
Diffusion for Fusion: Designing Stellarators with Generative AI

Misha Padidar , Teresa Huang , Andrew Giuliani , Marina Spivak 

Center for Computational Mathematics

Flatiron Institute

New York, NY 10010

(mpadidar, thuang, agiuliani, mspivak)@flatironinstitute.org

Abstract

Stellarators are a prospective class of fusion-based power plants that confine a hot plasma with three-dimensional magnetic fields. Typically framed as a PDE-constrained optimization problem, stellarator design is a time-consuming process that can take hours to solve on a computing cluster. Developing fast methods for designing stellarators is crucial for advancing fusion research. Given the recent development of large datasets of optimized stellarators, machine learning approaches have emerged as a potential candidate. Motivated by this, we present an *open inverse* problem to the machine learning community: to rapidly generate high-quality stellarator designs which have a set of desirable characteristics. As a case study in the problem space, we train a conditional diffusion model on data from the QUASR database to generate quasisymmetric stellarator designs with desirable characteristics (aspect ratio and mean rotational transform). The diffusion model is applied to design stellarators with characteristics not seen during training. We provide evaluation protocols and show that many of the generated stellarators exhibit solid performance: less than 5% deviation from quasisymmetry and the target characteristics. The modest deviation from quasisymmetry highlights an opportunity to reach the sub 1% target. Beyond the case study, we share multiple promising avenues for generative modeling to advance stellarator design.

1 Introduction

Stellarators are a prospective class of power plants that generate renewable energy with nuclear fusion [33]. The toroidal devices confine a hot plasma using strong magnetic fields generated by electromagnetic coils. Unlike their axisymmetric cousin, the tokamak, stellarators have magnetic fields with full three-dimensional shaping. Due to the complexity of the magnetic fields, stellarator design is typically framed as a PDE-constrained shape optimization problem, which requires running computationally intensive numerical solvers and evaluators [27, 7, 9, 51]. To advance research in fusion energy, it is crucial to accelerate stellarator design. Leveraging recent advances in simulation capabilities, the stellarator community has opened the doors to *data-driven* stellarator design by compiling large datasets of optimized stellarators [21, 20, 8, 41, 14, 19]. Using these datasets, machine learning models can uncover hidden relationships [42, 58], and rapidly generate stellarators with promising performance [14]. Given the past impact of generative machine learning methods on scientific applications, such as molecular design [26, 35], material design [46, 10], and astrophysics [50, 49, 36, 68] generative models are well positioned to have an impact on stellarator design.

The aim of this paper is to state an *open inverse problem* in stellarator design as a data-driven, generative modeling task, and to provide an entry point for the generative modeling community to advance stellarator research. The problem of interest is as follows,

Inverse Problem (IP). Given an array of desirable stellarator characteristics $\mathbf{y} \in \mathbb{R}^{n_y}$, generate a stellarator design which has properties \mathbf{y} .

Rapidly solving (IP) is the central goal of stellarator design. The main challenge being that the inverse mapping from properties, \mathbf{y} , to stellarator designs is not directly available. In fact, while the function, $\mathbf{F} : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_y}$, mapping design variables, \mathbf{x} , to stellarator properties, \mathbf{y} , is well defined, it may not be invertible since many distinct designs have the same properties. As a result, numerical optimization is typically used to vary \mathbf{x} until $\mathbf{F}(\mathbf{x})$ is close to a desired \mathbf{y} . Generative models, however, open the doors to a new probabilistic approach to this problem. Generative models can “invert” \mathbf{F} in a probabilistic sense, by finding a density with support over the set of \mathbf{x} that map to \mathbf{y} . In probabilistic terms, the generative model would learn a sampling mechanism (or conditional probability density) to produce features \mathbf{x} from conditions \mathbf{y} , $\mathbf{x} \sim p(\mathbf{x}|\mathbf{y})$, such that $\mathbf{F}(\mathbf{x}) \approx \mathbf{y}$.

To facilitate working with stellarator data, we provide an introduction to stellarator design, which discusses design variables \mathbf{x} , properties of stellarators \mathbf{y} , methods of evaluation \mathbf{F} , and current approaches to designing stellarators (see Sec. 2). In addition, we provide a concrete case study, in which we train denoising diffusion models [31] to generate stellarator designs with $n_x = 661$ features given $n_y = 4$ conditions, using data drawn from the QUASR database (see Sec. 3) [21, 20]. The diffusion model is applied to design stellarators with characteristics not seen during training. Many of the generated stellarators exhibit solid performance: a low deviation from quasisymmetry and good adherence to the conditions. To open the doors to further research, we discuss several promising research directions, from physics informed loss design, to dataset curation, model architecture design, sampling procedures, and multi-fidelity models (see Sec. 4).

2 Stellarator Design: Preliminaries, Problem Set-up, and Related Work

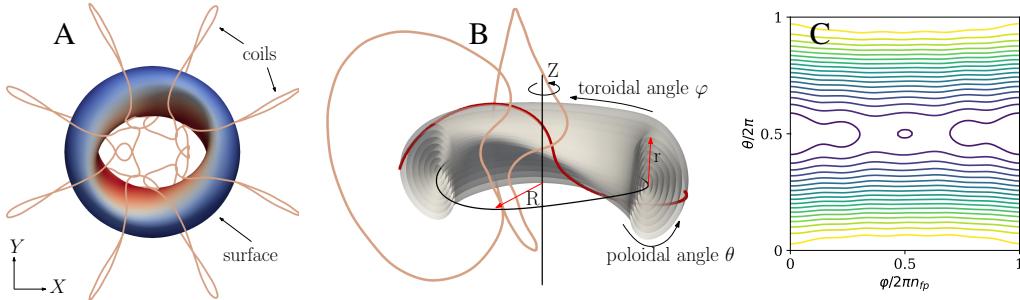


Figure 1: (A) Stellarator-0010198 from the QUASR database, encircled by the modular electromagnetic coils (gold) that generate its magnetic field (more information about this device can be found at [\[link\]](#)). The surface is colored by the strength of the magnetic field. (B) A slice of the device in (A), with internal magnetic surfaces plotted. The device’s minor radius, r , and major radius, R , are illustrated along with the directions in which the toroidal and poloidal angles increase on the magnetic surfaces. The device presents $n_{fp} = 2$ discrete rotational symmetry about the Z axis. The red curve indicates a hypothetical field line. (C) isolines of field strength $\|\mathbf{B}\|$ on a magnetic surface, where φ, θ are the toroidal and poloidal Boozer angles on the surface.

Stellarators use non-axisymmetric magnetic fields, \mathbf{B} , to confine plasma, a hot soup of ionized gas, to torus-like volumes [33]. These fields can be generated using electromagnetic coils (Figure 1A) and often form a region of nested magnetic surfaces (Figure 1B), which are everywhere tangent the magnetic field. A stellarator’s magnetic field satisfies the ideal MHD equations

$$\begin{aligned} \mu_0^{-1}(\nabla \times \mathbf{B}) \times \mathbf{B} &= \nabla p, \\ \nabla \cdot \mathbf{B} &= 0, \\ \mathbf{B} \cdot \mathbf{n} &= 0 \quad \text{on } S(\mathbf{x}), \end{aligned} \tag{1}$$

where S is the shape of the plasma boundary (a magnetic surface), p is the pressure profile, μ_0 is the permeability of free space, and \mathbf{n} is the surface normal. The geometry of S is given by the design variables \mathbf{x} , which are typically Fourier harmonics (see Appendix A for more information). In this work we only consider *vacuum* magnetic fields, those with zero current, $\nabla \times \mathbf{B} = 0$ – a common

simplification. Given the shape of the plasma boundary S , the ideal MHD equations can be solved to obtain \mathbf{B} with the VMEC, DESC, GVEC or SPEC codes [30, 18, 29, 32]; VMEC has long been the workhorse of the stellarator design community, but, unlike DESC, does not give access to derivatives with respect to design parameters.

Stellarator design is focused, in part, on manipulating properties of magnetic fields with nested surfaces. Points on a magnetic surface can be described with two angles: the poloidal angle, θ , and toroidal angle, φ , see Figure 1B; the toroidal angle varies the long way around the torus while the poloidal angle varies the short way around the torus. A magnetic surface has three basic attributes: the number of field periods n_{fp} , the aspect ratio A , and the rotational transform, ι . The number of field periods, n_{fp} , refers to a discrete rotational symmetry, i.e., rotating the device by $2\pi/n_{fp}$ about the Z -axis results in the same device. A device with n_{fp} -field periods can be sliced (like a pizza) into n_{fp} identical pieces; for example, Fig. 3 (top row) shows devices with $n_{fp} = 2, 3, 4, 5, 6, 7, 8$, from left to right. The aspect ratio, $A = R/r$, is a geometric quantity measuring the ratio of the major to minor radius of the device, see Figure 1B, and reflects how thin the enclosed plasma is. Low aspect ratio devices are often preferred, since they correspond to a high volume of plasma confined, relative to the footprint. A third property which characterizes magnetic surfaces is the rotational transform ι . The rotational transform is the average number of poloidal rotations a field line makes for every toroidal turn around the torus. Figure 1B illustrates a hypothetical red field line making, roughly, one poloidal turn as it transits one field period, leading to $\iota \approx n_{fp} = 2$. When given multiple nested magnetic surfaces with different values of ι , the average value, $\bar{\iota}$, is considered. When designing stellarators, the shape of S is varied to achieve desirable values of $\bar{\iota}, A$.

Without special shaping of the magnetic field, the plasma will not stay confined to the magnetic surfaces – plasma particles will drift off surfaces until they hit the chamber walls. To improve fusion performance, and prevent damage to plasma facing components, the magnetic field is optimized to confine the plasma particles. In addition to particle confinement, goals of stellarator design include turbulence and instability suppression, controlled heat load on the plasma facing components, and extraction of “ash” particles. Stellarator optimization problems can be stated as PDE-constrained optimization problem of the form,

$$\begin{aligned} & \min_{\mathbf{x}, \mathbf{B}} J(\mathbf{B}, \mathbf{x}), \\ & \text{subject to: (1),} \\ & \mathbf{c}(\mathbf{B}, \mathbf{x}) = 0, \end{aligned} \tag{2}$$

where the objective, J , and constraints, \mathbf{c} , encode desirable stellarator properties. Practically, (2) can be solved by eliminating \mathbf{B} via the PDE constraint (1), and minimizing the reduced objective $\hat{J}(\mathbf{x}) = J(\mathbf{B}(\mathbf{x}), \mathbf{x})$ subject to the reduced constraints $\hat{\mathbf{c}}(\mathbf{x}) = \mathbf{c}(\mathbf{B}(\mathbf{x}), \mathbf{x}) = 0$.

Due to the complexity of stellarator design, it is common to split the problem into two stages: in stage-I the shape of the plasma boundary S is optimized so that magnetic field has nested magnetic surfaces, good particle confinement, suppresses plasma instabilities, etc.; in stage-II the electromagnetic coils are designed to produce the magnetic field found in stage-I [70, 65]. The stage-II problem is a mechanical engineering problem, determined by constraints, \mathbf{c} , on the shape of the coils (length, curvature, torsion) – material properties and manufacturing practices place tolerances and limits on the shapes and currents which coils can sustain (see [70] for an introduction to stage-II). In certain settings, the stage-I and II are solved simultaneously as a *single-stage* – an end-to-end optimization [23, 24, 22, 20].

In this work, we focus on solutions to the stage-I problem. The decision variables, \mathbf{x} , parameterize the shape of the plasma boundary, S . Two common constraints, which we consider in this work, constrain the aspect ratio and $\bar{\iota}$ of the device to target values $A^*, \bar{\iota}^*$: $c_A(\mathbf{x}) = \frac{A(\mathbf{x}) - A^*}{A^*}$ and $c_{\bar{\iota}}(\mathbf{x}) = \frac{\bar{\iota}(\mathbf{x}) - \bar{\iota}^*}{\bar{\iota}^*}$ [44, 2, 3, 34, 41]. The objective function used in this work, $J = J_{QS}$ [24], measures the deviation of the magnetic field from the nearest quasisymmetric magnetic field, and is a proxy measure for particle confinement (see Appendix A). Quasisymmetry is a symmetry in the strength of the magnetic field, $\|\mathbf{B}\|$, that ensures that particle trajectories have no radial (outward) drift, thus ensuring good particle confinement [28, 44]. Quasisymmetric magnetic fields can be identified by plotting the contours of the field strength in a special coordinate system, Boozer coordinates – the field strength contours of quasisymmetric magnetic fields only depend on a linear combination of the Boozer angles, $\|\mathbf{B}\|(\theta, \varphi) = \|\mathbf{B}\|(\theta - Nn_{fp}\varphi)$, thus appearing straight in Boozer coordinates. For example, in Fig. 1(C) the fieldlines are approximately parallel to the x -axis (aside from the small curve near the

center). The type of quasisymmetry is parameterized by the normalized *helicity*, N : $N = 0$ denotes quasaxisymmetry (QA) in which $\|\mathbf{B}\|$ appears flat when plotted in Boozer coordinates e.g. Fig. 1(C); $N = 1$ denotes quasi-helical symmetry (QH) in which the contours of $\|\mathbf{B}\|$ form diagonal lines when plotted in Boozer coordinates i.e. Fig. 3(second row, column 3). Both types of quasisymmetry are desirable, however, there is an interplay between n_{fp} and N that makes it difficult to find high n_{fp} , QA devices, and low n_{fp} , QH devices. The metrics c_ι , c_A , J_{QS} can all be computed from the solution of (1), see Appendix A for details.

Challenges of the stage-I problem (2) include: high dimensionality ($n_x = \mathcal{O}(100-1,000)$); moderate computational cost (5 – 10 seconds to solve (1) or compute a gradient when using MPI-parallelism or GPU acceleration); multiple local minima. Due to the high dimensionality of x , gradient-based optimization algorithms are crucial to efficiently find solutions. Since (2) typically has multiple minima, globalization techniques can be useful for adequately exploring the landscape of solutions. Generating solutions comes at a high computational cost, and limits the extent of exploration.

Generative models are poised to overcome the computational challenges associated with sampling and optimization, since they can aggregate information across a variety of data sources and generate valuable solutions, rapidly. Generative models can also produce samples at new sets of conditions and provide a means of exploration of the stellarator landscape. To solve (2), a generative model would learn a sampling mechanism to produce x given conditions y . In this work, we characterize a (quasisymmetric) stellarator with the four conditions $y = (\iota, A, n_{fp}, N)$. A high quality generative model would consistently produce designs that minimize $J_{QS}(x)$ and satisfy $c_\iota(x) = 0, c_A(x) = 0$. Therefore, we use these metrics to evaluate the stellarators produced by a generative model.

2.1 Related Work

Fusion/stellarator optimization To our knowledge, the first work designing stellarators with generative models was [14], which used a mixture density network [6] to generate stellarators with quasisymmetry in their core. Yet challenges remain to generate stellarators with quasisymmetry throughout their volume, and to improve adherence to conditions. [8] trained a classifier to determine whether stellarator configurations satisfy a set of constraints and used Markov chain Monte Carlo methods to sample from this region. Another work, [15] took a regression approach, learning a nonlinear regressor from conditions to features. Appendix C covers additional work using ML and generative models in fusion research.

Generative models To our knowledge, this is the first use of diffusion models [31, 52] in stellarator design. Diffusion models have been applied successfully for solving inverse problems [11, 12, 16], and widely used for generative modeling in astrophysics, medical imaging [50, 56], image generation [57], audio synthesis [38], and protein design [64]. Generative adversarial networks and variational autoencoders are also a common choice for solving inverse problems [1, 69, 25], and generating solutions to problems in material science and molecular design [10, 26].

3 Case Study: Designing Stellarators With Diffusion Models

In this section, we share the methods and results of a case study in which diffusion models are trained to solve (IP). Specifically, our aim is to train a diffusion model to generate boundary shapes, S , of *quasisymmetric* stellarators. Generating boundary shapes for quasisymmetric stellarators would solve the stage-I stellarator optimization problem. Our code and data are available on Zenodo [54].

Training data We use data from the **QUAsisymmetric Stellarator Repository (QUASR)**. With over 300,000 stellarators, the QUASR repository is a large database of optimized stellarators including both coils and magnetic surfaces. QUASR was generated using Bayesian global optimization over the space of vacuum (current-free) stellarators, in search of quasisymmetric stellarators with a range of n_{fp} , aspect ratio A , and rotational transform ι [21, 20].

While the devices in QUASR often have multiple associated magnetic surfaces, we construct our dataset by just taking the outermost surface. The feature vector, x consists of $n_x = 661$ Fourier coefficients describing the Cartesian components of the surface S , $[x(\varphi, \theta), y(\varphi, \theta), z(\varphi, \theta)]$, through three separate two-dimensional Fourier series in the poloidal and toroidal Boozer angles (see Appendix A for more details). Fourier features have a natural scaling to them: features corresponding to high mode number Fourier coefficients have less impact on the surface geometry than those with low mode

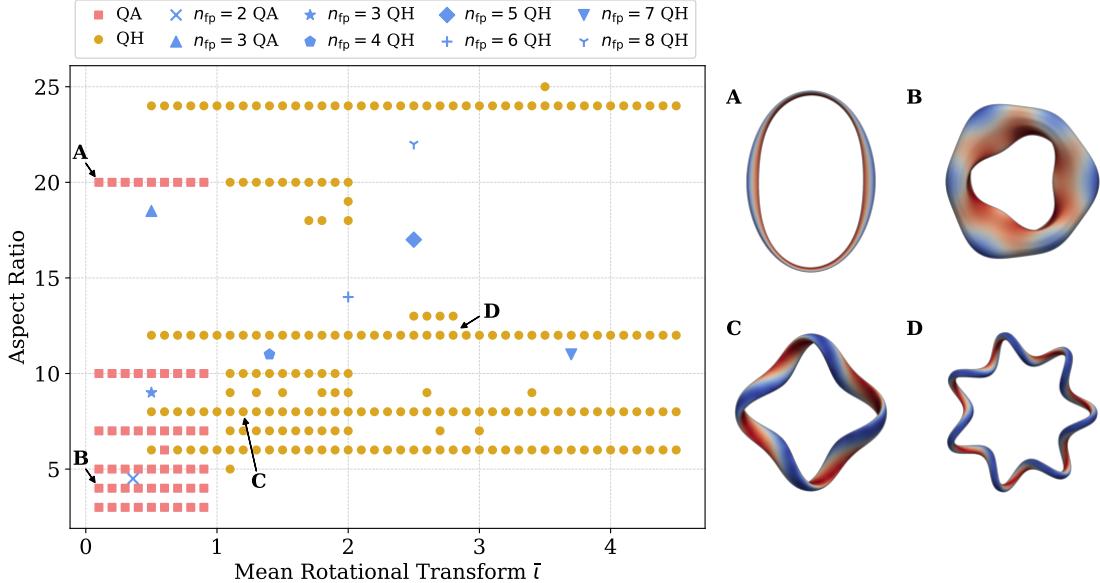


Figure 2: (Left) Values of the mean rotational transform and aspect ratio for devices in QUASR, colored by the type of quasisymmetry: QA (red square), QH (gold circle). The various blue markers indicate arbitrary conditions that are not within QUASR; in Figures 3 and 5, the diffusion model is used to generate stellarators with these conditions. (Right) Visualization of four devices in QUASR: (A) an $n_{fp} = 2$ -QA device; (B) an $n_{fp} = 3$ -QA device; (C) an $n_{fp} = 4$ -QH device; (D) an $n_{fp} = 7$ -QH device. The mean rotational transform and aspect ratio of each configuration can be seen on the left figure (labels A, B, C, D).

number. Since the feature representation of each surface is not unique, the surfaces are “canonicalized” to a unique form, using the procedure in [21]. Since the initial submission of this manuscript, it was found that less than 1% of the data points were affected by an error in the canonicalization procedure and roughly 8% of the data appeared twice in the dataset. Retraining the model on a corrected dataset showed that the main results were unaffected by these errors. For transparency, both datasets have been made available.

The diffusion model is given $n_y = 4$ conditions, $\mathbf{y} = (\bar{\iota}, A, N, n_{fp})$. J_{QS} is not treated as a condition since the stage-I problem seeks devices with $J_{QS} = 0$ and specifying a value of J_{QS} as a condition is unintuitive for domain specialists. The distribution of $\bar{\iota}$ and aspect ratio of the devices in QUASR is shown in Figure 2 as gold circles and red squares; the gold circles denote QH devices ($N = 1$), while the red square denotes QA devices ($N = 0$). Notably, the devices sparsely sample the space of conditions, providing plenty of opportunity for the diffusion model to generate novel designs. The 8 blue markers in Figure 2 highlight arbitrary conditions that are *not* in QUASR, which we use to evaluate the performance of the diffusion model in Figure 5.

Evaluation criteria A well-trained diffusion model should produce samples \mathbf{x} given conditions \mathbf{y} , such that $c_{\bar{\iota}}(\mathbf{x}) = 0, c_A(\mathbf{x}) = 0$. In addition, since the stellarators in the training data have low J_{QS} , we expect a well trained diffusion model to generate stellarators with similarly low values of J_{QS} . We quantify the performance of the diffusion model with three metrics: $c_{\bar{\iota}}, c_A, J_{QS}$. Since mean-rotational transform and aspect ratio are conditions, $c_{\bar{\iota}}, c_A$ measure the relative error in a $\bar{\iota}(\mathbf{x}), A(\mathbf{x})$ from the condition values. We consider $c_{\bar{\iota}}, c_A < 5\%$ to be satisfactory values of the metric. J_{QS} on the other hand, measures the quasisymmetry error of \mathbf{x} rather than error from a condition. Quasisymmetry must be achieved to a relatively high tolerance to achieve good particle confinement [66], so we consider $J_{QS} < 1\%$ to be a satisfactory value. See Appendix A for details on our evaluation procedure.

Diffusion model The *conditional* diffusion model is an instance of a Denoising Diffusion Probabilistic Model (DDPM) [31], modified to learn a conditional density $p(\mathbf{x}|\mathbf{y})$. Following [31], we modify the

neural network f_θ to accept the additional conditions \mathbf{y} , i.e. $f_\theta(\mathbf{x}_t, t, \mathbf{y})$. A detailed discussion of the training and sampling algorithms can be found in Appendix B.

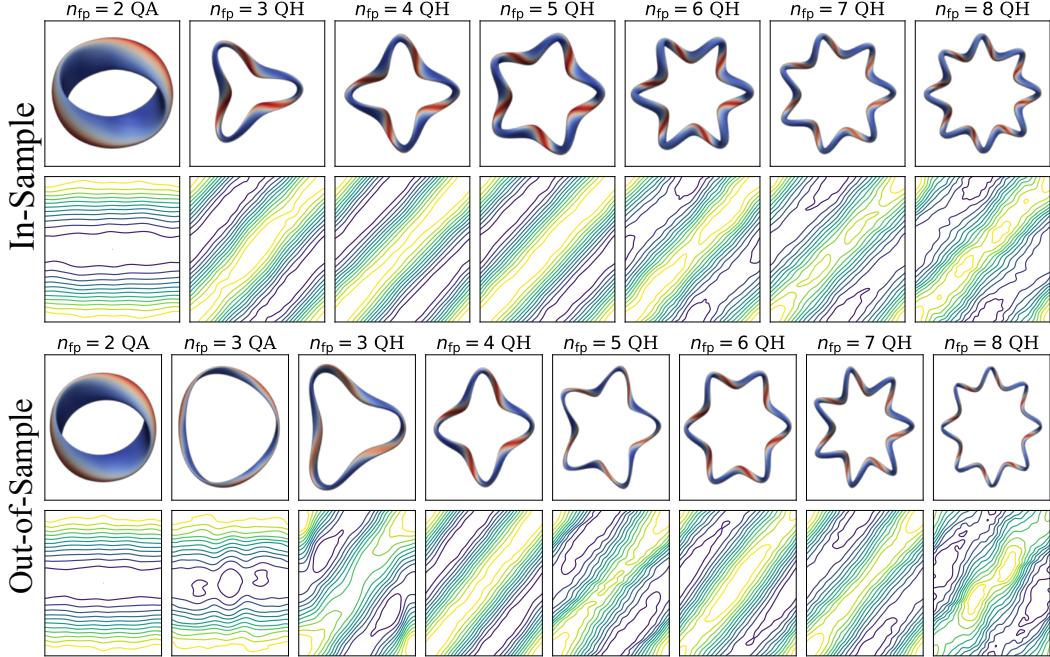


Figure 3: Stellarators generated by the diffusion model using in-sample conditions (first row) and out-of-sample conditions (third row). The contours of the magnetic field strength in Boozer coordinates of the in-sample devices (second row) and out-of-sample devices (fourth row), demonstrating the quality of quasisymmetry (perfectly straight contours would indicate perfect quasisymmetry) – all in-sample (out-of-sample) devices have less than 2.5% (6%) deviation from quasisymmetry, and 5% constraint violation, c_ℓ, c_A . Conditions used to generate the devices are shown in Table 1. Each out-of-sample device shown here corresponds to a blue marker on Figure 2 with identical label.

3.1 Generation Results

In this section, we examine the performance of the diffusion model by evaluating the metrics c_ℓ, c_A, J_{QS} on samples \mathbf{x} drawn from the diffusion model. We consider two settings: the “in-sample” setting where conditions \mathbf{y} are taken directly from the training data, and the “out-of-sample” setting where the conditions \mathbf{y} do *not* exist in the training data. Devices sampled with out-of-sample conditions are fundamentally new stellarator designs.

Figure 3 (top two rows) shows seven devices, each with a different number of field periods, generated by the diffusion model using conditions taken directly from the training data, while Figure 3 (bottom two rows) shows eight devices generated by the diffusion model using out-of-sample conditions. Figure 3 (rows 1 and 3) show top-down images of the generated stellarators, where the surfaces are colored by the field strength. Figure 3 (rows 2 and 4) show the contours of the magnetic field strength in Boozer coordinates on the surface. The near straight contours indicate good quasisymmetry in all devices: all in-sample devices have less than 2.5% deviation from quasisymmetry, and out-of-sample devices have less than 6% deviation from quasisymmetry. The good out-of-sample performance indicates the model learned the symmetry from the training data, and was able to extrapolate it to out-of-sample devices. Figure 3 illustrates that the diffusion model can generate sensible stellarators.

Figure 4 precisely quantifies the performance of the diffusion model in generating stellarators from *in-sample* conditions. Each frame shows box plots for a different metric: J_{QS} (left), c_A (middle), c_ℓ (right). In each frame, the tick labels on the x -axis are used to indicate different subsets of the data, e.g. the boxes above $n_{fp} = 2$ show the performance of DDPM when only given conditions with $n_{fp} = 2$. The left frame contrasts the distribution of J_{QS} for devices generated by the diffusion model (gold boxes labeled DDPM) to that of devices from the training data (red boxes labeled QUASR).

The devices generated by DDPM have similar, though somewhat worse quasisymmetry than the devices in QUASR, but are still near the target value of 1% (black dashed line). The middle and right frames show c_A , c_t for devices drawn from DDPM. For all subsets of the data, the aspect ratio of generated devices is typically less than 5% from the aspect ratio condition. On the other hand, \bar{t} tends to have higher error from the condition, coming within 12% of the condition, 75% of the time, for all subsets of the data. Overall, the in-sample performance of the diffusion model is satisfactory, hitting the target error values the majority of the time, while leaving room for improvement in future work.

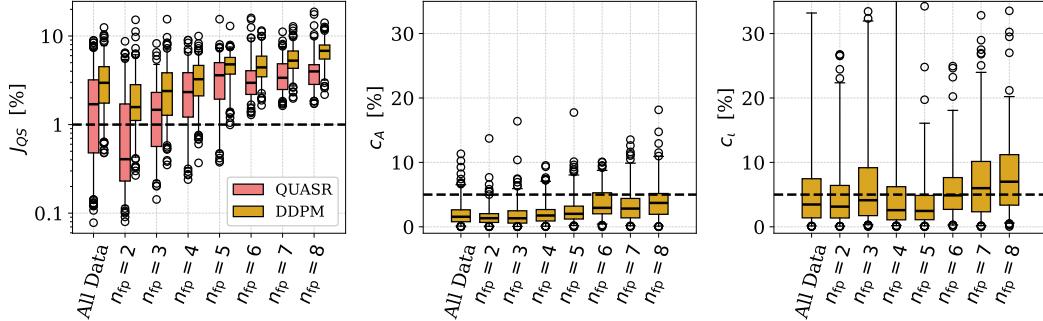


Figure 4: Performance of stellarators sampled by the diffusion model given in-sample conditions. The labels $n_{fp} = k$ shows the performance of the subset samples with the condition $n_{fp} = k$. (Left) Quasisymmetry error of samples from DDPM compared to the quasisymmetry of stellarators in QUASR (labeled Actual). DDPM generates stellarators with levels of J_{QS} similar to that of QUASR, even though quasisymmetry was not used as a condition to train the diffusion model. (Middle/Right) Percentage error in the aspect ratio (middle) and mean rotational transform (right) from the aspect ratio (mean rotational transform) condition used to generate the sample. In all panes, the black dashed line denotes a desirable value of the metric: $J_{QS} \leq 1\%$ and $c_t, c_A \leq 5\%$.

Figure 5 precisely quantifies the performance of the diffusion model in generating stellarators, with the *out-of-sample* conditions from Table 1. Each figure shows box plots for a different metric: J_{QS} (left), c_A (middle), c_t (right). In each frame, the tick labels on the x -axis are used to indicate which of the eight out-of-sample conditions was used, e.g. the boxes above ($n_{fp} = 2$ QA) show the performance of DDPM when only given ($n_{fp} = 2, N = 0$) out-of-sample condition from Table 1. The left panel shows that the deviation from quasisymmetry for each of the eight conditions is less than 10%, and in some cases less than the 1% target. The middle panel shows that for all conditions, the aspect ratio of generated devices is seldom more than 5% from the aspect ratio condition. Similarly the right panel shows that the \bar{t} of generated devices tends to be similar to the condition, coming within 10% of the condition, 75% of the time, for all conditions except $n_{fp} = 3$, QH. While there may be room to improve the precision of the diffusion model, Figures 3 and 5 concretely demonstrate that quasisymmetric devices can indeed be generated out-of-sample.

4 Open Research Problems

We now discuss variants and extensions of (IP) that are open areas for research.

Dataset and Benchmark Curation Despite the availability of large-scale stellarator datasets, no standard benchmark exists to unify them and provide consistent evaluation protocols for method comparison. A promising direction is to curate such benchmark by considering the following diverse stellarator datasets across stellarator classes, data representations, and properties: the QUASR database [link] contains coils and magnetic surfaces of quasisymmetric stellarators; the Constellation dataset [8] contains surfaces for quasi-isodynamic stellarators; the databases generated in [41, 14] contain the near-axis representation of quasisymmetric stellarators; the omnigenity database [19] contains surfaces for omnigenous stellarators; and the database from [42] contains surfaces for stellarators with measures of turbulence.

Physics-informed Architecture, Optimization, and Sampling In our case study (Sec. 3), the diffusion model learned to generate approximately quasisymmetric stellarators purely from the training data. A promising direction is to incorporate physical constraints explicitly in the design of the

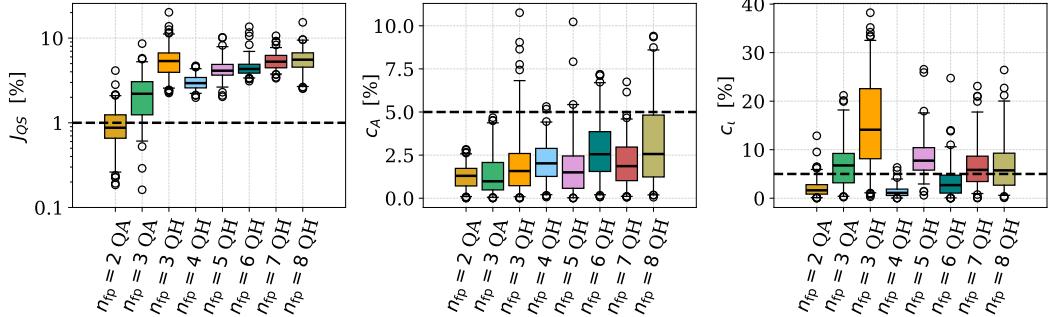


Figure 5: Performance of stellarators generated by the diffusion model given *out-of-sample* conditions, in terms of J_{QS} (left), c_A (middle), c_t (right). J_{QS} measures quasisymmetry deviation of devices, whereas c_A, c_t measure the relative error of A, \bar{t} from the condition used to generate the device. Black dashed lines indicate desirable value of the metric: $J_{QS} \leq 1\%$ and $c_t, c_A \leq 5\%$. Each box corresponds to a single condition from Table 1. The conditions are also plotted as a blue marker on Figure 2 with identical label.

generative model architecture (e.g. equivariant models), the loss function, or the sampling procedure. For example, replacing the non-differentiable evaluator VMEC, with a differentiable evaluator such as DESC [18] would enable the use of “physics informed loss functions” which penalize deviation from desired physical properties, i.e. penalize J_{QS}, c_A, c_t [5]. Similarly, derivatives of physics properties could be used in guided sampling procedures to ensure generated devices have, for instance, low J_{QS}, c_A, c_t [17]. Derivative information is also commonly available in the “single-stage” stellarator design setting [24, 23], where a generative model could be trained to generate a coil set-magnetic surface pair. Generation quality can also be improved by combining multiple data sources into a multi-fidelity generative model. For example, a multi-fidelity model for turbulent heat flux could combine low-fidelity analytic models, with sparse, high-fidelity data [42].

5 Acknowledgments

The author would like to thank the Flatiron Institute’s Scientific Computing Core for their support, in addition to Jiequn Han for his helpful feedback.

References

- [1] M. Asim, M. Daniels, O. Leong, A. Ahmed, and P. Hand. Invertible generative models for inverse problems: mitigating representation error and dataset bias. In *International Conference on Machine Learning*, pages 399–409. PMLR, 2020.
- [2] A. Baillod, J. Loizu, J. P. Graves, and M. Landreman. Stellarator optimization for nested magnetic surfaces at finite-beta and toroidal current. *Physics of Plasmas*, 29(4), 2022. Publisher: AIP Publishing.
- [3] A. Baillod, E. Paul, G. Rawlinson, M. Haque, S. Freiberger, and S. Thapa. Integrating novel stellarator single-stage optimization algorithms to design the Columbia stellarator experiment. *Nuclear Fusion*, 65(2):026046, 2025. Publisher: IOP Publishing.
- [4] S. Bannmann, O. Ford, U. Hoefel, P. Z. Poloskei, A. Pavone, S. Kwak, J. Svensson, S. Lazerson, P. McNeely, N. Rust, and others. Bayesian inference of electron density and ion temperature profiles from neutral beam and halo Balmer-alpha emission at Wendelstein 7-X. *Plasma Physics and Controlled Fusion*, 66(6):065001, 2024. Publisher: IOP Publishing.
- [5] J.-H. Bastek, W. Sun, and D. Kochmann. Physics-informed diffusion models. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu, editors, *International Conference on Representation Learning*, volume 2025, pages 3360–3385, 2025.

- [6] C. Bishop. *Pattern recognition and machine learning*. Information Science and Statistics. Springer New York, NY, Aug. 2006.
- [7] P. Bonofiglio, D. Dudit, and C. Swanson. Fast ion confinement in quasi-axisymmetric stellarator equilibria. *Nuclear Fusion*, 65(2):026050, 2025. Publisher: IOP Publishing.
- [8] S. A. Cadena, A. Merlo, E. Laude, A. Bauer, A. Agrawal, M. Pascu, M. Savtchouk, E. Guiraud, L. Bonauer, S. Hudson, and M. Kaiser. ConStellaration: A dataset of QI-like stellarator plasma boundaries and optimization benchmarks, June 2025. arXiv:2506.19583 [cs].
- [9] L. Carbajal, J. Varela, A. Bader, W. Guttenfelder, A. Cerfon, J. Schmitt, J. Morrissey, C. Hegna, J. Canik, N. Mandell, and others. Alpha-particle confinement in Infinity Two fusion pilot plant baseline plasma design. *Journal of Plasma Physics*, pages 1–30, 2025. Publisher: Cambridge University Press.
- [10] K. Choudhary, B. DeCost, C. Chen, A. Jain, F. Tavazza, R. Cohn, C. W. Park, A. Choudhary, A. Agrawal, S. J. Billinge, and others. Recent advances and applications of deep learning methods in materials science. *npj Computational Materials*, 8(1):59, 2022. Publisher: Nature Publishing Group UK London.
- [11] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, Sept. 2022.
- [12] H. Chung, B. Sim, D. Ryu, and J. C. Ye. Improving diffusion models for inverse problems using manifold constraints. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 25683–25696. Curran Associates, Inc., 2022.
- [13] B. Clavier, D. Zarzoso, D. del Castillo-Negrete, and E. Frénod. Generative-machine-learning surrogate model of plasma turbulence. *Physical Review E*, 111(1):L013202, 2025. Publisher: APS.
- [14] P. Curvo, D. Ferreira, and R. Jorge. Using deep learning to design high aspect ratio fusion devices. *Journal of Plasma Physics*, 91(1):E38, 2025. Publisher: Cambridge University Press.
- [15] J. Cândido and R. Jorge. Design of quasisymmetric fusion devices using novel machine learning methods. Master’s thesis, Instituto Superior Técnico, Universidade De Lisboa, 2023.
- [16] G. Daras, H. Chung, C.-H. Lai, Y. Mitsuishi, J. C. Ye, P. Milanfar, A. G. Dimakis, and M. Delbracio. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*, 2024.
- [17] P. Dhariwal and A. Nichol. Diffusion models beat GANs on image synthesis. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, NIPS ’21, pages 8780–8794, Red Hook, NY, USA, Dec. 2021. Curran Associates Inc.
- [18] D. Dudit and E. Kolemen. DESC: a stellarator equilibrium solver. *Physics of Plasmas*, 27(10), 2020. Publisher: AIP Publishing.
- [19] R. Gaur, R. Conlin, D. Dickinson, J. F. Parisi, D. Dudit, D. Panici, P. Kim, K. Unalmis, W. D. Dorland, and E. Kolemen. Omnidirectional stellarator equilibria with enhanced stability, Oct. 2024. arXiv:2410.04576 [physics].
- [20] A. 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. Publisher: Cambridge University Press.
- [21] A. Giuliani, E. Rodríguez, and M. Spivak. A comprehensive exploration of quasisymmetric stellarators and their coil sets, Apr. 2025. arXiv:2409.04826 [physics].
- [22] A. Giuliani, F. Wechsung, A. Cerfon, M. Landreman, and G. Stadler. Direct stellarator coil optimization for nested magnetic surfaces with precise quasi-symmetry. *Physics of Plasmas*, 30(4), 2023. Publisher: AIP Publishing.

[23] A. Giuliani, F. Wechsung, A. Cerfon, G. Stadler, and M. Landreman. Single-stage gradient-based stellarator coil design: optimization for near-axis quasi-symmetry. *Journal of Computational Physics*, 459:111147, 2022. Publisher: Academic Press.

[24] A. Giuliani, F. Wechsung, G. Stadler, A. Cerfon, and M. Landreman. Direct computation of magnetic surfaces in Boozer coordinates and coil optimization for quasisymmetry. *Journal of Plasma Physics*, 88(4):905880401, 2022. Publisher: Cambridge University Press.

[25] H. Goh, S. Sheriffdeen, J. Wittmer, and T. Bui-Thanh. Solving Bayesian inverse problems via variational autoencoders. *arXiv preprint arXiv:1912.04212*, 2019.

[26] R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018. Publisher: ACS Publications.

[27] C. Hegna, D. Anderson, E. Andrew, A. Ayilaran, A. Bader, T. Bohm, K. C. Mata, J. Canik, L. Carbajal, A. Cerfon, and others. The Infinity Two fusion pilot plant baseline plasma physics design. *Journal of Plasma Physics*, 91(3):E76, 2025. Publisher: Cambridge University Press.

[28] P. Helander. Theory of plasma confinement in non-axisymmetric magnetic fields. *Reports on Progress in Physics*, 77(8):087001, 2014. Publisher: IOP Publishing.

[29] F. Hindenlang, G. G. Plunk, and O. Maj. Computing MHD equilibria of stellarators with a flexible coordinate frame. *Plasma Physics and Controlled Fusion*, 67(4):045002, 2025. Publisher: IOP Publishing.

[30] S. Hirshman and O. Betancourt. Preconditioned descent algorithm for rapid calculations of magnetohydrodynamic equilibria. *Journal of Computational Physics*, 96(1):99–109, 1991. Publisher: Elsevier.

[31] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020.

[32] S. Hudson, R. Dewar, G. Dennis, M. Hole, M. McGann, G. Von Nessi, and S. Lazerson. Computation of multi-region relaxed magnetohydrodynamic equilibria. *Physics of Plasmas*, 19(11), 2012. Publisher: AIP Publishing.

[33] L.-M. Imbert-Gérard, E. J. Paul, and A. M. Wright. *An introduction to stellarators: from magnetic fields to symmetries and optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2024.

[34] R. Jorge, A. Goodman, M. Landreman, J. Rodrigues, and F. Wechsung. Single-stage stellarator optimization: combining coils with fixed boundary equilibria. *Plasma Physics and Controlled Fusion*, 65(7):074003, 2023. Publisher: IOP Publishing.

[35] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, and others. Highly accurate protein structure prediction with AlphaFold. *nature*, 596(7873):583–589, 2021. Publisher: Nature Publishing Group UK London.

[36] K. Karchev, N. A. Montel, A. Coogan, and C. Weniger. Strong-lensing source reconstruction with denoising diffusion restoration models, Nov. 2022. arXiv:2211.04365 [astro-ph].

[37] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, Jan. 2017. arXiv:1412.6980 [cs].

[38] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro. DiffWave: a versatile diffusion model for audio synthesis. In *International Conference on Learning Representations 2021*, Oct. 2020.

[39] S. Kwak, U. Hoefel, M. Krychowiak, A. Langenberg, J. Svensson, H. T. Mora, Y.-C. Ghim, and others. Bayesian modelling of multiple plasma diagnostics at Wendelstein 7-X. *Nuclear Fusion*, 64(10):106022, 2024. Publisher: IOP Publishing.

[40] R. Laia, R. Jorge, and G. Abreu. Data-driven approach to model the influence of magnetic geometry in the confinement of fusion devices, July 2025. arXiv:2507.03776 [physics].

[41] M. Landreman. Mapping the space of quasisymmetric stellarators using optimized near-axis expansion. *Journal of Plasma Physics*, 88(6):905880616, 2022. Publisher: Cambridge University Press.

[42] M. Landreman, J. Y. Choi, C. Alves, P. Balaprakash, M. Churchill, R. Conlin, and G. Roberg-Clark. How does ion temperature gradient turbulence depend on magnetic geometry? Insights from data and machine learning. *Journal of Plasma Physics*, 91(4):E120, Aug. 2025.

[43] M. Landreman, B. Medasani, F. Wechsung, A. Giuliani, R. Jorge, and C. Zhu. SIMSOPT: a flexible framework for stellarator optimization. *Journal of Open Source Software*, 6(65):3525, 2021.

[44] M. Landreman and E. Paul. Magnetic fields with precise quasisymmetry for plasma confinement. *Physical Review Letters*, 128(3):035001, 2022. Publisher: American Physical Society.

[45] C. Liu, C. Wu, S. Cao, M. Chen, J. C. Liang, A. Li, M. Huang, C. Ren, Y. N. Wu, D. Liu, and T. Geng. Diff-PIC: revolutionizing particle-in-cell nuclear fusion simulation with diffusion models. In *The Thirteenth International Conference on Learning Representations*, Oct. 2024.

[46] A. Merchant, S. Batzner, S. S. Schoenholz, M. Aykol, G. Cheon, and E. D. Cubuk. Scaling deep learning for materials discovery. *Nature*, 624(7990):80–85, 2023. Publisher: Nature Publishing Group UK London.

[47] A. Merlo, D. Böckenhoff, J. Schilling, U. Höfel, S. Kwak, J. Svensson, A. Pavone, S. A. Lazerson, T. S. Pedersen, and others. Proof of concept of a fast surrogate model of the VMEC code via neural networks in Wendelstein 7-X scenarios. *Nuclear Fusion*, 61(9):096039, 2021. Publisher: IOP Publishing.

[48] A. Merlo, D. Böckenhoff, J. Schilling, S. A. Lazerson, T. S. Pedersen, and others. Physics-regularized neural network of the ideal-MHD solution operator in Wendelstein 7-X configurations. *Nuclear Fusion*, 63(6):066020, 2023. Publisher: IOP Publishing.

[49] K. Moriwaki, N. Filipová, M. Shirasaki, and N. Yoshida. Deep learning for intensity mapping observations: component extraction. *Monthly Notices of the Royal Astronomical Society: Letters*, 496(1):L54–L58, 2020. Publisher: Oxford University Press.

[50] N. Mudur and D. P. Finkbeiner. Can denoising diffusion probabilistic models generate realistic astrophysical fields? *arXiv preprint arXiv:2211.12444*, 2022.

[51] F. Najmabadi, A. Raffray, S. Abdel-Khalik, L. Bromberg, L. Crosatti, L. El-Guebaly, P. Garabedian, A. Grossman, D. Henderson, A. Ibrahim, and others. The ARIES-CS compact stellarator fusion power plant. *Fusion Science and Technology*, 54(3):655–672, 2008. Publisher: Taylor & Francis.

[52] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8162–8171. PMLR, July 2021. ISSN: 2640-3498.

[53] D. Nieuwenhuizen. *Modeling nuclear fusion reactors using variational auto-encoders*. PhD thesis, Eindhoven University of Technology, 2021.

[54] M. Padidar, T. Huang, A. Giuliani, and M. Spivak. Diffusion for Fusion: Designing Stellarators with Generative AI. Sept. 2025. DOI:10.5281/zenodo.17713019.

[55] A. Pavone, A. Merlo, S. Kwak, and J. Svensson. Machine learning and Bayesian inference in nuclear fusion research: an overview. *Plasma Physics and Controlled Fusion*, 65(5):053001, 2023. Publisher: IOP Publishing.

[56] B. Remy. *Generative modeling for weak lensing inverse problems*. PhD Thesis, Université Paris-Saclay, 2023.

[57] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, New Orleans, LA, USA, June 2022. IEEE.

[58] W. Sengupta, R. Madan, S. Buller, N. Nikulsin, E. J. Paul, R. Nies, A. A. Kaptanoglu, S. R. Hudson, and A. Bhattacharjee. Periodic Korteweg-de Vries soliton potentials generate quasisymmetric magnetic field strength in a finite plasma-beta equilibrium, July 2025. arXiv:2507.06480 [physics].

[59] M. Szűcs, T. Szepesi, C. Biedermann, G. Cseh, M. Jakubowski, G. Kocsis, R. König, M. Krause, V. Perseo, A. Puig Sitjes, and others. Detecting plasma detachment in the Wendelstein 7-X stellarator using machine learning. *Applied Sciences*, 12(1):269, 2021. Publisher: MDPI.

[60] L. van Rijn. *Minimizing neoclassical transport in the Wendelstein 7-X stellarator using Variational autoencoders*. PhD Thesis, Eindhoven University of Technology, 2022.

[61] J. Vos. *Discovery of hidden neoclassical transport variables in Wendelstein 7-X through variational autoencoder latent space exploration*. PhD Thesis, Eindhoven University of Technology, 2024.

[62] A. Wakasa, S. Murakami, M. Itagaki, and S.-i. Oikawa. Construction of neoclassical transport database for large helical device plasma applying neural network method. *Japanese Journal of Applied Physics*, 46(3R):1157, 2007. Publisher: IOP Publishing.

[63] C. Wan, Y. Cho, Z. Qu, Y. Camenen, R. Varennes, K. Lim, K. Li, J. Li, Y. Li, and X. Garbet. A high-fidelity surrogate model for the ion temperature gradient (ITG) instability using a small expensive simulation dataset. *Nuclear Fusion*, 65(5):054001, 2025. Publisher: IOP Publishing.

[64] J. L. Watson, D. Juergens, N. R. Bennett, B. L. Trippe, J. Yim, H. E. Eisenach, W. Ahern, A. J. Borst, R. J. Ragotte, L. F. Milles, et al. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.

[65] F. Wechsung, M. Landreman, A. Giuliani, A. Cerfon, and G. Stadler. Precise stellarator quasi-symmetry can be achieved with electromagnetic coils. *Proceedings of the National Academy of Sciences*, 119(13):e2202084119, 2022. Publisher: National Academy of Sciences.

[66] A. Wiedman, S. Buller, and M. Landreman. Coil optimization for quasi-helically symmetric stellarator configurations. *Journal of Plasma Physics*, 90(3):905900307, 2024. Publisher: Cambridge University Press.

[67] D. Zarzoso, S. Pinches, A. Bustos, A. Cappa, and E. Ascasibar. A novel unsupervised machine learning algorithm for automatic Alfvénic activity detection in the TJ-II stellarator. *Nuclear Fusion*, 64(12):126057, 2024. Publisher: IOP Publishing.

[68] X. Zhao, Y.-S. Ting, K. Diao, and Y. Mao. Can diffusion model conditionally generate astrophysical images? *Monthly Notices of the Royal Astronomical Society*, 526(2):1699–1712, 2023. Publisher: Oxford University Press.

[69] Z. Zhao, J. C. Ye, and Y. Bresler. Generative models for inverse imaging problems: from mathematical foundations to physics-driven applications. *IEEE Signal Processing Magazine*, 40(1):148–163, 2023. Publisher: IEEE.

[70] C. Zhu, S. R. Hudson, Y. Song, and Y. Wan. New method to design stellarator coils without the winding surface. *Nuclear Fusion*, 58(1):016008, 2017. Publisher: IOP Publishing.

A Description and Evaluation of Stellarator Designs

In this section, we discuss the representation of surfaces and evaluation of J_{QS} , c_A , c_ℓ . Surfaces can be constructed with SIMSOPT [43] and evaluated with VMEC. Some stellarators in the training data, or generated by the diffusion model, have highly shaped boundaries. VMEC struggles to evaluate these, failing to return a solution for approximately 41% of stellarators from the training data, and 57% of stellarators generated in-sample by the diffusion model. In reporting results, we ignore these designs.

Surface Parameterization We follow [24, 20] in describing surfaces with the `SurfaceXYZTensorFourier` format in SIMSOPT. The vector \mathbf{x} is the set of Fourier coefficients x_{ij}, y_{ij}, z_{ij} used in Equations (5) to (7), with ordering dictated by SIMSOPT. Stellarator symmetry is assumed for all surfaces in training and evaluation. The Cartesian components of a n_{fp} -field period surface, $x(\varphi, \theta), y(\varphi, \theta), z(\varphi, \theta)$, can be constructed with the Fourier series,

$$\hat{x}(\varphi, \theta) = \sum_{i=0}^{m_{\text{pol}}} \sum_{j=0}^{n_{\text{tor}}} x_{ij} w_i(\theta) v_j(\varphi) + \sum_{i=m_{\text{pol}}+1}^{2m_{\text{pol}}} \sum_{j=n_{\text{tor}}+1}^{2n_{\text{tor}}} x_{ij} w_i(\theta) v_j(\varphi) \quad (3)$$

$$\hat{y}(\varphi, \theta) = \sum_{i=0}^{m_{\text{pol}}} \sum_{j=n_{\text{tor}}+1}^{2n_{\text{tor}}} y_{ij} w_i(\theta) v_j(\varphi) + \sum_{i=m_{\text{pol}}+1}^{2m_{\text{pol}}} \sum_{j=0}^{n_{\text{tor}}} y_{ij} w_i(\theta) v_j(\varphi) \quad (4)$$

$$x(\varphi, \theta) = \hat{x}(\varphi, \theta) \cos(\varphi) - \hat{y}(\varphi, \theta) \sin(\varphi) \quad (5)$$

$$y(\varphi, \theta) = \hat{x}(\varphi, \theta) \sin(\varphi) + \hat{y}(\varphi, \theta) \cos(\varphi) \quad (6)$$

$$z(\varphi, \theta) = \sum_{i=0}^{m_{\text{pol}}} \sum_{j=n_{\text{tor}}+1}^{2n_{\text{tor}}} z_{ij} w_i(\theta) v_j(\varphi) + \sum_{i=m_{\text{pol}}+1}^{2m_{\text{pol}}} \sum_{j=0}^{n_{\text{tor}}} z_{ij} w_i(\theta) v_j(\varphi) \quad (7)$$

where the basis functions v_j are,

$$\{1, \cos(1 nfp \varphi), \dots, \cos(n_{\text{tor}} nfp \varphi), \sin(1 nfp \varphi), \dots, \sin(n_{\text{tor}} nfp \varphi)\}, \quad (8)$$

and the basis functions w_i are,

$$\{1, \cos(1\theta), \dots, \cos(m_{\text{pol}}\theta), \sin(1\theta), \dots, \sin(m_{\text{pol}}\theta)\}. \quad (9)$$

The training data dictated that $m_{\text{pol}} = n_{\text{tor}} = 10$.

Quasisymmetry Objective The quasisymmetry objective, J_{QS} , can be evaluated from a magnetic field, \mathbf{B} , and surface parameterized in Boozer coordinates θ, φ . The J_{QS} objective from [24] computes the relative distance between the quasisymmetric and non-quasisymmetric components of the field strength $B(\varphi, \theta) = \|\mathbf{B}(\varphi, \theta)\|_2$,

$$J_{QS} = \sqrt{\frac{\int_{\Gamma_s} B_{\text{non-QS}}^2 dS}{\int_{\Gamma_s} B_{\text{QS}}^2 dS}}. \quad (10)$$

B_{QS} is the projection of B onto the space of quasisymmetric fields. For quasi-axisymmetry,

$$B_{\text{QS}} = \frac{\int_0^1 B \|\mathbf{n}\| d\varphi}{\int_0^1 \|\mathbf{n}\| d\varphi} \quad (11)$$

$$B_{\text{non-QS}} = B - B_{\text{QS}} \quad (12)$$

For quasihelical symmetry, the integral is taken over the helical angle $\theta - N n_{fp} \varphi$. When J_{QS} is zero, the magnetic field is perfectly quasisymmetric on the magnetic surface.

Evaluation The metrics J_{QS}, c_ℓ, c_A are evaluated with the following protocol. Given n_{fp} and \mathbf{x} a surface is object is constructed in SIMSOPT. The surface is evaluated with VMEC, from which \bar{A} and \bar{A} can be computed. Boozer coordinates are constructed using the `BoozXform` in SIMSOPT. B and the jacobian \sqrt{g} are constructed in Boozer coordinates, from which J_{QS} can be evaluated.

B Diffusion Model Details

The *conditional* diffusion model is an instance of the DDPM model [31] where the neural network model, f_θ , is modified to enable conditions. Training and sampling followed the standard DDPM training and sampling algorithms, Algorithms 1 and 2. Algorithm 1 requires as input a sequence of step sizes $\{\alpha_t\}_{t=1}^T$ and a neural network model $f_\theta(\mathbf{x}_t, t, \mathbf{y})$ which predicts a noisy sequence given a (noised) feature \mathbf{x}_t , timestep t , and condition \mathbf{y} . The conditions and noise-free features $\mathbf{y} \sim q(\mathbf{y})$, $\mathbf{x}_0 \sim q(\mathbf{x}_0 | \mathbf{y})$ are drawn from the training data.

The model was trained with $n_y = 4$ features, $\mathbf{y} = (\bar{t}, A, n_{fp}, N)$. The diffusion model used 200 timesteps, with a linear schedule. A standard feed-forward neural network architecture was used to describe $f_\theta(\mathbf{x}_t, t, \mathbf{y})$ with 4 hidden layers, each with width 2048, and Gaussian error linear unit (GELU) activation functions. A sinusoidal network head was used to embed $\mathbf{x}, t, \mathbf{y}$ before entering the standard feed-forward layers. The network head maps \mathbf{x} to dimension 64, t to dimension 128, and \mathbf{y} to dimension 128. The Adam optimizer [37] was used with a training batch size of 4096, an evaluation batch size of 128, for 250 epochs, with a learning rate of 0.0005.

Algorithm 1: Training f_θ

Input: Step size sequence $\{\alpha_t\}_{t=1}^T$

- 1 **repeat**
- 2 $\mathbf{y} \sim q(\mathbf{y})$,
 $\mathbf{x}_0 \sim q(\mathbf{x}_0 | \mathbf{y})$, $\mathbf{z} \sim \mathcal{N}(0, I)$;
- 3 $t \sim \text{Uniform}(\{1, \dots, T\})$;
- 4 Take a gradient descent step using

$$\nabla_\theta \|\mathbf{z} - f_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \mathbf{z}, t, \mathbf{y})\|^2$$
- 5 **until** converged;

Algorithm 2: Sampling from trained f_θ

Input: f_θ , step sizes $\{\alpha_t\}_{t=1}^T$, variances $\{\sigma_t\}_{t=1}^T$, conditions \mathbf{y} .

- 1 $\mathbf{x}_T \sim \mathcal{N}(0, I)$;
- 2 **for** $t = T, \dots, 1$ **do**
- 3 $\mathbf{z} \sim \mathcal{N}(0, I)$ if $t > 1$ else $\mathbf{z} = 0$;
- 4
$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} f_\theta(\mathbf{x}_t, t, \mathbf{y})) + \sigma_t \mathbf{z}$$
- 5 **end**
- 6 **return** \mathbf{x}_0

Dimensionality reduction Prior to training the diffusion model, we reduce the dimensionality of the raw training data to eliminate noise and high frequency oscillations in the dataset that do not have a meaningful connection to stellarator performance. Using principal component analysis (PCA), we project the data points, \mathbf{x} , onto the linear subspace of dimension $n_r < n_x$ which captures largest fraction of the data variance. As the dimension of the embedding reduces, so does the quality of the devices in the embedded dataset. Missing important information in the training set, the diffusion model is limited in what it can learn, and the quality of stellarators it will generate. Figure 6 shows the reduction in the quality of stellarators in the dataset as the embedding dimension decreases. There is not a significant loss in performance of devices for the embedding dimension used for training, $n_r = 50$. Prior to evaluating a sample from the diffusion model, the sample must be projected back up to the full n_x -dimensional space.

C Additional Related Work

Besides the work on (IP), there have been many approaches to use ML to understand fusion plasmas. [42, 58] used symbolic regression to understand turbulence and quasisymmetry in stellarators. [40] explored correlations in a dataset of stellarators and developed predictive models for quasisymmetry. [45] used diffusion models to accelerate particle-in-cell simulations. [13] used a variational autoencoder (VAE) to generate time-varying solutions to the Hasagawa-Wakatani equations for turbulence. [53] used VAEs to estimate hidden parameters that influence tokamak operation. [63] pre-trained a model for tokamak turbulence on low-fidelity data and fine-tune on high-fidelity data, developing an multi-fidelity surrogate for turbulence simulations. [60, 61] trained a VAE as a surrogate for particle and heat flux. [62] constructed a neural network surrogate for neoclassical transport coefficients in the LHD device. Physics-regularized neural networks have been used to predict solutions to the ideal MHD equations in W7-X geometries [48, 47]. [55] reviewed applications of Bayesian and ML methods in fusion research, spanning methods for plasma control and disruption prediction, to surrogates for diagnostics. [39, 4] used Bayesian modeling to improve plasma diagnostics, and [59, 67] built ML models for detect plasma activity in experiments.

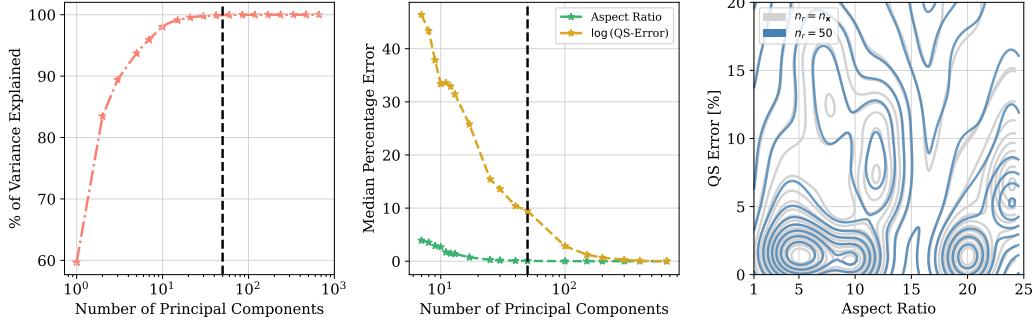


Figure 6: (Left) The fraction of the variance of \mathbf{x} explained by PCA with a given embedding dimension, n_r . The black dashed line marks an embedding dimension of 50. (Middle) The median percentage error of $\log(J_{QS})$ and aspect ratio between embedded devices and devices in the dataset. The properties of devices can change significantly as the dimension of the embedding is reduced – the grey region highlights areas where the *max* relative error in the aspect ratio is at least 20%. (Right) Joint probability density of the quasi-symmetry error and aspect ratio of devices embedded in $n_r = 50$ dimensions (blue) versus devices in the actual dataset (grey). Using $n_r = 50$ ensures that significant information is *not* lost by dimensionality reduction.

D Data for Numerical Experiments

This section contains additional data relevant to the numerical experiments. Table 1 contains the sampling conditions, \mathbf{y} , used to generate the samples in Fig. 3 and Fig. 5.

	In-sample conditions							Out-of-sample conditions							
	2	3	4	5	6	7	8	2	3	3	4	5	6	7	8
n_{fp}	2	3	4	5	6	7	8	2	3	3	4	5	6	7	8
N	0	1	1	1	1	1	1	0	0	1	1	1	1	1	1
A	4.0	8.0	8.0	8.0	8.0	12.0	12.0	4.5	18.5	9.0	11.0	17.0	14.0	11.0	22.0
\bar{t}	0.3	1.3	1.6	1.6	2.6	3.0	3.5	0.36	0.5	0.5	1.4	2.5	2.0	3.7	3.5

Table 1: In-sample conditions and out-of-sample conditions, \mathbf{y} , used for generating stellarators in Fig. 3 and 5.