
ConStellaration: A dataset of QI-like stellarator plasma boundaries and optimization benchmarks

Santiago A. Cadena Andrea Merlo Emanuel Laude Alexander Bauer
Atul Agrawal Maria Pascu Marija Savtchouk Enrico Guiraud
Lukas Bonauer Stuart Hudson Markus Kaiser

Proxima Fusion
{scadena, amerlo}@proximafusion.com

Abstract

Stellarators are magnetic confinement devices under active development to deliver steady-state carbon-free fusion energy. Their design involves a high-dimensional, constrained optimization problem that requires expensive physics simulations and significant domain expertise. Recent advances in plasma physics and open-source tools have made stellarator optimization more accessible. However, broader community progress is currently bottlenecked by the lack of standardized optimization problems with strong baselines and datasets that enable data-driven approaches, particularly for quasi-isodynamic (QI) stellarator configurations, considered as a promising path to commercial fusion due to their inherent resilience to current-driven disruptions. Here, we release an open dataset of diverse QI-like stellarator plasma boundary shapes, paired with their ideal magnetohydrodynamic (MHD) equilibria and performance metrics. We generated this dataset by sampling a variety of QI fields and optimizing corresponding stellarator plasma boundaries. We introduce three optimization benchmarks of increasing complexity: (1) a single-objective geometric optimization problem, (2) a “simple-to-build” QI stellarator, and (3) a multi-objective ideal-MHD stable QI stellarator that investigates trade-offs between compactness and coil simplicity. For every benchmark, we provide reference code, evaluation scripts, and strong baselines based on classical optimization techniques. Finally, we show how learned models trained on our dataset can efficiently generate novel, feasible configurations without querying expensive physics oracles. By openly releasing the dataset (<https://huggingface.co/datasets/proxima-fusion/constellaration>) along with benchmark problems and baselines (<https://github.com/proximafusion/constellaration>), we aim to lower the entry barrier for optimization and machine learning researchers to engage in stellarator design and to accelerate cross-disciplinary progress toward bringing fusion energy to the grid.

1 Introduction

Fusion energy promises virtually limitless, carbon-free power by harnessing the same process that powers the sun. Magnetic confinement fusion approaches trap a fully ionized gas (plasma) within magnetic fields to sustain the conditions required for fusion. Among these, *stellarators* confine the plasma solely through external coils, which produce three-dimensional, twisted magnetic flux surfaces (Figure 1). Unlike *tokamaks*, stellarators do not rely on large internal plasma currents, thereby avoiding associated instabilities [1, 2]. However, this advantage comes with a

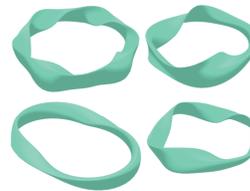


Figure 1: Examples of diverse stellarator plasma boundaries.

trade-off. Designing stellarators involves a significantly more complex parameter space: shaping the three-dimensional plasma boundary to satisfy multiple physics and engineering constraints is a high-dimensional, constrained, optimization problem.

Stellarator design has been mainly approached as a two-stage process [3]. In *stage one*, the magnetic field that confines the plasma is optimized; in *stage two*, electromagnetic coils are designed to reproduce this field. *Stage one*, the focus of this work, optimizes a three-dimensional surface that defines the boundary condition for the plasma equilibrium magnetic field. The surface is commonly parameterized by a truncated Fourier series in cylindrical coordinates (Figure 2, left). A solution to the ideal-magnetohydrodynamics (MHD) equations is then computed to determine the magnetic field inside the plasma [4]. VMEC [5] and its recent C++ re-implementation [6] are classical physics codes that compute a solution to the ideal-MHD model (Figure 2). From the MHD solution, we can compute multiple magnetic field properties, e.g. the *rotational transform*, that we can iteratively optimize to target a desired value in an outer optimization loop by updating the plasma boundary.

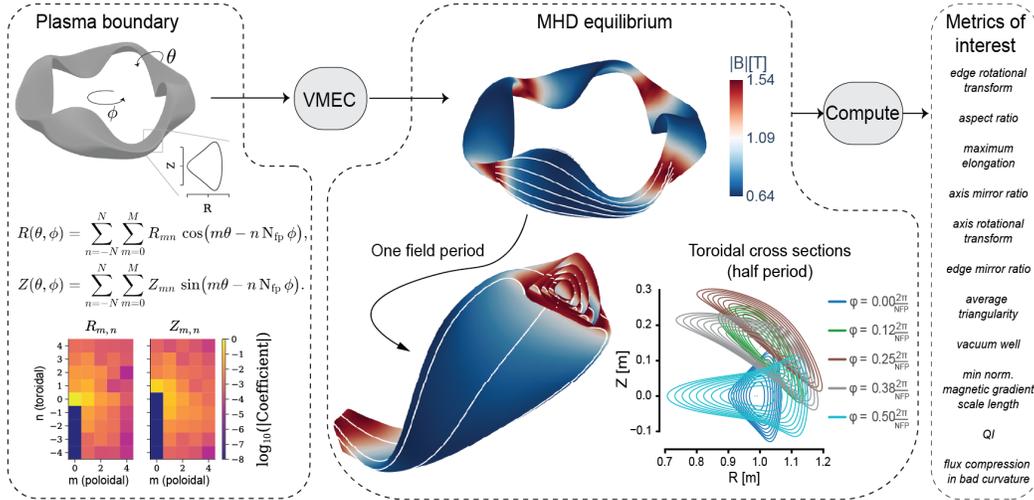


Figure 2: A plasma boundary is defined by the coefficients R_{mn} and Z_{mn} of a truncated Fourier series in cylindrical coordinates, parametrized by the lab-frame poloidal angle θ and toroidal angle ϕ . This boundary is passed to the VMEC++ code [5, 6] to compute an ideal-MHD equilibrium. In this example, the configuration is stellarator symmetric, meaning that $R(\theta, \phi) = R(-\theta, -\phi)$ and $Z(\theta, \phi) = -Z(-\theta, -\phi)$, and the number of repeated field periods (N_{fp}) is four. The ideal-MHD equilibrium defines the magnetic field throughout the plasma volume, comprising nested magnetic flux surfaces on which magnetic field lines (depicted in white) lie. We can then compute various metrics of interest from the equilibrium field.

Unlike tokamaks, classical or unoptimized stellarators lack toroidal symmetry and inherently suffer from poor confinement of high-energy particles: such fusion-born particles often escape the plasma volume, striking plasma facing components before depositing their energy back into the plasma. This prevents a self-sustained fusion process. The root cause lies in the behavior of particles trapped in poloidal, toroidal, or helical magnetic wells, which fail to sample the entire magnetic flux surface, experiencing a net radial drift that leads to gradually losing confinement. These challenges are a direct consequence of the non-axisymmetric magnetic geometry of stellarators. A particularly effective strategy to suppress these drifts is to optimize stellarators imposing the condition of *omnigenity* [7, 8], which requires only that the average radial drift of trapped particles vanishes. Among omnigenous fields, quasi-isodynamicity (QI) fields have poloidally closed contours of the magnetic field strength [9–11], which results in a vanishing net plasma toroidal current. The advantages of even approximate QI fields have been validated in laboratory experiments, most notably in the Wendelstein-7X (W7-X) stellarator [12]. These compelling benefits have made the QI symmetry a target in the design of next-generation stellarator-based fusion power plants [13, 14].

Major advances in open-source software frameworks for stellarator design have been presented in recent years. For example, SIMSOPT [15] provides high-level interfaces to link plasma equilibrium solvers such as VMEC [5] or SPEC [16] with numerical optimizers. Moreover, tools like DESC [17]

have leveraged end-to-end automatic differentiation [18] to simultaneously compute MHD equilibria and target desired properties. However, these tools still present a high entry barrier for practitioners in the optimization and machine learning communities, as they require substantial domain knowledge to make meaningful contributions.

Although significant progress has been made in defining what to target in stellarator design, there remains a lack of standardized benchmark problems and evaluation protocols to address stage one optimization. This contrasts to other areas of machine learning, where well-defined challenges have driven rapid and measurable progress [19]. Establishing such benchmarks in stellarator research would offer significant value by enabling systematic comparisons of optimization methods across a range of problem formulations. For instance, different representations (parameterizations) of the plasma boundary may vary in their effectiveness: some may better avoid local minima, while others may facilitate faster or more reliable convergence to feasible solutions. Our contributions are as follows.

- We release a diverse dataset of about 158,000 QI-like stellarator plasma boundaries with their associated ideal-MHD equilibria (in vacuum and for five different levels of plasma beta, the ratio between the plasma thermal pressure and the magnetic pressure) computed with VMEC++ [6] and corresponding figures of merit.
- We propose three optimization problems of varying complexity and kind, and release associated code.
- We provide a set of baselines for these optimization problems using classical optimization approaches.
- We show that models trained on our dataset can generate novel configurations that satisfy optimization constraints, even when only a handful of training examples do.

1.1 Related work

Stellarator datasets. Beyond works releasing a small set of plasma configurations [20, 21], large-scale datasets have focused on stellarators with quasi-axisymmetry (QA) or quasi-helical symmetry (QH) [22–24], but not QI. These studies rely on an expansion about the magnetic axis [25–30] (the field line representing the innermost flux surface) that reduces the 3D MHD equations into a 1D ordinary differential equation, which is much faster to solve. Landreman [22] sampled $\sim 500k$ QA and QH configurations, while Giuliani [23], Giuliani et al. [24] sampled $\sim 370k$ QA and QH configurations as part of the QUASR dataset¹. To the best of our knowledge, none of these datasets include publicly available computed ideal-MHD equilibria.

Stellarator optimization benchmarks. While several studies have proposed sets of optimization problems to test optimization strategies or shape parameterizations (e.g.[31]), and others have surveyed optimization approaches [32], there are no standardized benchmarks for stellarator optimization.

1.2 Background

In *Boozer* coordinates [33] (Figure 3), magnetic field strength contours of QI fields exhibit three characteristic properties: (i) the contours close poloidally, appearing as vertically closed loops in a Boozer plot; (ii) the magnetic field strength maxima align along straight vertical lines; and (iii) the arc length between points of equal magnetic field strength along a field line depends only on the flux surface (i.e., it is invariant across field lines) [11]. The targets in Figure 4 and Figure 5 are examples of precise QI fields.

2 A diverse dataset of QI-like plasma boundaries and ideal-MHD equilibria

Directly sampling the Fourier coefficients representing the plasma boundary (Figure 2) would very rarely lead to good (or even valid) stellarator fields [34]. To generate a large and diverse dataset of stellarator configurations that are approximately QI, we aim to sample diverse QI fields and other geometrical properties, and search for plasma boundaries that produce those target fields. These target

¹<https://quasr.flatironinstitute.org>

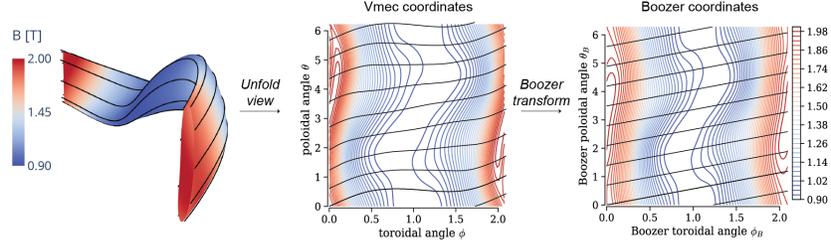


Figure 3: Visualization of the iso-contours of the magnetic field strength B and a few magnetic field lines (black). In the *Boozer* coordinate system [33], the original poloidal and toroidal angles are transformed into Boozer angles θ_B and ϕ_B , respectively, to straighten the magnetic field lines (black).

generative factors include the aspect ratio A (the ratio between the major and minor toroidal radii: R_0/a), the edge rotational transform ι_{edge} (how far a field line moves around the “short” (poloidal) way along the torus each time it goes once around the “long” (toroidal) way), the mirror ratio Δ_{edge} (defined as $(B_{max} - B_{min})/(B_{max} + B_{min})$), and the maximum elongation ϵ_{max} (the largest cross-section elongation across toroidal angles [11]).

For a given target QI field and set of properties, we generated surfaces either through physics-informed heuristics (Section 3 of Goodman et al. [11]), fast near-axis expansion models [29, 35] (using pyQSC²), or through *stage-one* optimization runs. We passed all resulting surfaces to our forward model running VMCC++ at high fidelity (Section 3) to obtain ideal-MHD equilibria and metrics of interest (Figure 2). All configurations are limited to poloidal and toroidal mode numbers of at most four. Assuming stellarator symmetry, $R_{m,n} = 0; m = 0, n < 0$ and $Z_{m,n} = 0; m = 0, n \leq 0$, and fixing the major radius $R_{0,0} = 1$, the total number of degrees of freedom is 80 (Figure 2, left).

Sampling targets. To sample diverse QI fields, we used the parameterization for an omnigenous field from Dudt et al. [8], imposing stellarator symmetry (Section A.1). Notably, our fields span a diverse range of magnetic well shapes and show variation in how these wells are stretched along field lines (Figure 4). The other target properties were drawn from a uniform distribution spanning a range of sensible values (Table 6).

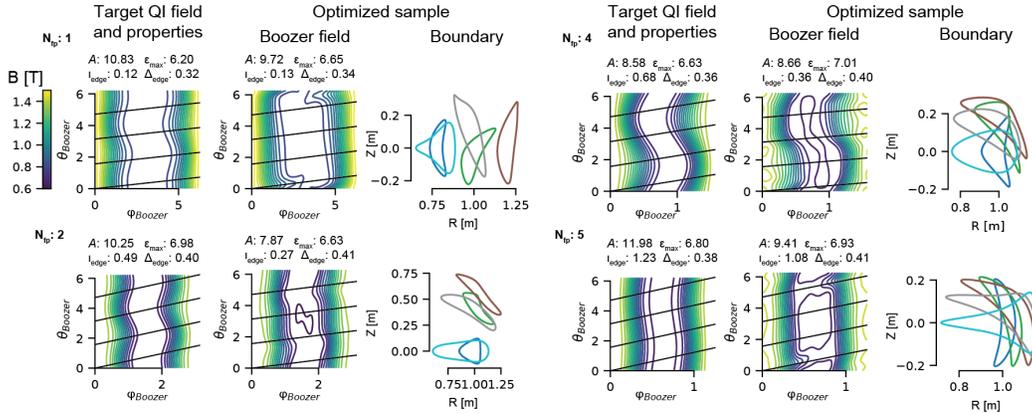


Figure 4: Four optimized samples from our dataset with 1, 2, 4, and 5 field periods. A finite computational budget for each sample generation leads to an approximate QI field at the plasma boundary. All field plots share the same color bar and the boundary cross-section labels correspond to those in Figure 2.

Optimization. We implemented *stage-one* optimization approaches seeded with heuristic or near-axis expansion models using DESC or VMCC++-based frameworks and varying objective settings (Section A.2). Multiple optimization approaches with finite budget increased diversity in the resulting

²<https://github.com/rogeriojorge/pyQIC>

boundaries, even for the same set of target field and properties (Figure 5). Each DESC run took three minutes on average on a 32 vCPU 128GB RAM machine, while each VMEC++ run took around one hour on average on a 32 vCPU 32GB RAM machine.

Results. We began by sampling 100k target sets. From this pool, we generated 30k and 49k plasma boundary candidates using our heuristic and near-axis-expansion models, respectively. We then applied the DESC optimizer twice to each target—once per initialization strategy—yielding an additional 88k optimized boundaries. A subset of 15k targets was also optimized with VMEC++ in the loop, seeded by rotating ellipses. Altogether, this produced roughly 182k candidate configurations, and we evaluated equilibria and metrics with the high-fidelity forward model on 158k of them without errors. Among these successful cases, 15k, 20k, 68k, 27k, and 28k configurations have 1, 2, 3, 4, and 5 field periods, respectively. Our resulting dataset spans a broad range of target metrics (Fig. 6, left) and reveals strong correlations between prescribed targets and the achieved values (Figure 6, right). To enable investigations of equilibrium properties at finite pressure profiles (i.e., beyond vacuum), we also made available the ideal-MHD equilibria of the boundaries at five different volume-averaged β values³ (1, 2, 3, 4, and 5%) and their correspondent metrics.

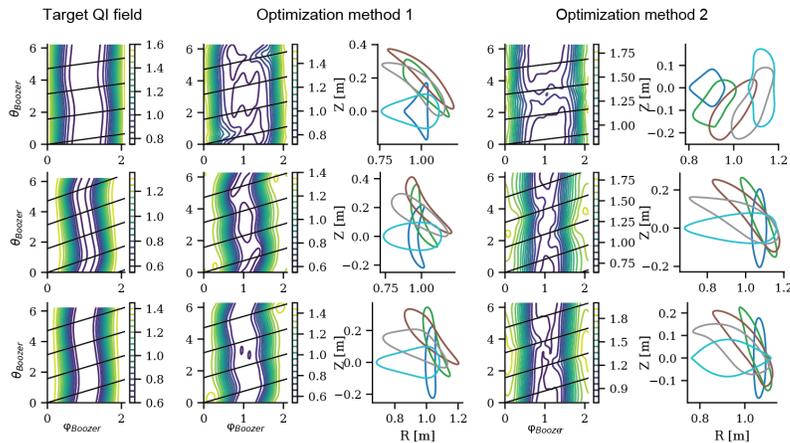


Figure 5: Diverse plasma configurations obtained for the same targets. Optimization methods vary in initialization strategy, framework, and settings. While some runs favor matching the target QI field and mirror ratio, other runs better match the remaining target properties.

Statistical consistency of the dataset. To assess the degree to which the target metrics can be inferred from the available boundary coefficients, we trained an ensemble of multilayer perceptrons (MLP) model on the optimized dataset samples and evaluated its performance on a held-out test set. The model achieved good predictive accuracy ($R^2 > 0.97$ and $RMSE < 0.1$ for all metrics; Section A.4), indicating that the metrics are fairly learnable in-domain with an expressive enough model.

3 Optimization benchmark

Stellarator design can be naturally formulated as a multi-objective constrained optimization problem [32]. The objectives and constraints arise from both engineering and economic considerations (e.g., limiting the aspect ratio to achieve a compact device) as well as physics-based requirements (e.g., ensuring a stable MHD plasma). The design process involves translating stakeholder expectations into a consistent set of feasible requirements, and navigating the trade-offs among conflicting objectives in a manner that aligns with the overarching design goals.

We introduce three prototypical stellarator design tasks with increasing complexity each involving different subsets of design metrics (Table 1): (1) *Geometric*, (2) *Simple-to-Build QI*, and (3) *MHD-stable QI*, detailed in Table 2.

³ β is defined as the ratio of the thermal plasma pressure p to the magnetic field pressure that has to be externally applied: $\beta = 2\mu_0 p / B^2$ where μ_0 is the vacuum permeability. To set these beta values, we assumed a radial linear pressure profile and scaled the pressure at the axis to match the target volume-averaged beta for each boundary.

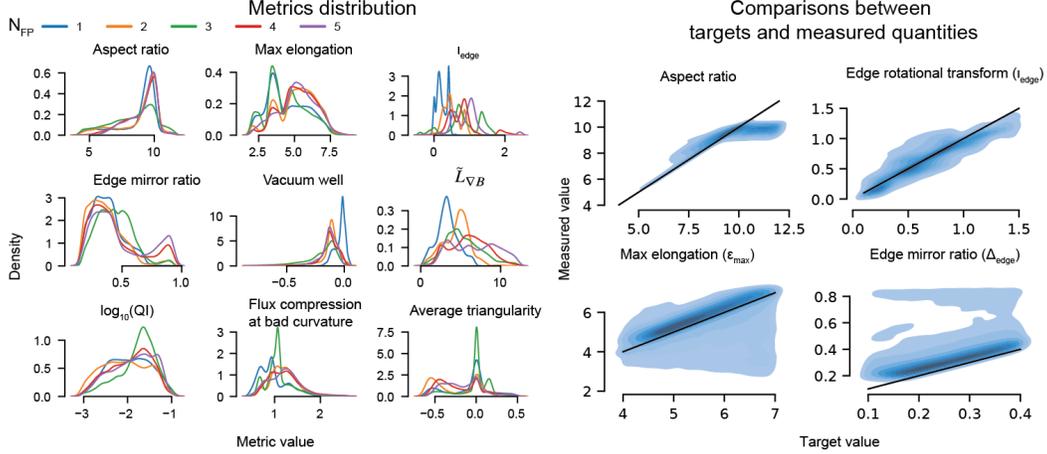


Figure 6: Distribution of metrics and comparisons between targets and outcomes. Pair plots show only optimized configurations. Black lines represent the identity. A , ϵ_{max} , and Δ_{edge} were used as upper-bound constraints during optimization, while the rotational transform was enforced as equality constraint.

Forward Model We leverage VMEC++ [6] to compute vacuum 3D ideal-MHD equilibria, scaled to $R_0 = 1$ m, $B_0 \approx 1$ T. Each vacuum equilibrium is fully defined by a single flux-surface mapping $\Sigma_\Theta : (\theta, \varphi) \mapsto (R, \phi, Z)$, where θ and φ are generic poloidal and toroidal angles, respectively, and Θ denotes the set of surface parameters, and (R, ϕ, Z) are cylindrical coordinates. In VMEC++, the truncated Fourier series presented in Figure 2 is used for Σ_Θ . However, for the purpose of these optimization benchmark problems, we make no assumptions about the functional form of Σ_Θ . All optimization problems have the form:

$$\begin{aligned} \min_{\Theta} (f_1(\Theta), f_2(\Theta), \dots) \\ \text{subject to } c_i(\Theta) \leq c_i^*, \forall i, \end{aligned} \quad (1)$$

where $f_i : \mathbb{R}^D \rightarrow \mathbb{R}$ are objective functions, $c_i : \mathbb{R}^D \rightarrow \mathbb{R}$ are constraint functions, and c_i^* are constraint targets. Each objective and constraint depends directly on the magnetic field, which in turn is determined by the surface mapping that defines the boundary condition of the ideal-MHD model.

(a) **Geometric problem**

$$\begin{aligned} \min_{\Theta} \epsilon_{max} \\ \text{s.t. } A \leq A^*, \\ \bar{\delta} \leq \bar{\delta}^*, \\ \bar{\tau} \geq \bar{\tau}^*. \end{aligned}$$

(b) **Simple-to-build QI**

$$\begin{aligned} \min_{\Theta} -\tilde{L}_{\nabla B} \\ \text{s.t. } \tilde{\tau} \geq \tilde{\tau}^*, QI \leq QI^* \\ \Delta \leq \Delta^*, A \leq A^* \\ \epsilon_{max} \leq \epsilon_{max}^* \end{aligned}$$

(c) **MHD-stable QI**

$$\begin{aligned} \min_{\Theta} (-\tilde{L}_{\nabla B}, A) \\ \text{s.t. } \tilde{\tau} \geq \tilde{\tau}^*, QI \leq QI^* \\ \Delta \leq \Delta^*, W_{MHD} \geq 0 \\ \langle \chi_{\nabla r} \rangle \leq \langle \chi_{\nabla r} \rangle^* \end{aligned}$$

Table 2: Constrained optimization problem formulations. See Table 1 for semantic associations to the symbols. All metrics are a function of the boundary

3.1 Problem 1: Geometric

To onboard contributors to stellarator optimization, we propose an intuitive, purely geometric problem (Table 2) where we look for stellarators that minimize the maximum elongation ϵ_{max} for a given aspect ratio A , edge rotational transform $\bar{\tau}$, and average triangularity $\bar{\delta}$. $\bar{\delta}$ averages the triangularity between the two

| Metric | Acronym |
|--|-----------------------------------|
| minimum normalized magnetic gradient scale length | $\tilde{L}_{\nabla B}$ |
| edge rotational transform over number of field periods | $\bar{\tau}$ |
| aspect ratio | A |
| max elongation | ϵ_{max} |
| edge magnetic mirror ratio | Δ |
| quasi isodynamicity residual | QI |
| vacuum well | W_{MHD} |
| flux compression in regions of bad curvature | $\langle \chi_{\nabla r} \rangle$ |
| average triangularity | $\bar{\delta}$ |

Table 1: Equilibrium field metrics and their acronyms.

stellarator-symmetric cross-sections ($\phi = 0$ and $\phi = \pi/N_{\text{fp}}$), and \tilde{t} is the edge rotational transform per field period.

3.2 Problem 2: Single-objective simple-to-build QI stellarator

Stellarators are notoriously challenging to construct due to their inherently three-dimensional magnetic geometry. Optimized designs like W7-X demand millimeter coil tolerances [36]. Moreover, the development and assembly of such devices can run into cost and schedule overruns driven by manufacturing complexity, potentially leading to the cancellation of entire projects as it was the case for the NCSX stellarator [37, 38]. This raises a key question: *Can optimized QI stellarators be realized using simpler, easier-to-manufacture coils?*

In a fusion reactor, the spatial region between the plasma and coils must accommodate a divertor, first wall (plasma-facing material components), neutron shielding, tritium-breeding blanket, and magnets structural support. These layers, together with the magnet superconducting technologies (e.g., low-temperature superconductors (LTS) or high-temperature superconductors (HTS)), impose geometric and engineering constraints on coil design. The feasibility of a stellarator configuration depends not only on plasma performance but also on how easily the required magnetic fields can be generated using manufacturable coils.

Not all magnetic fields are equally *coil-friendly*. We colloquially refer to *coil simplicity* as the ease with which modular coils can be placed and shaped to produce the desired field. For example, surfaces with high coil simplicity allow coils to be located further from the plasma and require lower curvature and fewer tight bends. We quantify coil simplicity using the normalized magnetic field gradient scale length on the plasma boundary, following Kappel et al. [39]. This metric has proven effective in guiding optimization towards configurations with simpler, more feasible coil designs [13, 14].

Historically, QI stellarators have required particularly complex coil geometries compared to other quasi-symmetric configurations [40–43]. This benchmark problem challenges that assumption by optimizing for *precise* QI fields that can be generated with *simple* coils.

Table 2 introduces the problem definition, where $\tilde{L}_{\nabla\mathbf{B}}$ is magnetic field gradient scale length [39] normalized by a/N_{fp} , $QI = \frac{1}{4\pi^2} \int \int r_{QI}^2 d\theta d\phi$ quantifies deviation from a *precise* QI field following Goodman et al. [11], and Δ is the magnetic mirror ratio at the plasma boundary. We normalize the objective by a/N_{fp} to ensure scale invariance across configurations with varying field period numbers. Since a QI field is easier to achieve for large aspect ratio configurations, highly elongated flux surfaces, and large mirror ratios [11], we explicitly control these quantities through inequality constraints.

3.3 Problem 3: Multi-objective ideal-MHD stable QI stellarators

This optimization problem introduces two new critical constraints for reactor relevant stellarator design: ideal-MHD plasma stability and mitigation of turbulent transport.

Despite the fact that QI configurations eliminate current-driven instabilities (“disruptions”) that often affect tokamak designs, pressure-driven instabilities persist [9], thus limiting access to high fusion power density regimes. To optimize for ideal-MHD stability, we adopt the vacuum magnetic well W_{MHD} as a proxy [25, 44]

Turbulent transport, expected to be dominated by ion-temperature gradient (ITG) turbulence in QI stellarators [12, 45, 46], limits the achievable fusion gain. Landreman et al. [47] demonstrated how purely geometrical quantities correlate strongly with the turbulence heat flux. As a constraint, we compute the “flux-surface compression in regions of *bad curvature*” given by $\chi_{\nabla r} = \mathcal{H}(\mathbf{B} \times \kappa \cdot \nabla\alpha) \|\nabla r\|_2^2$ as a simple geometric proxy. Here \mathcal{H} is the Heavyside step function, $\mathbf{B} \times \kappa \cdot \nabla\alpha$ is the curvature drift [48], κ is the magnetic field curvature, α is the field line label, and ∇r is the flux compression. A positive curvature drift represents regions of *bad curvature*. This quantity is evaluated on a single-flux surface at $\rho = r/a = 0.7$.

In quasi-poloidal (QP) and QI stellarators, $L_{\nabla\mathbf{B}} \propto R_0/N_{\text{fp}} = aA/N_{\text{fp}}^4$. More compact devices (low A) reduce capital cost per unit output power [49, 50] but increase coil complexity (proxied by $L_{\nabla\mathbf{B}}$). This

⁴Assuming that the characteristic length scale of the magnetic field gradient satisfies $L_{\nabla\mathbf{B}} \propto L_{\nabla B}$, and considering a QP magnetic field where the magnetic field strength forms a single well, i.e., $B(\varphi) = B_0 \cos(N_{\text{fp}}\varphi)$, where φ is a field-aligned coordinate.

trade-off motivates a Pareto-optimal search [51] between coil simplicity and compactness. Table 2 introduces the problem definition, where $\langle \cdot \rangle$ denotes flux-surface averaging.

3.4 Evaluation metric

We release evaluation code that scores candidate plasma boundaries across benchmarks. Our evaluation code requires the plasma boundaries to be represented by the truncated Fourier series in cylindrical coordinates (see Figure 2).

Single-objective scoring For single-objective problems, we map each design point to a bounded scalar score value $s(\Theta)$ given by:

$$s(\Theta) = \begin{cases} h(f(\Theta)) & \text{if } \tilde{c}_i(\Theta) \leq \varepsilon, \forall i, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where $f(\Theta)$ is the objective value, $h : \mathbb{R} \rightarrow [0, 1]$ is a linear map that rescales objectives into the $[0, 1]$ interval (higher is better), $\tilde{c}_i = (c_i - c_i^*)/c_i^*$ is the i -th normalized constraint violation, and ε is a relative tolerance.

Multi-objective scoring For multi-objective problems, we compute the hypervolume (HV) indicator [52, 53] over feasible solutions (i.e., those with $\tilde{c}_i(\Theta) \leq \varepsilon, \forall i$) using a fixed reference point in objectives space.

4 Optimization baselines

We now provide baselines for the three optimization problems. For all experiments, we target stellarators with three field periods and seed optimizations from rotating ellipse configurations. For the single-objective case (problem 1 and 2), we benchmark three approaches: a) gradient-based (where the gradient of the objective and constraint functions is approximated via forward finite-differences) trust-region interior point constrained optimizer [54] (`scipy-trust-constr`); b) gradient-free COBYQA [55] algorithm (`scipy-COBYQA`); and c) Augmented Lagrangian method (ALM) [56, 57] with a non-Euclidean proximal regularization [58–60] employing the NGOpt gradient-free meta-algorithm from Nevergrad [61] (ALM-NGOpt), to solve the subproblem. Implementation specifics are provided in the Section A.5.

Only ALM-NGOpt obtains feasible solutions, while both `scipy-trust-constr` and `scipy-COBYQA` did not (Table 4 and 3). Consequently, our leaderboard (Table 3) reports results exclusively for ALM-NGOpt. Figure 9 shows the optimized QI field and a representative coilset for the simple-to-build problem.

| Problem | Score \uparrow |
|-----------------|------------------|
| Geometrical | 0.969 |
| Simple-to-build | 0.431 |
| MHD-stable | 130.0 |

Table 3: ALM-NGOpt scores.

The multi-objective problem is decomposed into multiple single-objective problems by treating the aspect ratio as an inequality constraint. Using ALM-NGOpt, we found solutions for four of these instances. A sparse Pareto front is provided in Figure 7.

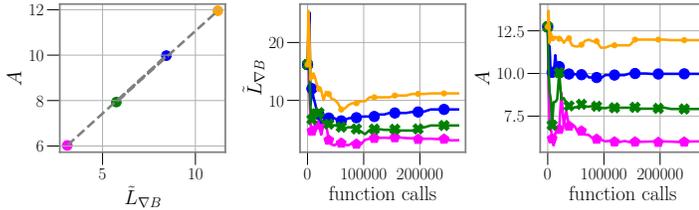


Figure 7: Pareto-front for the multi-objective optimization problem of MHD stable QI stellarators obtained with ALM-NGOpt.

5 Generative modeling of feasible domains without access to the oracle

We present a method to generate feasible configurations using learning-based models trained on the dataset, without relying on a zero-order oracle (e.g., VMEC++) and with limited feasible examples.

| Method | Simple-to-build | | Geometric problem | |
|--------------------|---------------------------------|---------------------|------------------------------|---------------------|
| | $\tilde{L}_{\nabla B} \uparrow$ | norm. constr. viol. | $\epsilon_{\max} \downarrow$ | norm. constr. viol. |
| scipy-trust-constr | 2.10* | 3.25* | 15.0* | 0.301* |
| scipy-COBYQA | 14.4* | 2.04* | 1.27 | 0.953 |
| ALM-NGOpt | 8.61 | 0.009 | 1.27 | 0.0002 |

Table 4: Comparison of baselines for the simple-to-build and the geometric problem. \uparrow means that a quantity is maximized and \downarrow means that a quantity is minimized. Final optimized boundaries for which VMCC++ failed to converge at high fidelity (i.e., the fidelity with which we score a plasma boundary) are represented with *; for them, we report the objective and constraint values from a lower fidelity equilibrium computation. `scipy-trust-constr` and `scipy-COBYQA` do not produce feasible solutions. SciPy-based optimizers ran for ~ 40 hours on a machine with 4 vCPUs. `ALM-NGOpt` ran on a 96 vCPU machine for 18 hours (geometric problem) and 34 hours (simple-to-build).

| $A \downarrow$ | $\tilde{L}_{\nabla B} \uparrow$ | norm. constr. viol. |
|----------------|---------------------------------|---------------------|
| 6.02 | 2.98 | 0.104 |
| 7.93 | 5.60 | 0.00130 |
| 9.98 | 8.45 | 0.0 |
| 11.9 | 11.1 | 0.00210 |

Table 5: Objectives and constraint violations for `ALM-NGOpt` on the multi-objective problem. Optimization was carried out by solving a sequence of single-objective problems, converting one objective into a constraint $A \leq A^*$ with $A^* \in \{6, 8, 10, 12\}$. All instances were run on a 96-vCPU machine for 15–24 h.

We test whether this method can produce many valid configurations to support downstream tasks like optimization.

We reduce the input dimensionality using Principal Component Analysis (PCA) [62] to obtain a low-dimensional latent space. In this space, Random Forest classifiers [63, 64] estimate the probability that a configuration is feasible. Thresholding this probability (e.g., above 0.8) defines a soft feasible region. Within this region, we fit a Gaussian mixture model (GMM) to capture the distribution of feasible points. Treating the GMM as a prior and the classifier output as a quasi-likelihood, we use adaptive Markov chain Monte Carlo (MCMC) [9, 65] to sample from the posterior. This allows us to generate several new configurations that are likely to satisfy constraints without querying the oracle (Figure 8). Details are given in Section A.6 with full algorithmic details in Algorithm 2.

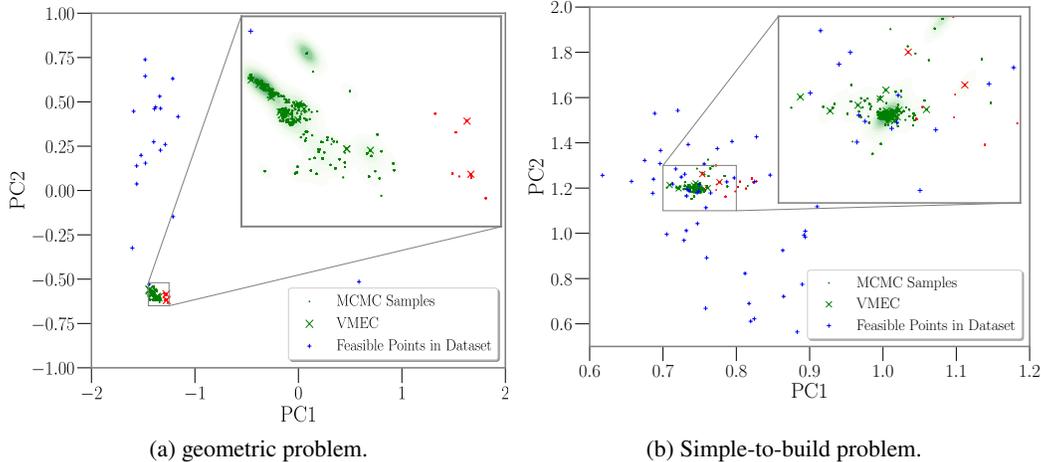


Figure 8: Posterior estimate of the feasible region in the first two PCA dimensions for two constraint-relaxed problems. Blue crosses represent feasible configurations from the dataset. Green dots show MCMC samples predicted to be feasible (classifier confidence ≥ 0.99), and red dots indicate predicted infeasible samples. Green contours reflect the estimated density of feasible samples. Oracle validation of randomly selected MCMC points are marked with green (feasible) and red (infeasible) crosses. Both the geometric and simple-to-build problems are initially relaxed, with 41 and 52 feasible points available in the dataset (out of $\sim 160k$).

When applied to relaxed versions of both the Geometric and Simple-to-build problems, our method successfully identifies regions of design space in which sample points are judged feasible by both the Random Forest classifier and the oracle model (i.e., using VMEC++) (Figure 8).

6 Discussion

We released a diverse dataset of approximately 158k QI-like stellarator plasma boundaries, associated metrics, and ideal-MHD equilibria in vacuum and at five levels of plasma beta values. Alongside the dataset, we introduced a set of stellarator optimization tasks with strong classical baselines, designed to facilitate rigorous and reproducible evaluation of stellarator optimization strategies. We further demonstrated a data-driven generative approach that can produce feasible plasma configurations without querying an expensive physics oracle. Nonetheless, several limitations remain. First, the degree of QI in the dataset is inherently limited by the finite-budget, optimization-based sampling process used during generation. Second, the dataset is limited to plasma boundaries; while these are usually the seeds in stellarator design, a consistent design also requires many additional systems (e.g., electromagnetic coils).

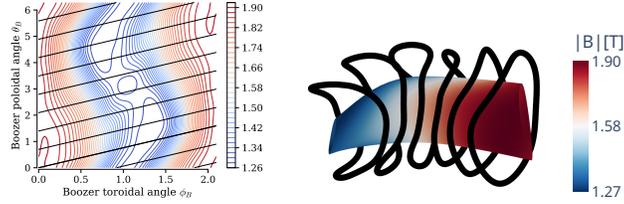


Figure 9: Left: Optimized QI magnetic field contours at the plasma boundary in Boozer coordinates for the simple-to-build optimization problem. Right: A representative coilset designed to reproduce the target magnetic field.

Acknowledgments

This work was independently funded by Proxima Fusion, and supported by the BMBF grant FUSKI (FKZ: 13F1012A).

References

- [1] John Wesson and David J Campbell. *Tokamaks*, volume 149. Oxford university press, 2011.
- [2] TC Hender, JC Wesley, J Bialek, A Bondeson, AH Boozer, RJ Buttery, A Garofalo, TP Goodman, RS Granetz, Y Gribov, et al. Mhd stability, operational limits and disruptions. *Nuclear fusion*, 47(6):S128, 2007.
- [3] S. A. Henneberg, S. R. Hudson, D. Pfefferlé, and P. Helander. Combined plasma–coil optimization algorithms. *Journal of Plasma Physics*, 87(2):905870226, 2021. doi: 10.1017/S0022377821000271.
- [4] Jeffrey P Freidberg. *ideal MHD*. Cambridge University Press, 2014.
- [5] Steven P Hirshman and JC Whitson. Steepest descent moment method for three-dimensional magnetohydrodynamic equilibria. Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 1983.
- [6] Jonathan Schilling. The numerics of vmec++. *arXiv preprint arXiv:2502.04374*, 2025.
- [7] John R Cary and Svetlana G Shasharina. Omnigenity and quasihelicity in helical plasma confinement systems. *Physics of Plasmas*, 4(9):3323–3333, 1997.
- [8] Daniel W Dudt, Alan G Goodman, Rory Conlin, Dario Panici, and Egemen Kolemen. Magnetic fields with general omnigenity. *Journal of Plasma Physics*, 90(1):905900120, 2024.
- [9] Per Helander and J Nührenberg. Bootstrap current and neoclassical transport in quasi-isodynamic stellarators. *Plasma Physics and Controlled Fusion*, 51(5):055004, 2009.
- [10] Per Helander. Theory of plasma confinement in non-axisymmetric magnetic fields. *Reports on Progress in Physics*, 77(8):087001, 2014.
- [11] Alan G Goodman, K Camacho Mata, Sophia A Henneberg, Rogerio Jorge, Matt Landreman, GG Plunk, HM Smith, RJJ Mackenbach, CD Beidler, and P Helander. Constructing precisely quasi-isodynamic magnetic fields. *Journal of Plasma Physics*, 89(5):905890504, 2023.

- [12] CD Beidler, HM Smith, A Alonso, T Andreeva, J Baldzuhn, MNA Beurskens, Matthias Borchardt, SA Bozhnikov, Kai Jakob Brunner, Hannes Damm, et al. Demonstration of reduced neoclassical energy transport in wendelstein 7-x. *Nature*, 596(7871):221–226, 2021.
- [13] J Lion, J-C Anglès, L Bonauer, A Bañón Navarro, SA Cadena Ceron, R Davies, M Drevlak, N Foppiani, J Geiger, A Goodman, et al. Stellaris: A high-field quasi-isodynamic stellarator for a prototypical fusion power plant. *Fusion Engineering and Design*, 214:114868, 2025.
- [14] CC Hegna, DT Anderson, EC Andrew, A Ayilaran, A Bader, TD Bohm, K Camacho Mata, JM Canik, L Carbajal, A Cerfon, et al. The infinity two fusion pilot plant baseline plasma physics design. *Journal of Plasma Physics*, pages 1–44, 2025.
- [15] Matt Landreman, Bharat Medasani, Florian Wechsung, Andrew Giuliani, Rogerio Jorge, and Caoxiang Zhu. Simsop: a flexible framework for stellarator optimization. *Journal of Open Source Software*, 6(65):3525, 2021.
- [16] S. R. Hudson, R. L. Dewar, G. Dennis, M. J. Hole, M. McGann, G. von Nessi, and S. Lazerson. Computation of multi-region relaxed magnetohydrodynamic equilibria. *Physics of Plasmas*, 19(11):112502, 11 2012. ISSN 1070-664X. doi: 10.1063/1.4765691. URL <https://doi.org/10.1063/1.4765691>.
- [17] DW Dudt and E Kolemen. Desc: a stellarator equilibrium solver. *Physics of Plasmas*, 27(10), 2020.
- [18] Mathieu Blondel and Vincent Roulet. The Elements of Differentiable Programming, March 2024. URL <http://arxiv.org/abs/2403.14606>. arXiv:2403.14606 [cs].
- [19] Moritz Hardt. The emerging science of machine learning benchmarks. Online at <https://mlbenchmarks.org>, 2025. Manuscript.
- [20] Richard Nies, Elizabeth J Paul, Dario Panici, Stuart R Hudson, and Amitava Bhattacharjee. Exploration of the parameter space of quasisymmetric stellarator vacuum fields through adjoint optimisation. *Journal of Plasma Physics*, 90(6):905900620, 2024.
- [21] Stefan Buller, Matt Landreman, John Kappel, and Rahul Gaur. A family of quasi-axisymmetric stellarators with varied rotational transform. *Journal of Plasma Physics*, 91(1):E18, 2025.
- [22] Matt Landreman. Mapping the space of quasisymmetric stellarators using optimized near-axis expansion. *Journal of Plasma Physics*, 88(6):905880616, 2022.
- [23] Andrew Giuliani. Direct stellarator coil design using global optimization: application to a comprehensive exploration of quasi-axisymmetric devices. *Journal of Plasma Physics*, 90(3):905900303, 2024.
- [24] Andrew Giuliani, Eduardo Rodríguez, and Marina Spivak. A comprehensive exploration of quasisymmetric stellarators and their coil sets. *arXiv preprint arXiv:2409.04826*, 2024.
- [25] Claude Mercier. Equilibrium and stability of a toroidal magnetohydrodynamic system in the neighbourhood of a magnetic axis. *Nuclear Fusion*, 4(3):213, 1964.
- [26] DA Garren and Allen H Boozer. Existence of quasihelically symmetric stellarators. *Physics of Fluids B*, 3(10):2822–2834, 1991.
- [27] David Alan Garren and AH Boozer. Magnetic field strength of toroidal plasma equilibria. *Physics of Fluids B: Plasma Physics*, 3(10):2805–2821, 1991.
- [28] Matt Landreman. Optimized quasisymmetric stellarators are consistent with the garren–boozer construction. *Plasma Physics and Controlled Fusion*, 61(7):075001, 2019.
- [29] Matt Landreman and Wrick Sengupta. Constructing stellarators with quasisymmetry to high order. *Journal of Plasma Physics*, 85(6):815850601, 2019.
- [30] Eduardo Rodríguez, Wrick Sengupta, and Amitava Bhattacharjee. Constructing the space of quasisymmetric stellarators through near-axis expansion. *Plasma Physics and Controlled Fusion*, 65(9):095004, 2023.
- [31] Sophia A Henneberg, Per Helander, and Michael Drevlak. Representing the boundary of stellarator plasmas. *Journal of Plasma Physics*, 87(5):905870503, 2021.
- [32] Rory Conlin, Patrick Kim, Daniel W Dudt, Dario Panici, and Egemen Kolemen. Stellarator optimization with constraints. *Journal of Plasma Physics*, 90(5):905900501, 2024.

- [33] Allen H Boozer. Plasma equilibrium with rational magnetic surfaces. Technical report, Princeton Plasma Physics Lab.(PPPL), Princeton, NJ (United States), 1981.
- [34] P Curvo, DR Ferreira, and R Jorge. Using deep learning to design high aspect ratio fusion devices. *Journal of Plasma Physics*, 91(1):E38, 2025.
- [35] R Jorge, W Sengupta, and M Landreman. Near-axis expansion of stellarator equilibrium at arbitrary order in the distance to the axis. *Journal of Plasma Physics*, 86(1):905860106, 2020.
- [36] Thomas Rummel, Konrad Risse, Holger Viebke, Torsten Braeuer, and Johann Kisslinger. Accuracy of the construction of the superconducting coils for wendelstein 7-x. *IEEE transactions on applied superconductivity*, 14(2):1394–1398, 2004.
- [37] National compact stellarator experiment (ncsx) project closeout report. Technical report, Princeton Plasma Physics Laboratory, August 2009. URL https://ncsx.ppp1.gov/NCSX_Engineering/CloseOut_Documentation/Closeout_Report_20080910.pdf. Prepared for the U.S. Department of Energy Office of Science.
- [38] GH Neilson, CO Gruber, Jeffrey H Harris, DJ Rej, RT Simmons, and RL Strykowski. Lessons learned in risk management on ncsx. *IEEE transactions on plasma science*, 38(3):320–327, 2010.
- [39] John Kappel, Matt Landreman, and Dhairya Malhotra. The magnetic gradient scale length explains why certain plasmas require close external magnetic coils. *Plasma Physics and Controlled Fusion*, 66(2):025018, 2024.
- [40] Haifeng Liu, Akihiro Shimizu, Mitsutaka Isobe, Shoichi Okamura, Shin Nishimura, Chihiro Suzuki, Yuhong Xu, Xin Zhang, Bing Liu, Jie Huang, et al. Magnetic configuration and modular coil design for the chinese first quasi-axisymmetric stellarator. *Plasma and Fusion Research*, 13: 3405067–3405067, 2018.
- [41] Rogerio Jorge, GG Plunk, M Drevlak, M Landreman, J-F Lobsien, K Camacho Mata, and P Helander. A single-field-period quasi-isodynamic stellarator. *Journal of Plasma Physics*, 88(5):175880504, 2022.
- [42] Florian Wechsung, Matt Landreman, Andrew Giuliani, Antoine Cerfon, and Georg Stadler. Precise stellarator quasi-symmetry can be achieved with electromagnetic coils. *Proceedings of the National Academy of Sciences*, 119(13):e2202084119, 2022.
- [43] Alexander Vyacheslav Wiedman, Stefan Buller, and Matt Landreman. Coil optimization for quasi-helically symmetric stellarator configurations. *arXiv preprint arXiv:2311.16386*, 2023.
- [44] John M Greene. A brief review of magnetic wells. *Comments on Plasma Physics and Controlled Fusion*, 17:389–402, 1997.
- [45] Marc NA Beurskens, Sergey A Bozhenkov, O Ford, Pavlos Xanthopoulos, Alessandro Zocco, Yuri Turkin, A Alonso, Craig Beidler, I Calvo, Daniel Carralero, et al. Ion temperature clamping in wendelstein 7-x electron cyclotron heated plasmas. *Nuclear Fusion*, 61(11):116072, 2021.
- [46] Alan G Goodman, Pavlos Xanthopoulos, Gabriel G Plunk, Håkan Smith, Carolin Nührenberg, Craig D Beidler, Sophia A Henneberg, Gareth Roberg-Clark, Michael Drevlak, and Per Helander. Quasi-isodynamic stellarators with low turbulence as fusion reactor candidates. *PRX Energy*, 3(2):023010, 2024.
- [47] Matt Landreman, Jong Youl Choi, Caio Alves, Prasanna Balaprakash, R Michael Churchill, Rory Conlin, and Gareth Roberg-Clark. How does ion temperature gradient turbulence depend on magnetic geometry? insights from data and machine learning. *arXiv preprint arXiv:2502.11657*, 2025.
- [48] GT Roberg-Clark, P Xanthopoulos, and GG Plunk. Reduction of electrostatic turbulence in a quasi-helically symmetric stellarator via critical gradient optimization. *Journal of Plasma Physics*, 90(3):175900301, 2024.
- [49] WR Spears and JA Wesson. Scaling of tokamak reactor costs. *Nuclear Fusion*, 20(12):1525, 1980.
- [50] JP Freidberg, FJ Mangiarotti, and J Minervini. Designing a tokamak fusion reactor—how does plasma physics fit in? *Physics of Plasmas*, 22(7), 2015.
- [51] Joaquim RRA Martins and Andrew Ning. *Engineering design optimization*. Cambridge University Press, 2021.

- [52] Eckart Zitzler and Lothar Thiele. Multiobjective optimization using evolutionary algorithms—a comparative case study. In *International conference on parallel problem solving from nature*, pages 292–301. Springer, 1998.
- [53] Miqing Li and Xin Yao. Quality evaluation of solution sets in multiobjective optimisation: A survey. *ACM Computing Surveys (CSUR)*, 52(2):1–38, 2019.
- [54] Richard H Byrd, Mary E Hribar, and Jorge Nocedal. An interior point algorithm for large-scale nonlinear programming. *SIAM Journal on Optimization*, 9(4):877–900, 1999.
- [55] Tom M Ragonneau. Model-based derivative-free optimization methods and software. *arXiv preprint arXiv:2210.12018*, 2022.
- [56] Magnus R Hestenes. Multiplier and gradient methods. *Journal of optimization theory and applications*, 4(5):303–320, 1969.
- [57] Michael JD Powell. A method for nonlinear constraints in minimization problems. *Optimization*, pages 283–298, 1969.
- [58] R Tyrrell Rockafellar. Augmented lagrangians and applications of the proximal point algorithm in convex programming. *Mathematics of operations research*, 1(2):97–116, 1976.
- [59] Emanuel Laude and Panagiotis Patrinos. Anisotropic proximal point algorithm. *arXiv preprint arXiv:2312.09834*, 2023.
- [60] Paul Dufossé and Nikolaus Hansen. Augmented lagrangian, penalty techniques and surrogate modeling for constrained optimization with cma-es. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 519–527, 2021.
- [61] J. Rapin and O. Teytaud. Nevergrad - A gradient-free optimization platform. <https://GitHub.com/FacebookResearch/Nevergrad>, 2018.
- [62] Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459, 2010.
- [63] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- [64] Kevin P Murphy. *Probabilistic machine learning: Advanced topics*. MIT press, 2023.
- [65] Russel E Caflisch. Monte Carlo and quasi-Monte Carlo methods. *Acta numerica*, 7:1–49, 1998.
- [66] Sergey Shirobokov, Vladislav Belavin, Michael Kagan, Andrei Ustyuzhanin, and Atilim Gunes Baydin. Black-box optimization with local generative surrogates. *Advances in Neural Information Processing Systems*, 33:14650–14662, 2020.
- [67] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(3):88, 2022.
- [68] Nathan Baker, Frank Alexander, Timo Bremer, Aric Hagberg, Yannis Kevrekidis, Habib Najm, Manish Parashar, Abani Patra, James Sethian, Stefan Wild, et al. Workshop report on basic research needs for scientific machine learning: Core technologies for artificial intelligence. Technical report, USDOE Office of Science (SC), Washington, DC (United States), 2019.
- [69] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [70] W Keith Hastings. Monte Carlo sampling methods using markov chains and their applications. 1970.
- [71] Heikki Haario, Eero Saksman, and Johanna Tamminen. An adaptive metropolis algorithm. *Bernoulli*, pages 223–242, 2001.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We made claims about generating a diverse dataset of plasma configurations (see Section 2). We also made claims about three optimization benchmark problems (see Section 4 and code), and strong baselines (see Section 4). Finally we made claims about a data-driven machine learning model trained on our data for finding novel configurations (see Section 5).

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: See limitations in the Discussion (Section 6).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We did not include any theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Our code (<https://github.com/proximafusion/constellation>) and the explanations found in (Sections 2, 4, 5) and in the Appendix, provide sufficient detail to reproduce the main claims.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code is open and available at <https://github.com/proximafusion/constellation> and the dataset is also open source at <https://huggingface.co/datasets/proxima-fusion/constellation>

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We went in great detail to explain what optimizers we used and referenced relevant packages. See Section 4, Section A.5, Section 5, and Section A.6).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: We did not make experiments requiring error bars as we focused on 1) generating plasma boundaries, and 2) generating unique configurations that serve as baselines meeting all the constraints of our optimization problems.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The computational cost of generating the dataset was outline on a per-optimization method approach in Section 2. For the optimization baseline, we provided the compute resources (number of vCPUs and time of execution) in the captions of Tables 4 and 5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We reviewed the NeurIPS Code of Ethics and believe that we adhere to it.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In the Abstract and Introduction we reflect on the potential of fusion energy and the impact that stellarator optimization can have for a future of clean energy.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We don't believe our code and dataset of plasma boundaries and ideal MHD equilibria carries any risk.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Both the code and the dataset we release use the MIT license as specified in their corresponding repositories. The list of dependencies also have very permissive licenses.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The code and the dataset have thorough README files. Moreover, the paper describes the data generation process in detail (see Section 2).

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No crowdsourcing involved.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Did not involve any crowdsourcing.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA] .

Justification: Only used LLMs for text editing.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Technical Appendices

A.1 Data generation: sampling omnigenous poloidal fields

We leverage the parameterization of an *omnigenous poloidal* field from Dudt et al. [8] in which the 1D magnetic well on each flux surface is represented by a spline on the interval $[-\pi/2, \pi/2]$. The well is symmetric about its minimum, so it can simply be parameterized between B_{min} and B_{max} . The full omnigenous field is then built by “morphing” this one-dimensional well across magnetic field lines via a computational coordinate h , which is expanded in a Chebyshev basis (radial index l) and Fourier bases (poloidal index m , toroidal index n) with coefficients x_{lmn} [8]. In practice, we generate new omnigenous-poloidal fields by sampling both the spline knots and x_{lmn} coefficients. To enforce stellarator symmetry (invariance under simultaneous flips of the poloidal and toroidal Boozer angles), we set the coefficients of the odd terms of the Fourier basis along the toroidal direction to zero, namely $x_{lmn} = 0 \forall n \geq 0$.

To produce a variety of monotonically increasing 1D well shapes, we draw knot positions from $\text{Beta}(\alpha, \beta)$ cumulative distribution functions and then rescaled them to lie between B_{min} and B_{max} . Finally, we fix the mean magnetic field at 1 T and sample the mirror ratio Δ to determine the pair (B_{min}, B_{max}) .

The ranges from which we sample these parameters can be found in Table 6.

| Parameter | Min | Max |
|------------------|-----|------|
| N_{fp} | 1 | 5 |
| $\tilde{\iota}$ | 0.1 | 0.3 |
| A | 4.0 | 12.0 |
| ϵ_{max} | 4.0 | 7.0 |
| α_{Beta} | 2.0 | 6.0 |
| β_{Beta} | 2.0 | 6.0 |
| Δ_{edge} | 0.1 | 0.4 |

Table 6: Ranges of sampling parameters with both minimum and maximum values.

A.2 Data generation: stage one optimizations using DESC [17]

Given a set of target quantities:

$$T = (\iota^*, A^*, E^*, O^*)$$

where

- ι^* is the desired edge rotational transform,
- A^* is the target aspect ratio,
- E^* is the maximum elongation,
- O^* is the target omnigenous field,

we ran numerical optimizations to find a toroidal boundary surface Σ (parameterized in a Fourier-*RZ* basis) that simultaneously matches these goals. Note that the mirror ratio Δ is defined within O^* .

A.2.1 Initial Guess Generation

An initial boundary Σ_0 is generated either by

1. **Heuristic QP model (Section III from Goodman et al. [11]):** prescribing average major radius R_0 , aspect ratio A^* , elongation E^* , mirror ratio, torsion, and field periods; or
2. **Near-Axis Expansion (NAE) using pyQSC**⁵: specifying A^* , E^* , ι^* , mirror ratio, field periods, and mode cutoffs.

This yields a smooth Σ_0 expressed in the `FourierRZToroidalSurface` format of DESC.

⁵<https://github.com/rogeriojorge/pyQIC>

A.2.2 Equilibrium Solve

Starting from Σ_0 , we form the DESC equilibrium object and solve the force balance

$$\mathcal{E}(\Sigma) = \text{Equilibrium}(\Psi, \Sigma, M, N)$$

and solve the magnetostatic force-balance equations using

$$\mathcal{E} \xrightarrow{\text{solve(force)}} \mathcal{E}^{\text{sol}}.$$

A.2.3 Objective Function

On the solved equilibrium \mathcal{E}^{sol} , we define individual objective terms:

$$J_A(\Sigma) = \frac{R_0(\Sigma)}{a(\Sigma)}, \quad f_A = w_A (J_A - A^*)^2, \quad (3)$$

$$J_E(\Sigma) = \max_{\varphi} \frac{b(\varphi; \Sigma)}{a(\varphi; \Sigma)}, \quad f_E = w_E (J_E - E^*)^2, \quad (4)$$

$$J_l(\Sigma) = \iota [\mathcal{E}^{\text{sol}}], \quad f_l = w_l (J_l - \iota^*)^2, \quad (5)$$

$$\mathbf{J}_O(\Sigma) = \mathcal{O}[\mathcal{E}^{\text{sol}}, \mathcal{O}^*], \quad f_O = w_O \|\mathbf{J}_O\|_2^2, \quad (6)$$

where a, b are the minor/major half-axes of the cross-section, φ is the toroidal angle, and the omnigenous residual \mathbf{J}_O is computed by the DESC `Omnigenity` objective using the target field \mathcal{O}^* .

The omnigenity contribution $\|\mathbf{J}_O\|_2^2$ is given by

$$\|\mathbf{J}_O\|_2^2 = \sum_{i=1}^{N_\eta} \sum_{j=1}^{N_\alpha} w(\eta_i) [B_{\text{eq}}(\rho_0, \eta_i, \alpha_j) - B^*(\rho_0, \eta_i, \alpha_j)]^2,$$

with the poloidal weight

$$w(\eta) = \frac{\eta_{\text{weight}} + 1}{2} + \frac{\eta_{\text{weight}} - 1}{2} \cos \eta \quad (\text{so } w \equiv 1 \text{ if } \eta_{\text{weight}} = 1),$$

where B_{eq} is the field strength of the equilibrium in (ρ, η, α) coordinates and B^* is the perfectly-omnigenous target field generated by `OmnigenousField` as in [8]. The residuals $r_{ij} = \sqrt{w(\eta_i)}(B_{\text{eq}} - B^*)_{ij}$ are evaluated on the same (η, α) -grid used by the target field.

On a solved equilibrium \mathcal{E}^{sol} at a fixed flux surface $\rho = \rho_0$, we assemble a least-squares objective

$$\mathcal{L}(\Sigma) = f_A + f_E + f_l + f_O.$$

Internally, DESC invokes `JAX` to compute residuals, leveraging automatic differentiation to compute gradients.

The objective is then wrapped in an augmented-Lagrangian least-squares optimizer (`lsq-auglag`) [32] to minimize $\|r\|_2^2$ alongside the other terms.

A.2.4 Constraints

To enforce vacuum equilibrium and fix global invariants, the following constraints are imposed:

$$\begin{aligned} R_{0,0}(\Sigma) &= 1, & (\text{FixBoundaryR}) \\ j_{\parallel}(\Sigma) &= 0, & (\text{CurrentDensity}) \\ p(\Sigma) &= 0, & (\text{FixPressure}) \\ J_{\text{tor}}(\Sigma) &= 0, & (\text{FixCurrent}) \\ \Psi(\Sigma) &= \text{const.}, & (\text{FixPsi}) \end{aligned}$$

where each is implemented via the corresponding DESC linear-objective wrapper.

A.2.5 Nonlinear Optimization

We employ DESC’s `lsq-auglag` optimizer [32] to solve

$$\min_{\Sigma} \mathcal{L}(\Sigma) \quad \text{s.t. all linear constraints,}$$

using automatic differentiation and a trust-region least-squares augmented-Lagrangian scheme. Iterations continue until convergence (up to 200 iterations by default), yielding the optimized boundary Σ^* .

Our exact implementation is available at <https://github.com/proximafusion/constellation>.

A.3 Stage one optimizations using VMEC++ [6] in the loop

We carried out optimizations using the `NGOpt` algorithm from the `Nevergrad`⁶ library. To improve convergence, we preconditioned the problem using a diagonal scaling matrix as detailed in Section A.2.1. We parameterized the boundary with up to four poloidal and toroidal Fourier modes and ran the optimization on a single machine equipped with 32 vCPUs and 32GB of RAM. Each run is allocated a time budget of approximately 1 h.

The optimization minimizes the following objective function:

$$\begin{aligned} f(\Theta) = & \int_0^{2\pi} \int_0^{\pi/N_{\text{fp}}} (B(\theta, \phi) - B^*(\theta, \phi))^2 d\theta d\phi \\ & + \int_0^{2\pi} \left(\max_{\phi} B(\theta, \phi) - B(\theta, \phi = 0) \right)^2 d\theta \\ & + \left(\frac{A - A^*}{A^*} \right)^2 \\ & + \left(\frac{\iota_{\text{edge}} - \iota_{\text{edge}}^*}{\iota_{\text{edge}}^*} \right)^2 \\ & + \left(\max \left(0, \frac{\epsilon_{\text{max}} - \epsilon_{\text{max}}^*}{\epsilon_{\text{max}}^*} \right) \right)^2. \end{aligned} \quad (7)$$

where B denotes the magnetic field strength from the ideal-MHD equilibrium in Boozer coordinates, and B^* represents the target omnigenous magnetic field strength. The quantities A , ι_{edge} , and ϵ_{max} correspond to the aspect ratio, edge rotational transform, and maximum elongation, respectively, with asterisks denoting their target values. The additional target on the maxima of the magnetic field strength guides the optimizer towards more QI fields.

In the optimization loop, we used `VMEC++` within the forward model. To speed up the generation of the optimized boundary, we run `VMEC++` at a lower resolution than the one used to score plasma boundaries in optimization benchmarks (e.g. reduced number of flux surfaces, higher required force tolerance to converge).

Due to the constrained time budget, the optimization may not fully minimize the objective function but added the desired diversity to the dataset.

Our exact implementation is available at <https://github.com/proximafusion/constellation>.

A.4 In-domain predictability of the dataset.

We filtered the vacuum data for three field period configurations that were also the result of either the DESC or VMEC optimizations. Then we filtered outliers (0.05% tails) for each of the metrics, resulting in $\sim 23k$ data points. Finally, we split the data into training and test (20%) sets.

⁶<https://github.com/facebookresearch/nevergrad>

To facilitate training, we also Z-scored the output metrics, keeping track of the training set statistics for later inference.

Using Bayesian optimization to sweep over hyperparameters like network depth, width, type of activation, and learning rate; we converged to an ensemble model of ten multi-layer perceptrons (MLPs) with three layers, 256 hidden units, and \tanh activations. The MLPs mapped Fourier-boundary coefficients to target key metrics by minimizing mean squared error.

We obtained fairly good in-domain generalization results in terms of root mean squared error (RMSE), Pearson’s correlation coefficient (R^2), normalized root mean squared error (NRMSE), and signal-to-noise ratio (SNR) (Table 7).

Table 7: Test-set performances of an MLP ensemble model trained to predict target metrics from boundary coefficients

| Metric | RMSE | R^2 | NRMSE | SNR |
|---|-------|-------|-------|---------|
| aspect_ratio | 0.090 | 0.997 | 0.050 | 808.54 |
| aspect_ratio_over_edge_rotational_transform | 0.581 | 0.993 | 0.080 | 507.73 |
| max_elongation | 0.161 | 0.994 | 0.072 | 731.78 |
| axis_rotational_transform_over_n_field_periods | 0.006 | 0.994 | 0.071 | 556.74 |
| edge_rotational_transform_over_n_field_periods | 0.006 | 0.997 | 0.052 | 985.74 |
| axis_magnetic_mirror_ratio | 0.010 | 0.974 | 0.159 | 79.72 |
| edge_magnetic_mirror_ratio | 0.013 | 0.989 | 0.102 | 244.02 |
| average_triangularity | 0.018 | 0.995 | 0.065 | 806.58 |
| vacuum_well | 0.006 | 0.998 | 0.062 | 1934.06 |
| minimum_normalized_magnetic_gradient_scale_length | 0.330 | 0.990 | 0.101 | 342.57 |
| flux_compression_in_regions_of_bad_curvature | 0.034 | 0.990 | 0.059 | 432.36 |
| log_10_qi | 0.051 | 0.982 | 0.134 | 132.22 |

We highlight that such surrogate models are prone to extrapolation errors, particularly when queried far from the data distribution they were trained on [66]. Uncertainty calibration, active learning, and physics informed strategies (among others) could be considered moving forward for effective surrogate-based optimizations [67, 68].

A.5 Optimization baselines

A.5.1 Implementation details and hyperparameters

In this section we provide implementation details for the optimization baseline. For the SciPy-based optimizers, we use default parameters, and set the maximum number of iterations to a large value.

We implement a variant of the proximal ALM [58] where the quadratic proximal term is replaced by a trust-region constraint. This can be seen as an instance of the anisotropic proximal ALM [59]. The modification is essential for improving convergence when using evolutionary algorithms (such as NGOpt), as it restricts the sampling of new candidate solutions to a region around the current iterate [60].

As the degrees of freedom Θ operate on different scale, we precondition the problem with a diagonal matrix $\text{diag}(\Lambda)$ where the entries Λ decay exponentially. We define the rescaled variables as $\tilde{\Theta} := \text{diag}(\Lambda)^{-1}\Theta$ and $\tilde{f}(\tilde{\Theta}) := f(\text{diag}(\Lambda)\tilde{\Theta})$ and $\tilde{c}_i(\tilde{\Theta}) := c_i(\text{diag}(\Lambda)\tilde{\Theta})$. In addition, we apply a base-10 logarithmic transformation to the QI constraint.

In each iteration, the algorithm alternates between primal and dual updates. For each constraint \tilde{c}_i , it tracks a penalty parameter ρ_i^k and a Lagrange multiplier y_i^k . The complete algorithm is given in Algorithm 1.

Algorithm 1 non-Euclidean proximal augmented Lagrangian method

Require: $\Theta^0 \in \mathbb{R}^D$, $\rho^0 \in \mathbb{R}_{++}^m$, $y^0 \in \mathbb{R}_+^m$, $\delta_0 > 0$, $0 < \tau, \gamma < 1$, $\sigma > 1$ and $\delta_{\min}, \rho_{\max} > 0$

1: **for** $k \in \{0, 1, \dots, N\}$ **do**

2: Primal update

$$\tilde{\Theta}^{k+1} = \underset{\tilde{\Theta} \in B(\tilde{\Theta}^k, \delta_k)}{\arg \min} \tilde{f}(\tilde{\Theta}) + \frac{1}{2} \sum_{i=1}^m \frac{1}{\rho_i^k} \left(\max\{0, y_i^k + \rho_i^k \tilde{c}_i(\tilde{\Theta}^k)\}^2 - (y_i^k)^2 \right) \quad (8)$$

3: dual update

$$y_i^{k+1} = \max\{0, y_i^k + \rho_i^k \tilde{c}_i(\tilde{\Theta}^{k+1})\}$$

4: update penalty parameters

$$\rho_i^{k+1} = \begin{cases} \rho_i^k & \text{if } \tilde{c}_i(\tilde{\Theta}^{k+1}) \leq \tau \tilde{c}_i(\tilde{\Theta}^k) \\ \min\{\rho_{\max}, \sigma \rho_i^k\} & \text{otherwise.} \end{cases}$$

5: decrease trust-region

$$\delta_{k+1} = \max\{\delta_{\min}, \gamma \delta_k\}$$

6: **end for**

For the geometric problem we choose $\rho_i^0 = 10$, $\rho_{\max} = 1e9$, $\delta_0 = 0.5$, $\gamma = 0.9$, $\delta_{\min} = 0.05$, $\tau = 0.8$, $\sigma = 5$. The subproblem (8) is solved with NGOpt with a budget of $\min\{20.000, 1500 + k \cdot 260\}$ forward-model calls.

For the simple-to-build problem we choose $\rho_i^0 = 10$, $\rho_{\max} = 1e9$, $\delta_0 = 0.5$, $\gamma = 0.95$, $\delta_{\min} = 0.05$, $\tau = 0.8$, $\sigma = 5$. The subproblem (8) is solved with NGOpt with a budget of $\min\{20.000, 1500 + k \cdot 260\}$ forward-model calls.

For the MHD-stable problems we choose $\rho_i^0 = 10$, $\rho_{\max} = 1e8$, $\delta_0 = 0.33$, $\gamma = 0.95$, $\delta_{\min} = 0.05$, $\tau = 0.8$, $\sigma = 5$. The subproblem (8) is solved with NGOpt with a budget of $\min\{20.000, 1500 + k \cdot 300\}$ forward-model calls.

For all problems, Θ^0 is a rotating ellipse configuration. We optimize up to four poloidal and toroidal Fourier modes, which results in $D = 80$ degrees of freedom. During the optimization, we run VMEC++ at low fidelity.

A.5.2 Additional experimental results

We provide convergence plots for the three problems obtained with ALM-NGOpt. Green curves represent metrics that are constrained. Red colored metrics are maximized and blue colored metrics are minimized. Gray curves correspond to metrics that are not part of the optimization problem. The blue dashed lines indicate lower bounds and the red dashed lines indicate upper bounds. Figure 11 provides plots corresponding to the single-objective problem, while Figure 12 provides a plot for one instance ($A \leq 8$) of the sequence of single-objective problems corresponding to the multi-objective problem.

In Figure 10 we show the initial and final plasma configurations for the different problems.

A.6 Generative modeling details

We use the Random Forest classifier and the GMM implementations from Scikit-learn [69]. We use the Random walk Metropolis-Hastings algorithm [70] with adaptive proposal distribution [71] as the MCMC sampler. To monitor the convergence of the MCMC sampler, Figure 13 presents the log-probability of the posterior distribution evaluated at each sampled point. The rising and stabilizing log-probability indicates convergence to high-density regions. Algorithm 2 summarizes the formulation discussed in Section 5.

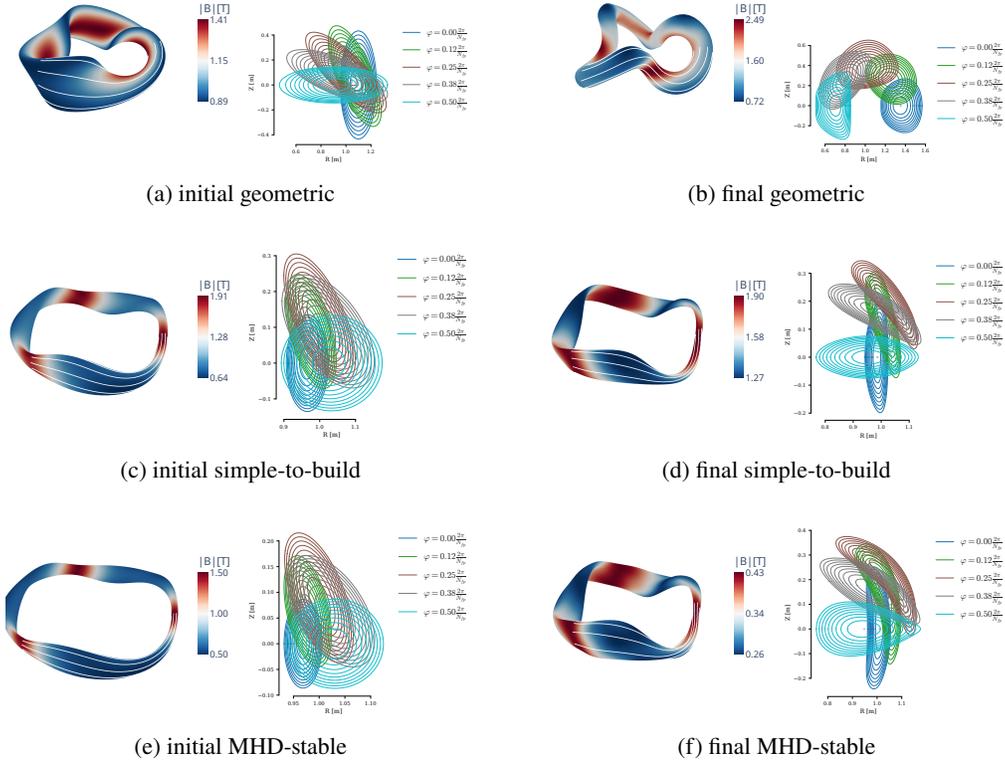
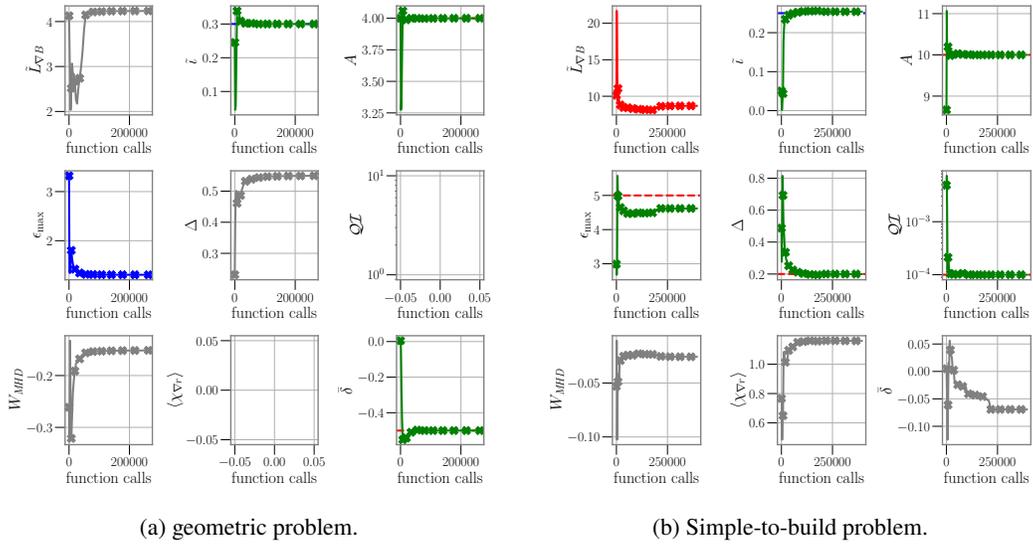


Figure 10: Initial guesses and final plasma configurations optimized with ALM-NGOpt. We selected a low aspect ratio configuration from the Pareto Front of solutions for the multi-objective problem.



(a) geometric problem.

(b) Simple-to-build problem.

Figure 11: Single-objective problem optimization traces.

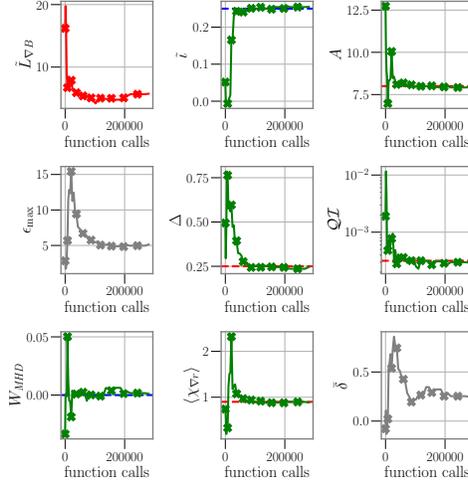
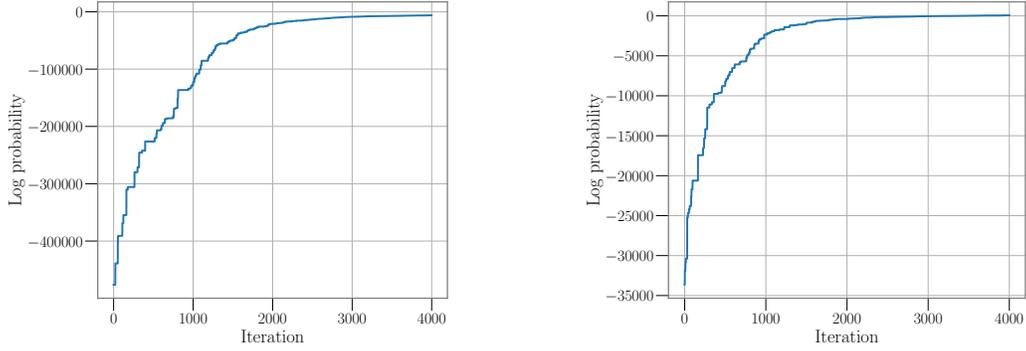


Figure 12: Multi-objective problem with $A \leq 8$.



(a) geometric problem.

(b) Simple-to-build QI problem.

Figure 13: Trace plot of the log-posterior probability values over MCMC iterations.

Algorithm 2 Generative Inference of Feasible Configurations without Oracle Access

Require: Dataset $\mathcal{D} = \{x_1, \dots, x_N\} \subset \mathbb{R}^D$, constraint definition

Ensure: Set of configurations $\{x^*\}$ predicted to lie in the feasible domain

// Dimensionality Reduction:

- 1: Compute PCA mapping $\Phi : \mathbb{R}^D \rightarrow \mathbb{R}^d$, where $d \ll D$
- 2: Project dataset to latent space: $Z \leftarrow \{\mathbf{z}_i = \Phi(x_i)\}_{i=1}^N$

// Feasibility Classification:

- 3: Train Random Forest classifiers $\{C_i(\mathbf{z})\}_{i=1}^{N_c}$ to predict feasibility label $y \in \{0, 1\}$
- 4: Define soft-feasible region: $\tilde{\mathcal{F}} \leftarrow \bigcap_{i=1}^{N_c} \{p(C_i(\mathbf{z}) = 1) \geq \tau_i\}$, where $\tau_i = 0.8 \forall i$

// Density Estimation:

- 5: Fit Gaussian Mixture Model $\text{GMM}(\mathbf{z})$ on data restricted to $\tilde{\mathcal{F}}$

// Bayesian Refinement:

- 6: Define prior: $p(\mathbf{z}) \leftarrow \text{GMM}(\mathbf{z})$
- 7: Define quasi-likelihood: $\ell(\mathbf{z}) \leftarrow \sum_{i=1}^{N_c} \log C_i(\mathbf{z})$
- 8: Compute posterior using MCMC: $\{\mathbf{z}^*\} \sim p(\mathbf{z} \mid \text{feasible}) \propto \ell(\mathbf{z}) \cdot p(\mathbf{z})$
- 9: Inverse transform to original space: $x^* \leftarrow \Phi^{-1}(\mathbf{z}^*)$

// Oracle Validation:

- 10: Evaluate x^* using VMEC++ oracle to confirm feasibility

11: **return** $\{x^*\}$
