization methods (Ehrgott 2005; Miettinen 2012) and other gradient-based MOO methods (Fliege & Svaiter 2000; Schäffler, Schultz & Weinzierl 2002; Désidéri 2012; Gkaragkounis *et al.* 2018; Peitz & Dellnitz 2018) are an appropriate choice since they allow the use of derivatives and for the user to dictate which part of the Pareto front is explored. For problems with more than three objectives, parallel algorithmic approaches will find promising configurations most efficiently (Deb *et al.* 2002; Knowles 2006; Wagner *et al.* 2010; Daulton *et al.* 2022; Chang & Wild 2023).

Looking beyond the two trade-offs considered here, there are a host of other trade-offs that appear in stellarator design that deserve attention. There is a natural trade-off, for instance, between coil complexity and quasi-symmetry that should be explored using a direct single-state optimization method, such as with the near-axis expansion (Giuliani *et al.* 2022a). In addition, it is not well understood how stability criteria trade off with particle confinement criteria, or how design flexibility for multipurpose coils trades off with volume and quasi-symmetry (Lee *et al.* 2022). Ideally, these problems should be solved and the Pareto optimal solutions should be tabulated in a way that allows practitioners to easily survey a multitude of configurations and analyse their strengths and weaknesses as a holistic set.

### Acknowledgements

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

### Funding

This work was generously supported by a grant from the Simons Foundation (No. 560651, D.B.).

### Declaration of interest

The authors report no conflict of interest.

### Data availability

Code and data can be found at <https://doi.org/10.5281/zenodo.7838063>.

### Appendix A. Quadratic penalty methods

The quadratic penalty method (QPM) (Nocedal & Wright 1999) solves inequality constrained optimization problems, such as (2.2), by solving a sequence of unconstrained optimization problems. The unconstrained optimization problems minimize a ‘merit function’, which captures the effects of the constraints through a smooth penalty on constraint violation.

For example, to solve (2.2), the QPM solves a sequence of unconstrained problems of the form

$$\min_x P_k(\mathbf{x}) := F_j(\mathbf{x}) + \frac{\mu_k}{2} \sum_{i \neq j}^{n_F} \max(F_i(\mathbf{x}) - \epsilon_i, 0)^2, \quad (\text{A1})$$

where  $\mu_k > 0$  is the ‘penalty parameter’. As  $\mu_k \rightarrow \infty$ , the sequence of minima  $\{\mathbf{x}_k\}$  of the penalty function can converge to minima of the  $\epsilon$ -constraint problem, (2.2). Heuristically,  $\mu_0$  can be set to 1 and increased by a factor of 10 after each optimization, until the solution of (A1) reaches a desired level of constraint violation. Solutions of one iteration of (A1) should be used to warm start the optimization of subsequent iterations.

A practical drawback of the quadratic penalty approach is that achieving highly accurate solutions of (2.2) requires increasing  $\mu_k$  to large values, and thus solving ill-conditioned optimization problems. This can be circumvented by using augmented Lagrangian methods (Nocedal & Wright 1999), in place of the QPM. While somewhat inefficient in terms of sample complexity, the QPM is easy to implement and finds approximate solutions of the constrained optimization problem quite well in practice.

## REFERENCES

- BADER, A., DREVLAK, M., ANDERSON, D.T., FABER, B.J., HEGNA, C.C., LIKIN, K.M., SCHMITT, J.C. & TALMADGE, J.N. 2019 Stellarator equilibria with reactor relevant energetic particle losses. *J. Plasma Phys.* **85** (5).
- CHANG, T.H. & WILD, S.M. 2023 Parmoo: a python library for parallel multiobjective simulation optimization. *J. Open Source Softw.* **8** (82), 4468.
- CLARKE, F.H. 1990 *Optimization and Nonsmooth Analysis*. SIAM.
- DAS, I. & DENNIS, J.E. JR. 1996 A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems. *Tech. Rep.*
- DAULTON, S., ERIKSSON, D., BALANDAT, M. & BAKSHY, E. 2022 Multi-objective bayesian optimization over high-dimensional search spaces. In *Uncertainty in Artificial Intelligence*, pp. 507–517. PMLR.
- DEB, K., PRATAP, A., AGARWAL, S. & MEYARIVAN, T.A.M.T. 2002 A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **6** (2), 182–197.
- DÉSIDÉRI, J.-A. 2012 Multiple-gradient descent algorithm (MGDA) for multiobjective optimization. *C. R. Math.* **350** (5-6), 313–318.
- DEWAR, R.L., HUDSON, S.R. & PRICE, P.F. 1994 Almost invariant manifolds for divergence-free fields. *Phys. Lett. A* **194** (1–2), 49–56.
- EHRGOTT, M. 2005 *Multicriteria Optimization*, vol. 491. Springer Science & Business Media.
- EMMERICH, M.T.M. & DEUTZ, A.H. 2018 A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Nat. Comput.* **17** (3), 585–609.
- FLIEGE, J. & SVAITER, B.F. 2000 Steepest descent methods for multicriteria optimization. *Math. Meth. Oper. Res.* **51**, 479–494.
- GARREN, D.A. & BOOZER, A.H. 1991 Existence of quasihelically symmetric stellarators. *Phys. Fluids B: Plasma Phys.* **3**, 2822–2834.
- 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., LANDREMAN, M., STADLER, G. & CERFON, A. 2022b *Direct computation of magnetic surfaces in boozer coordinates and coil optimization for quasi-symmetry*. Preprint, [arXiv:2203.03753](https://arxiv.org/abs/2203.03753).
- GKARAGKOUNIS, K., PAPOUTSIS-KIACHAGIAS, E., ASOUTI, V., GIANNAKOGLU, K. & OPTIMIZATION UNIT 2018 Adjoint-based pareto front tracing in aerodynamic shape optimization. In *10th International Conference on Computational Fluid Dynamics (ICCFD10), Barcelona, Spain*.
- GLAS, S., PADIDAR, M., KELLISON, A. & BINDEL, D. 2022 Global stochastic optimization of stellarator coil configurations. *J. Plasma Phys.* **88** (2), 905880208.
- HILLERMEIER, C., *et al.* 2001 *Nonlinear Multiobjective Optimization: A Generalized Homotopy Approach*, vol. 135. Springer Science & Business Media.
- JORGE, R., GOODMAN, A., LANDREMAN, M., RODRIGUES, J. & WECHSUNG, F. 2023 *Single-stage stellarator optimization: combining coils with fixed boundary equilibria*. Preprint, [arXiv:2302.10622](https://arxiv.org/abs/2302.10622).
- KNOWLES, J. 2006 Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Trans. Evol. Comput.* **10** (1), 50–66.
- KRANTZ, S.G. & PARKS, H.R. 2002 *The Implicit Function Theorem: History, Theory, and Applications*. Springer Science & Business Media.
- KRUGER, T.G., ZHU, C., BADER, A., ANDERSON, D.T. & SINGH, L. 2021 Constrained stellarator coil curvature optimization with focus. *J. Plasma Phys.* **87** (2), 175870201.
- LANDREMAN, M. 2017 An improved current potential method for fast computation of stellarator coil shapes. *Nucl. Fusion* **57** (4), 046003.
- LANDREMAN, M. 2022 Mapping the space of quasisymmetric stellarators using optimized near-axis expansion. *J. Plasma Phys.* **88**, 905880616.
- 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. & PAUL, E. 2022 Magnetic fields with precise quasisymmetry for plasma confinement. *Phys. Rev. Lett.* **128** (3), 035001.
- LEE, B.F., PAUL, E.J., STADLER, G. & LANDREMAN, M. 2022 Stellarator coil optimization supporting multiple magnetic configurations. *Nucl. Fusion* **63** (1), 014002.
- MERKEL, P. 1987 Solution of stellarator boundary value problems with external currents. *Nucl. Fusion* **27**, 867.
- MIETTINEN, K. 2012 *Nonlinear Multiobjective Optimization*, vol. 12. Springer Science & Business Media.
- NAJMABADI, F., RAFFRAY, A.R., ABDEL-KHALIK, S.I., BROMBERG, L., CROSATTI, L., EL-GUEBALY, L., GARABEDIAN, P.R., GROSSMAN, A.A., HENDERSON, D., IBRAHIM, A., *et al.* 2008 The ARIES-CS compact stellarator fusion power plant. *Fusion Sci. Technol.* **54** (3), 655–672.
- NOCEDAL, J. & WRIGHT, S.J. 1999 *Numerical Optimization*. Springer.
- PEITZ, S. & DELLNITZ, M. 2018 Gradient-based multiobjective optimization with uncertainties. In *NEO 2016: Results of the Numerical and Evolutionary Optimization Workshop NEO 2016 and the NEO Cities 2016 Workshop held on September 20–24, 2016 in Tlalnepantla, Mexico*, pp. 159–182. Springer.
- PEITZ, S., OBER-BLÖBAUM, S. & DELLNITZ, M. 2019 Multiobjective optimal control methods for the Navier-Stokes equations using reduced order modeling. *Acta Appl. Math.* **161** (1), 171–199.
- SCHÄFFLER, S., SCHULTZ, R. & WEINZIERL, K. 2002 Stochastic method for the solution of unconstrained vector optimization problems. *J. Optim. Theor. Appl.* **114**, 209–222.
- SCHÜTZE, O., DELL'AERE, A. & DELLNITZ, M. 2005 On continuation methods for the numerical treatment of multi-objective optimization problems. In *Dagstuhl Seminar Proceedings*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- SINGH, L., KRUGER, T.G., BADER, A., ZHU, C., HUDSON, S.R. & ANDERSON, D.T. 2020 Optimization of finite-build stellarator coils. *J. Plasma Phys.* **86** (4), 905860404.
- VASILOPOULOS, I., ASOUTI, V.G., GIANNAKOGLU, K.C. & MEYER, M. 2021 Gradient-based pareto front approximation applied to turbomachinery shape optimization. *Engng Comput.* **37** (1), 449–459.
- WAGNER, T., EMMERICH, M., DEUTZ, A. & PONWEISER, W. 2010 On expected-improvement criteria for model-based multi-objective optimization. In *Parallel Problem Solving from Nature, PPSN XI: 11th International Conference, Kraków, Poland, September 11–15, 2010, Proceedings, Part I 11*, pp. 718–727. Springer.
- WECHSUNG, F., GIULIANI, A., LANDREMAN, M., CERFON, A. & STADLER, G. 2022a Stochastic and a posteriori optimization to mitigate coil manufacturing errors in stellarator design. *Plasma Phys. Control. Fusion* **64** (10), 105021.
- WECHSUNG, F., LANDREMAN, M., GIULIANI, A., CERFON, A. & STADLER, G. 2022b Precise stellarator quasi-symmetry can be achieved with electromagnetic coils. *Proc. Natl Acad. Sci.* **119** (13), e2202084119.
- ZHU, C., BYRD, R.H., LU, P. & NOCEDAL, J. 1997 Algorithm 778: L-BFGS-B: fortran subroutines for large-scale bound-constrained optimization. *ACM Trans. Math. Softw.* **23** (4), 550–560.
- 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.