

A comprehensive exploration of quasisymmetric stellarators and their coil sets

Andrew Giuliani¹ , Eduardo Rodríguez²  and Marina Spivak¹

¹Center for Computational Mathematics, Flatiron Institute, 162 Fifth Avenue, New York 10128, USA

²Max Planck Institute of Plasma Physics at Greifswald, Wendelsteinstraße 1, Greifswald 17491, Germany

Corresponding author: Andrew Giuliani; agiuliani@flatironinstitute.org

(Received 7 September 2024; revision received 21 April 2025; accepted 22 April 2025)

We augment the ‘quasisymmetric stellarator repository’ (QUASR) to include vacuum field stellarators with quasihelical symmetry using a globalized optimization workflow. The database now has over 300 000 quasisaxisymmetry and quasihelically symmetric devices along with coil sets, optimized for a variety of aspect ratios, rotational transforms and discrete rotational symmetries. This paper outlines a couple of ways to explore and characterize the data set. We plot devices on a near-axis quasisymmetry (QS) landscape, revealing close correspondence to this predicted landscape. We also use principal component analysis (PCA) to reduce the dimensionality of the data so that it can easily be visualized in two or three dimensions. The PCA also gives a mechanism to compare the new devices here with previously published ones in the literature. We are able to characterize the structure of the data, observe clusters and visualize the progression of devices in these clusters. The topology of the data are governed by the interplay of the design constraints and valleys of the QS objective. These techniques reveal that the data has structure, and that typically one, two or three principal components are sufficient to characterize it. The latest version of QUASR is archived at <https://zenodo.org/doi/10.5281/zenodo.10050655> and can be explored online at quasr.flatironinstitute.org.

Key words: plasma confinement, plasma devices, plasma simulation

1. Introduction

In recent years, the stellarator community has developed novel numerical optimization algorithms to quickly design stellarators that are quasisymmetric. Quasisymmetry (QS) is a property of the magnetic field, where the field strength on a given magnetic surface can be written as a linear combination of the Boozer angles (Boozer 1981*a*) on the surface (Helander 2014). It can be shown that such fields confine guiding centre orbits (Boozer 1981*b*; Nührenberg & Zille 1988; Burby, Kallinikos & MacKay 2020; Rodríguez, Helander & Bhattacharjee 2020). Gradient-based optimization algorithms have facilitated the design of magnetic fields with the

QS property (Landreman & Paul 2022), as well as the electromagnetic coils that generate them (Wechsung *et al.* 2022b).

There is interest in the stellarator community to explore the landscape of QS stellarators. Examples of such surveys can be found in the literature, including a set of four stellarators with highly precise volumetric QS in Landreman & Paul (2022), a similarly generated data set of more than 150 devices by Buller *et al.* (2024) and one of over 30 generated exploiting adjoint-based algorithms by Nies *et al.* (2024). An even larger data set of 500 000 stellarators was presented in Landreman (2022), which leverages a model of quasisymmetric stellarators near their magnetic axes (Garren & Boozer 1991; Landreman, Sengupta & Plunk 2019; Landreman & Sengupta 2019; Rodríguez *et al.* 2023). The devices in the data sets mentioned above all were released without the coil sets that generate them. For data sets that include their coils we must turn to the examples of over 40 stellarator coil sets studied by Kappel, Landreman & Malhotra (2024) or the work of Jorge, Giuliani & Loizu (2024), which includes a set of over 10 devices with simple and flexible coils.

A globalized coil design workflow was recently described in Giuliani (2024), and was used to generate a large database of quasisymmetric (QA) stellarators and their associated coil sets. The database is called QUASR for ‘quasisymmetric stellarator repository’ and can be explored online at quasr.flatironinstitute.org. In this work, we extend these algorithms to search for quasihelically (QH) symmetric devices and augment QUASR so that it now contains over 300 000 devices with associated coil sets. Using this large data set, we explore the landscape of stellarators that are generated by modular electromagnetic coils. To start, we present a few devices in QUASR and analyse their physics properties. We then apply a variety of techniques to visualize and analyse this large data set. The first uses the near-axis model and its landscape of excellent and poor QS (Rodríguez, Sengupta & Bhattacharjee 2022b; Rodríguez *et al.* 2023), which we call the near-axis QS landscape, to frame some configurations in the data set. The second, and as an alternative, we employ a dimensionality reduction approach called principal component analysis (PCA) (Jolliffe 1990). This is a well-known technique frequently used in data science for dimensionality reduction and visualization (Bishop 2006). It enables the practitioner to project the high-dimensional data set onto a linear manifold, e.g. a two-dimensional plane or three-dimensional volume, to visualize relationships and possible formation of clusters. We compute two quantitative scores that measure how faithful the lower-dimensional representation is to the higher-dimensional data.

To summarize, the main goals of this paper are twofold: (i) we extend QUASR to include both QA and QH stellarators (§§ 2 and 3); and (ii) we use a couple of visualization techniques to explore the data set and gain insight on the devices in it (§§ 4 and 5).

2. The coil design workflow

Our goal is to find coil geometries and currents that produce magnetic fields with the QS property (either QA or QH). To achieve this, we use a globalized coil design workflow that comprises three phases, wrapped in a globalization algorithm. Since this workflow and its constituent algorithms have been described before (Eriksson *et al.* 2019; Giuliani *et al.* 2022a,b; Giuliani *et al.* 2023, Giuliani 2024), we only provide a brief overview here, summarized in figure 1. The first phase searches for stellarator coils that are quasisymmetric near their axis. This algorithm is robust, but may only find devices with nested flux surfaces on a small region in the




FIGURE 1. Globalized coil design workflow.

neighbourhood of the magnetic axis. These devices are then provided as an initial guess to Phases II and III which heal generalized chaos, optimize for nested flux surfaces and polish for precise QS. The three phases are wrapped in a globalization algorithm (Eriksson *et al.* 2019), ensuring that the objective landscape is sufficiently well-explored.

To proceed with the volumetric refinement of QS of Phases II and III, we must define some measure f_{QS} for the field generated by the electromagnetic coils to minimize. To do so, we must take into account the two flavours of QS we are interested in: quasisymmetry and quasihelical symmetry. The approaches for both symmetries are very similar, so we explain the process for quasisymmetry, and then show how it can be extended to quasihelical symmetry.

Quasisymmetry is characterized by θ isolines of $B(\varphi, \theta) = \|\mathbf{B}(\boldsymbol{\Sigma}_s(\varphi, \theta))\|$, on magnetic flux surfaces $\boldsymbol{\Sigma}_s$ parametrized in Boozer coordinates (φ, θ) (Boozer 1981a). The field's deviation from exact quasisymmetry on the magnetic surface can be computed by evaluating

$$f_{\text{QS}} = \frac{\int_0^{1/n_{\text{fp}}} \int_0^1 (B(\varphi, \theta) - B_{\text{QA}}(\theta))^2 \|((\partial \boldsymbol{\Sigma}_s)/(\partial \varphi)) \times ((\partial \boldsymbol{\Sigma}_s)/(\partial \theta))\| \, d\theta \, d\varphi}{\int_0^{1/n_{\text{fp}}} \int_0^1 B_{\text{QA}}(\varphi, \theta)^2 \|((\partial \boldsymbol{\Sigma}_s)/(\partial \varphi)) \times ((\partial \boldsymbol{\Sigma}_s)/(\partial \theta))\| \, d\theta \, d\varphi}, \quad (2.1)$$

where

$$B_{\text{QA}}(\theta) = \frac{\int_0^{1/n_{\text{fp}}} B(\varphi, \theta) \|((\partial \boldsymbol{\Sigma}_s)/(\partial \varphi)) \times ((\partial \boldsymbol{\Sigma}_s)/(\partial \theta))\| \, d\varphi}{\int_0^{1/n_{\text{fp}}} \|((\partial \boldsymbol{\Sigma}_s)/(\partial \varphi)) \times ((\partial \boldsymbol{\Sigma}_s)/(\partial \theta))\| \, d\varphi} \quad (2.2)$$

is a least squares projection of $B(\varphi, \theta)$ onto the space of functions that do not depend on φ . These integrals can easily be evaluated using the quadrature points on the tensor product grid $(\varphi_i, \theta_j) \in [0, 1/n_{\text{fp}}) \times [0, 1)$, where $\varphi_i = i/(N_\varphi n_{\text{fp}})$ and $\theta_j = j/N_\theta$ for $i = 0, \dots, N_\varphi - 1$, $j = 0, \dots, N_\theta - 1$ and N_φ, N_θ are the number of quadrature points in the toroidal and poloidal directions, respectively. These points correspond to the periodic trapezoidal rule, which is spectrally accurate (Trefethen & Weideman 2014). This measure of QS is similar to the commonly used sum over the symmetry breaking components of B (defined as f_B in Rodriguez, Paul & Bhattacharjee (2022a)). To extend f_{QS} to the QH symmetric case, we note that the difference is in the isolines of $B(\varphi, \theta)$, which now correspond to lines of constant $\theta - N n_{\text{fp}} \varphi$, for $N \in \mathbb{Z}_{\neq 0}$. Most devices of interest have isolines with a slope exactly equal to $\pm n_{\text{fp}}$ (Landreman 2022; Rodriguez *et al.* 2022b; Rodríguez *et al.* 2023). The same formulae and quadrature points used for QA can also be used in the case of QH by applying (2.1)–(2.2) in a rotated frame.

To generate multiple devices through the phases described above, we scan over a number of physics and engineering target characteristics. In terms of physics properties, we target various values of mean rotational transform ι , aspect ratio and total coil length, while keeping the major radius of 1 m. The engineering characteristics on the coils impose limits on their curvature and mean squared curvature, which

may not exceed 5 and 5 m^{-2} , respectively. The coils are also designed to have uniform incremental arclength. We refer to these as engineering constraints. The goal of these parameter scans is to explore the space of quasisymmetric stellarators that can be generated by coil sets, and observe possible trade-offs between competing physics and engineering characteristics. Whenever ‘quality of QS’ for a device is provided in the manuscript, it is calculated by taking the square root of the average of (2.1) on a number of surfaces of the magnetic field, which we may refer to as f_{QS} . Note that by construction, the surfaces on which we optimize for QS in the final stage of the workflow are tangent the magnetic field at a fixed number of quadrature points on the surface. As a result, $\mathbf{B} \cdot \mathbf{n}$ error is numerically zero when queried at those same quadrature points. This is notably different from the approach in (Jorge *et al.* 2023), where the $\mathbf{B} \cdot \mathbf{n}$ error is explicitly minimized.

The stellarator design problem amounts to a constrained optimization problem, where the ODE (Phase I) or PDE (Phases II and III) constraint is satisfied to numerical precision at each iteration of the optimization loop. This was not necessarily the case for the physics and engineering constraints. However, all the target physics and engineering constraints were satisfied to a precision of 0.1 % by using a penalty method. After a fixed budget of iterations of the minimization algorithm, we increased by a factor of 10 the penalty weights associated with the constraints that were violated by more than this accuracy threshold. Then, with these new set of weights, the minimization algorithm is restarted. This ensures that these constraints are satisfied to the desired accuracy at the end of the optimization workflow. After a fixed number of restarts, devices that do not satisfy the engineering and physics constraints to 0.1 % accuracy, even after the penalty reweighting, are discarded.

3. A few QH designs in QUASR

In this section, we present a few devices in the data set, summarized in table 1 and plotted in figure 2. The table also includes a link to a device manifest page, where additional physics and engineering information is provided. We estimate the collisionless guiding-centre (de Blank 2004) losses of alpha particles in vacuum magnetic generated the coils of these devices after scaling them up to the ARIES-CS (Najmabadi *et al.* Najmabadi, 2008) minor radius and spawning 5000 α particles isotropically at half-radius in SIMSOPT (Landreman *et al.* 2021) then tracing until 0.2 seconds. All devices have favourable particle loss estimates despite deviations from QS, most likely owing to the narrow orbit widths in QH devices (Landreman & Paul 2022; Paul *et al.* 2022; Buller *et al.* 2024). Except for the two field period device, all presented devices have reasonable elongation of flux surfaces in the RZ plane. Note that this measure of shaping (especially when large) can be somewhat misleading in devices with a large helical excursion of the magnetic axis (see § 5.2 for additional discussion on this). In such cases it is more instructive to measure the elongation in the plane perpendicular to the magnetic axis; for the $n_{\text{fp}} = 2$ device, the maximum elongation in such plane reduces to 6.85, which is still significant. This quantity was numerically computed by fitting a near-axis model to the field (Appendix A), and providing the parameters to pyQSC (Landreman *et al.* 2024).

The first device is a two field period QH field with a ‘crossing point’ with little space for coils, for which coils may nevertheless be found. Although it has the highest particle losses compared with the other devices presented here, the device is reminiscent of the original figure-eight stellarator design (Spitzer 1958).

ID	n_{fp}	Mean ι	Aspect ratio	Max. elongation	Losses (%)	Quality of QS	Device manifest
2229316	2	0.9	6	8.75	3.26	4.03×10^{-4}	[link]
1655332	3	1.2	6	3.66	0.00	4.22×10^{-4}	[link]
1630198	4	1.2	8	3.62	0.00	7.54×10^{-4}	[link]
1641080	5	1.5	8	2.16	0.00	2.26×10^{-3}	[link]
2008429	6	2.5	12	2.31	0.14	8.09×10^{-4}	[link]

TABLE 1. Characteristics of some QH devices from the QUASR data set. The maximum elongation is computed in the RZ plane. The final column of the table contains links to a device manifest page that contain more physics and engineering information about each design.




FIGURE 2. Notable QH devices and their coil sets in the QUASR.

The $n_{fp} = 4$ device is quite similar to the Landreman–Paul precise QH configuration (Landreman & Paul 2022), but was discovered here independently with coils. Coils for the precise QH device have been presented in Wiedman, Buller & Landreman (2024), and found using a stage II algorithm (Zhu *et al.* 2017). Both device ID = 1630198 and the Wiedman *et al.* (2024) device use five coils per half-period and a total coil length of approximately 104 m, when scaled to a major radius of 1 m. All the coil sets in QUASR were designed to have a minimum coil–surface and coil–coil distance of 0.1 m, and maximum coil curvature and mean

squared curvature of 5 m^{-1} and 35 m^{-2} , respectively. These were arbitrary choices, and (Wiedman *et al.* 2024) used different thresholds (some more strict and some less). Both designs though have remarkable estimated alpha particle losses (0.00 %) at reactor scale. Modular coils that wrap poloidally around the magnetic surfaces are just one coil topology and one might conceive of designs for the Landreman–Paul precise QH device with dipole (Kaptanoglu *et al.* 2025) or helical (Kaptanoglu, Langlois & Landreman 2024) coils.

The $n_{\text{fp}} = 6$ device shows an example of a configuration with a high number of field periods, which are configurations that have proved challenging to find and provide a set of coils for (Landreman 2022).

A note must be made on the coils presented for these configurations here. The filamentary coils sometimes appear quite long, which might not make the coil sets particularly practical. However, having such coils remains useful at least as an initial starting point for further refinement.

4. The near-axis QS landscape

The landscape of QS quality in terms of the coil degrees of freedom is complex and high-dimensional, which can be difficult to visualize. Using a near-axis framework, though, this landscape gains significant structure when parametrized in terms of the magnetic axis shape, owing to the key role played by topological properties of these (Rodríguez *et al.* 2022b). A representation of a representative subset of these axes described only by a few parameters (and thus lower dimensionality) can then provide a simple way to assess and visualize QS fields (Rodríguez *et al.* 2023).

To construct such a space, we start by providing the elements necessary to characterize a near-axis quasisymmetric vacuum field (Landreman & Sengupta 2019; Rodríguez *et al.* 2023). Such a field is defined by providing the geometry of the magnetic axis (customarily given as a set of Fourier harmonics \mathbf{R} , \mathbf{Z} describing the axis shape in cylindrical coordinates), and two scalars $\bar{\eta}$ and B_{2c} . Given such inputs, one can evaluate a function $\Delta B_{20}(\mathbf{R}, \mathbf{Z}, \bar{\eta}, B_{2c})$ which measures the magnitude of the symmetry breaking of B at second order in the distance from the magnetic axis (Landreman & Sengupta 2019; Rodríguez *et al.* 2023). That is, if $\Delta B_{20} = 0$, then the near-axis vacuum field considered has ‘perfect’ QS to second order. This scalar function of $(\mathbf{R}, \mathbf{Z}, \bar{\eta}, B_{2c}) \in \mathbb{R}^{2N+1}$ variables is too high-dimensional for straightforward plotting. Using the approach in Rodríguez *et al.* (2023), one may nevertheless reduce the dimensionality in such a way that the critical structure of the problem is retained. This reduction of the space draws from the following logic.

- (i) The main structural element in the field description is the axis shape, which will remain our free parameters. For visualization purposes, the axis geometry is restricted to a few harmonics (this is key only to visualization, but not to the concept itself), in this case $\mathbf{R} = (1, R_{n_{\text{fp}}}, R_{2n_{\text{fp}}}, R_{3n_{\text{fp}}})$, and $\mathbf{Z} = (Z_{n_{\text{fp}}}, Z_{2n_{\text{fp}}}, Z_{3n_{\text{fp}}})$.
- (ii) For a given axis shape, $\bar{\eta} = \bar{\eta}^\dagger(\mathbf{R}, \mathbf{Z})$ is chosen uniquely to roughly minimize elongation of surfaces, formally

$$\bar{\eta}^\dagger = \underset{\bar{\eta}}{\operatorname{argmax}} |\iota(\bar{\eta}) - N|, \quad (4.1)$$

where ι is the on-axis rotational transform and N is the helicity (see Rodríguez *et al.* (2023) for further details and proofs). Changing the value of $\bar{\eta}$ will directly affect elongation of surfaces as well as the rotational transform.

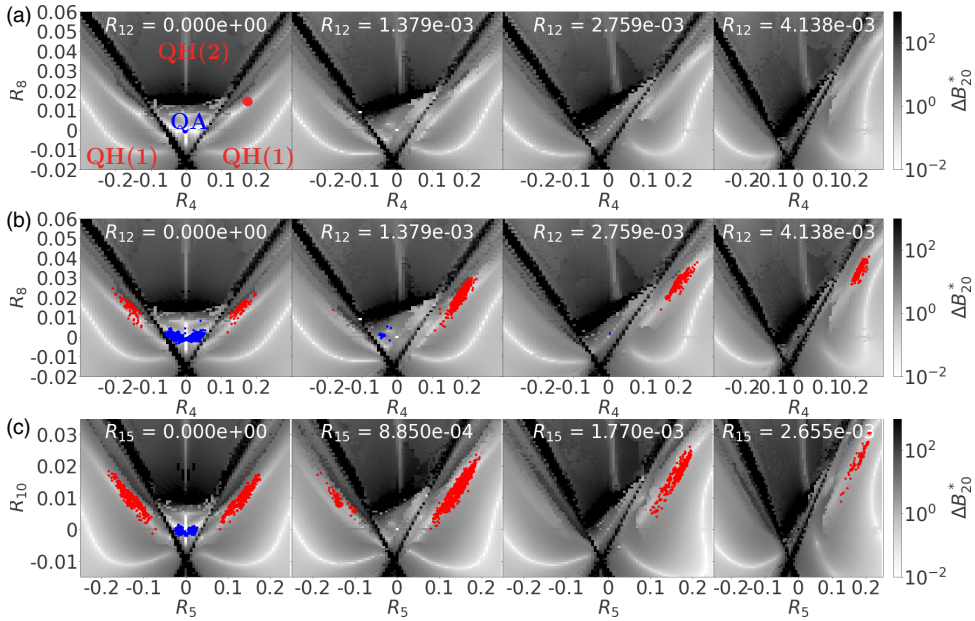


FIGURE 3. Panel (a) shows the landscape of quasisymmetric stellarators for $n_{fp} = 4$, the red dot in (a i) corresponds to the ‘Helically Symmetric eXperiment’ (HSX) device. There are four QS phases delineated by poorly quasisymmetric (dark colour) devices. Devices optimized for QA and QH in QUASR plotted on top of the landscape with blue and red dots, respectively. Panels (b) and (c) correspond to $n_{fp} = 4$ and 5, respectively.

(iii) The \mathbf{Z} harmonics and B_{2c} are then optimized over to minimize ΔB_{20} .

This results in a scalar function

$$\Delta B_{20}^*(\mathbf{R}) = \min_{\mathbf{Z}, B_{2c}} \Delta B_{20}(\mathbf{R}, \mathbf{Z}, \bar{\eta}^\dagger(\mathbf{R}, \mathbf{Z}), B_{2c}) \quad (4.2)$$

which is only a function of three independent variables ($R_{n_{fp}}, R_{2n_{fp}}, R_{3n_{fp}}$), and it can easily be visualized by evaluating it on a three-dimensional tensor product grid using the `pyQSC` code (Landreman *et al.* 2024). The minimization problem over the \mathbf{Z} harmonics (\mathbf{Z}) uses local optimization algorithms, where the \mathbf{Z} harmonics provided to the optimization are initialized to have $\mathbf{Z} := \mathbf{R}$. This results in magnetic axes with \mathbf{Z} harmonics that scale with R as the optimization algorithm does not stray far from the initial guess. The resulting plots have structure thanks to the topological properties of the magnetic axis (Rodríguez *et al.* 2022b). This landscape can be divided into phases (or regions) where either QA and QH devices are found, figure 3, corresponding to axis shapes with different self-linking numbers (Rodríguez *et al.* 2022b; Moffatt & Ricca 1992; Fuller Jr 1999; Oberti & Ricca 2016). At the origin, we have the QA phase, surrounded by QH phases. As the magnitude of the third harmonic increases (and thus the axis shaping increases), the QA region shrinks; there is a limit on how much an axis can be shaped before it gains a non-zero helicity and represent a QH device. The east and west QH regions have helicity n_{fp} , and the north and south regions have helicity $2n_{fp}$.

In Rodríguez *et al.* (2023), two branches of QH stellarators with favourable QS were observed: the HSX branch, as well as a new QH branch that had not yet been

explored. The HSX branch is the one on which the HSX device lies, and the ‘new QH’ branch is the one below it (see [figure 3a](#)). To visualize the devices in QUASR on this landscape, we fit the near-axis model to the magnetic fields generated by the coil sets. The algorithm to do this is described in [Appendix A](#).

4.1. The QUASR on the near-axis QS landscape

Following the procedure in [Appendix A](#), QUASR devices may be then approximated by near-axis fields (a form of reduced representation of the configuration) and represented in the reduced quasisymmetric phase-space discussed above. However, we must note that some of the simplifying assumptions behind such a reduced phase space plot in [figure 3](#) do not necessarily hold for all the devices in QUASR.

First, note that the optimization that was used to design the devices in QUASR did not directly target the devices’ elongation. As a result, we are not guaranteed that $\bar{\eta}^*$ will be close to $\bar{\eta}^\dagger$ as the phase space construction assumes. Around 40 % of $n_{fp} = 4, 5$ devices have $|\bar{\eta}^* - \bar{\eta}^\dagger|/|\bar{\eta}^\dagger| < 0.1$ (roughly a change of ~ 20 % in elongation) and this increases to around 70 % when the upper bound is relaxed to 0.2. This illustrates that there is a non-negligible portion of devices in QUASR with possibly unrealistic elongations, which motivates the addition of constraints on surface elongation to the design workflow in [figure 1](#).

Second, for illustration purposes, we focused on configurations whose axes had three dominant harmonics. A device is deemed to have N dominant Fourier harmonics if $R_{k n_{fp}}(1 + k^2 n_{fp}^2), Z_{k n_{fp}}(1 + k^2 n_{fp}^2) < 0.1$ for $k \geq N + 1$. These conditions rely on rescaled Fourier harmonics that are dimensionless and naturally measure relevance of axis harmonics, introduced in Rodríguez *et al.* (2022b). Around 50 % of devices have three dominant harmonics, and this number increases to around 70 % when four dominant harmonics are allowed. During the globalization of the near-axis phase of the workflow, the axis is described only using two Fourier harmonics. Then, the solution is subsequently polished and the number of harmonics used to describe the axis is increased. This continuation procedure may have influenced these statistics. Additionally, note that the choice to represent magnetic axes in cylindrical coordinates and Fourier harmonics might not be the optimal choice, and an alternative parametrization of the phase space could be better suited (even within the near-axis framework).

With the above in mind, we present on the phase space in [figure 3](#) those devices in QUASR that satisfy

- (i) the device’s $\bar{\eta}^*$ has a relative error of at most 10 % with respect to $\bar{\eta}^\dagger$, i.e. $|\bar{\eta}^* - \bar{\eta}^\dagger|/|\bar{\eta}^\dagger| < 0.1$;
- (ii) the magnetic axis has only three dominant harmonics, i.e. $R_{k n_{fp}}(1 + k^2 n_{fp}^2), Z_{k n_{fp}}(1 + k^2 n_{fp}^2) < 0.1$ for $k \geq 4$.

It is also ensured that all plotted devices have $R_{3n_{fp}} > 0$. If this is not the case, then the device is rotated by a half-period, which flips the sign of the odd harmonics: $R_{n_{fp}}, R_{3n_{fp}}$. A device’s coordinate on the landscape is computed from the harmonics of its magnetic axis.

We observe that the QA devices lie as expected in the QA region. This QA region shrinks as the magnitude of the third harmonic increases, which corroborates the observation that QA devices typically do not have highly shaped magnetic axes. In the QH region, QUASR configurations appear to cluster around the HSX branch,

noticeably leaving the ‘new QH’ branch, which at the time of publication the authors in Rodríguez *et al.* (2023) noted that had not yet been explored. This is so even when optimization efforts were launched directly from the new branch, cases in which optimization trajectories left the branch to find better QS elsewhere. Whether the struggle to find such configurations comes from the difficulty to find appropriate coils under the current engineering constraints, some intrinsic limitations on volumetric QS or the existence of a higher-dimensional path in the phase space towards better QS, remains a subject of study.

While this comparison with the near-axis model is interesting and insightful, this approach has some important drawbacks. The landscape plots we obtain here are restricted to two or three dimensions before the visualizations become laborious. Despite this, it might still be possible to measure a device’s distance to a favourable branch in higher dimensions, permitting some additional higher-dimensional insights. A better representation of the axes could also be fitting. In addition, the filters on $|\bar{\eta}^* - \bar{\eta}^\dagger|/|\bar{\eta}^\dagger|$ may also be unnecessarily restrictive, although as we have argued, this filter is not completely artificial if elongation of flux surfaces is taken into consideration. In the following section, we will use a different approach that applies to the entire data set.

5. Principal component analysis

Principal component analysis is a dimensionality reduction technique that projects D -dimensional data points $\mathbf{x}_i \in \mathbb{R}^D$ onto a lower-dimensional manifold (Bishop 2006). The vector \mathbf{x}_i is also called a feature vector. For the following explanation, we focus on reducing the dimensionality of the stellarators to two-dimensional points on the plane, though one is free to retain as many dimensions as desired.

The PCA finds a plane that lies closest to the data points, and then projects the data points on this plane, thereby reducing its dimensionality. Any projected point can be expressed as $\alpha_1 \mathbf{d}_1 + \alpha_2 \mathbf{d}_2 + \boldsymbol{\gamma}$, for some $\alpha_1, \alpha_2 \in \mathbb{R}$, where \mathbf{d}_1 and \mathbf{d}_2 are orthonormal. The vectors \mathbf{d}_1 , \mathbf{d}_2 and $\boldsymbol{\gamma}$ are chosen to minimize the error,

$$\boldsymbol{\gamma}^*, \mathbf{d}_1^*, \mathbf{d}_2^* = \underset{\substack{\boldsymbol{\gamma}, \mathbf{d}_1, \mathbf{d}_2 \\ \text{s.t. } \|\mathbf{d}_1\|, \|\mathbf{d}_2\|=1}}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \tilde{\mathbf{x}}_i\|^2, \quad (5.1)$$

where $\tilde{\mathbf{x}}_i$ is the projection of the original point \mathbf{x}_i onto the plane of the form

$$\tilde{\mathbf{x}}_i = \boldsymbol{\gamma} + \alpha_{1,i} \mathbf{d}_1 + \alpha_{2,i} \mathbf{d}_2. \quad (5.2)$$

From the orthonormality of \mathbf{d}_1 and \mathbf{d}_2 , it can be shown that $\alpha_{1,i} = \mathbf{d}_1 \cdot (\mathbf{x}_i - \boldsymbol{\gamma})$ and $\alpha_{2,i} = \mathbf{d}_2 \cdot (\mathbf{x}_i - \boldsymbol{\gamma})$, and the optimal choice of $\boldsymbol{\gamma}$ is the mean of the data, i.e. $\boldsymbol{\gamma}^* = \bar{\mathbf{x}} = 1/N \sum_{i=1}^N \mathbf{x}_i$. The Karush–Kuhn–Tucker conditions of the minimization problem (Bishop 2006) in (5.1) reveal that the optimal choice of vectors $\mathbf{d}_1^*, \mathbf{d}_2^*$ are given by the first two eigenvectors of the covariance matrix

$$\frac{1}{N} \sum_{i=1}^N (\mathbf{x}_i - \bar{\mathbf{x}})^T (\mathbf{x}_i - \bar{\mathbf{x}}). \quad (5.3)$$

The vectors $\mathbf{d}_1^*, \mathbf{d}_2^*$ are called, respectively, the first and second principal components. For this choice of $\boldsymbol{\gamma}^*, \mathbf{d}_1^*, \mathbf{d}_2^*$, the reconstruction error is

$$\frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \tilde{\mathbf{x}}_i\|^2 = \lambda_3 + \cdots + \lambda_D, \quad (5.4)$$

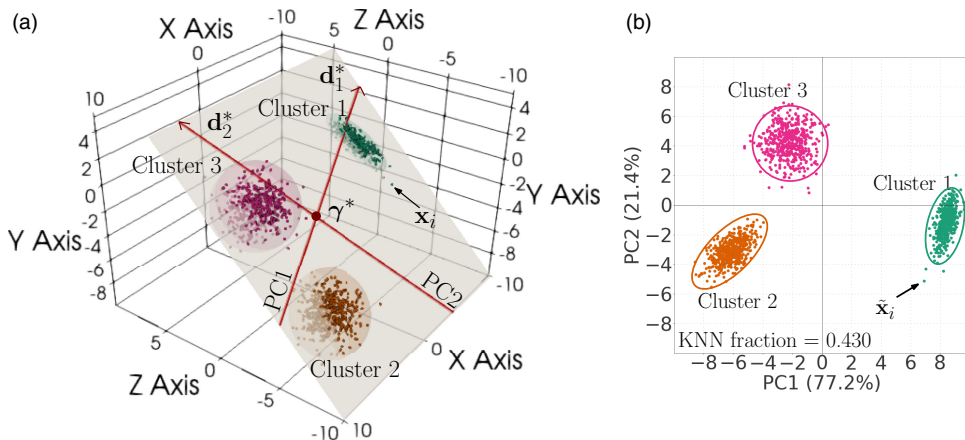


FIGURE 4. Data $\mathbf{x}_i \in \mathbb{R}^3$ from three Gaussians are shown in (a), and the optimal projection $\tilde{\mathbf{x}}_i$ onto a two-dimensional plane is shown in (b). The transparent ellipsoids correspond to the 95 % confidence interval of each Gaussian.

which is the sum of the $D - 2$ remaining eigenvalues of the covariance matrix. The ratios $\lambda_1/(\lambda_1 + \dots + \lambda_D)$ and $\lambda_2/(\lambda_1 + \dots + \lambda_D)$ measure by how much the first two principal components reduced the reconstruction error. If the sum of these two quantities is close to unity, then little information was lost in the dimensionality reduction. One is free to include as many principal components as desired, but for simplicity of plotting and ease of understanding, we restrict ourselves to two or three principal components.

Once a PCA is completed, each high-dimensional data point \mathbf{x}_i will have a corresponding two-dimensional coordinate $(\mathbf{d}_1^* \cdot (\mathbf{x}_i - \boldsymbol{\gamma}^*), \mathbf{d}_2^* \cdot (\mathbf{x}_i - \boldsymbol{\gamma}^*)) \in \mathbb{R}^2$. We illustrate this procedure with a three-dimensional set of data points $\mathbf{x}_i \in \mathbb{R}^3$ in figure 4(a), projected down to a two-dimensional plane on the right. Projected points far away from each other on the plane will also be far away from each other in the higher dimensional space. In other words, dissimilar devices in the low-dimensional representation are dissimilar in higher dimensions. Note that the converse of this statement is not necessarily true.

The advantages of PCA is that it is an unsupervised technique that does not require labels or prior knowledge about the data. It identifies important features in a low-dimensional subspace that describe the data set and permit its visualization. By doing so, it reduces the complexity and noise in the data and gives an idea of its intrinsic dimensionality, which is revealed by examining the eigenvalues of the covariance matrix. The PCA is often used as a preprocessing step for machine learning and optimization algorithms to expose the most important components of the data (Kobak & Berens 2019). It also does not have any parameters to adjust, which makes it straightforward to apply to a data set.

However, PCA can also have disadvantages. By reducing the dimensionality of the data, we may lose some important information. The transformation to low-dimensional space may distort the geometric relationships that existed in the original space. In addition, PCA is useful when the data are well approximated on a linear manifold of dimension less than D , which is certainly not always the case. For example, consider a data set that lies on a three-dimensional helix. The intrinsic dimension of the data are one, since the helix can be unravelled to parameterize any point on

it by a single coordinate. But applying PCA to the data set distorts it, since its topology is ill-approximated on a linear manifold of dimension less than three. In cases like this, nonlinear dimensionality reduction techniques can be useful, and the literature on these approaches is rich. For example, isomap (Tenenbaum, Silva & Langford 2000), and t-distributed stochastic neighbour embedding (Van der Maaten & Hinton, 2008) are all well-known algorithms to accomplish this task. The disadvantage of nonlinear techniques is that they may have user-defined tuning parameters, which can be difficult to select. Since the underlying topology of our data set is not known, and to avoid the difficulties associated with nonlinear techniques, we use PCA in this work. Auto encoders have also been proposed as interesting methods for compression or dimensionality reduction. Bourlard & Kamp (1988) showed that auto encoders perform a similar operation to singular value decomposition or low rank matrix approximation. Since PCA is a type of low rank matrix approximation, this is an interesting connection to what is done in this work.

In light of the above discussion, we use a global and local measure of accuracy to evaluate how faithful the data in lower dimensions is to the original.

- (i) *The cumulative principal component (PC) ratios.* As mentioned above, this score quantifies how much information was lost in the projection onto the lower-dimensional linear manifold, corresponding to the sum of the PC ratios: $\lambda_1/(\lambda_1 + \dots + \lambda_D)$ when projecting onto a one-dimensional manifold, $(\lambda_1 + \lambda_2)/(\lambda_1 + \dots + \lambda_D)$ when projecting onto two-dimensional manifold, and $(\lambda_1 + \lambda_2 + \lambda_3)/(\lambda_1 + \dots + \lambda_D)$ when projecting onto a three-dimensional manifold. The closer this value is to unity, the less information is lost in the projection. If more PCs are included, then the corresponding eigenvalues are added to the numerator. This is a global measure of accuracy.
- (ii) *The k-nearest neighbours (KNN) fraction.* The fraction of KNN in the high-dimensional data set that remain in the two-dimensional projected data points (Lee & Verleysen 2009; Kobak & Berens 2019). The purpose of this quantity is to measure how accurately local relationships are preserved after dimensionality reduction. Note that for very large k , this fraction approaches unity, so some reasonable choice of k has to be picked. This is a local measure of accuracy.

We present these at work in a simple example of Gaussian processes in [Appendix B](#). In what follows, we define how we represent the different devices in the QUASR dataset, and then analyse them using the tools above.

5.1. Feature vectors

In order to apply PCA to our set of stellarators, we must define how a device's feature vector is constructed. A device's feature vector is defined to have $D = 663$ entries, where each entry is a Fourier harmonic of a stellarator symmetric surface in the device's vacuum magnetic field, using $m_{\text{pol}}, n_{\text{tor}} = 10$, where $m_{\text{pol}}, n_{\text{tor}}$ are, respectively, the number of poloidal and toroidal harmonics used to represent the surface geometry. We use the special surface representation (Giuliani *et al.* 2022b):

$$x(\varphi, \theta) = \cos(\varphi)\hat{x}(\varphi, \theta) - \sin(\varphi)\hat{y}(\varphi, \theta), \quad (5.5)$$

$$y(\varphi, \theta) = \sin(\varphi)\hat{x}(\varphi, \theta) + \cos(\varphi)\hat{y}(\varphi, \theta), \quad (5.6)$$

$$z(\varphi, \theta) = \sum_i \sum_j z_{i,j} u_i(\varphi) v_j(\theta), \quad (5.7)$$

where

$$\hat{x}(\varphi, \theta) = \sum_i \sum_j \hat{x}_{i,j} u_i(\varphi) v_j(\theta), \quad (5.8)$$

$$\hat{y}(\varphi, \theta) = \sum_i \sum_j \hat{y}_{i,j} u_i(\varphi) v_j(\theta) \quad (5.9)$$

and

$$\begin{aligned} u(\varphi) &= (1, \cos(2\pi n_{\text{fp}}\varphi), \sin(2\pi n_{\text{fp}}\varphi), \dots, \cos(2\pi n_{\text{fp}}n_{\text{tor}}\varphi), \sin(2\pi n_{\text{fp}}n_{\text{tor}}\varphi)), \\ v(\theta) &= (1, \cos(2\pi\theta), \sin(2\pi\theta), \dots, \cos(2\pi m_{\text{pol}}\theta), \sin(2\pi m_{\text{pol}}\theta)). \end{aligned} \quad (5.10)$$

Since we are only working with stellarator symmetric devices, the symmetry-breaking basis functions in the sums of (5.5)–(5.9) are skipped. This representation uses three truncated Fourier series to represent the Cartesian coordinates of the surface: $(x(\varphi, \theta), y(\varphi, \theta), z(\varphi, \theta))$. This is in contrast to ‘Variational Moments Equilibrium Code’ (VMEC) (Hirshman & Whitson 1983a) and DESC (Panici *et al.* 2023), which use a cylindrical representation $(R(\phi, \theta), \phi, Z(\phi, \theta))$, with ϕ the standard cylindrical toroidal angle.

The feature vector associated with a device contains all the Fourier harmonics $\hat{x}_{i,j}$, $\hat{y}_{i,j}$ and $\hat{z}_{i,j}$ associated with the surface, without any additional normalization. We emphasize this lack of normalization because in the context of PCA it is usually recommended to normalize the data across samples. One common normalization is to rescale each component of the feature vector to a Gaussian with mean zero and covariance one to ensure no feature dominates the analysis due to it being much larger. However, this kind of normalization only makes sense if the different components of the feature vector are not meaningfully related to each other and their values have different orders of magnitude because they come from inherently different quantities or phenomena. In our case, the feature vectors contain Fourier coefficients, which decay as the harmonic increases. Therefore, we chose to forgo the normalization of different Fourier components across the devices because the values of the Fourier coefficients have a very important meaning: they are related to each other to describe a single surface. For the vacuum field devices in QUASR, the magnetic field of the device is fully specified by the geometry of a single magnetic surface, regardless of how this is parametrized. Additional parameters would be required in the feature vector for devices with plasma pressure, and normalization might be indicated there. This leaves some non-uniqueness in our chosen representation, as different angle choices for the defining surface can lead to different Fourier coefficients defining the same magnetic surface geometry. To remove this indeterminacy, we choose the angles on the surface (φ, θ) to correspond to Boozer coordinates. Even then, there remains certain non-uniqueness; for example, the same configuration rotated by a half-period about the Z -axis, or reflected about the XY plane would correspond to different feature vectors. Thus, to ensure that geometrically similar devices are close to each other in Fourier space, in addition to

- (i) making the toroidal and poloidal angles on the surface Boozer angles (φ, θ) ,

we check that the magnetic axis in cylindrical coordinates $\Gamma(\phi) = (R(\phi), \phi, Z(\phi))$ satisfies

$$(ii) \quad Z'(0) \geq 0,$$

$$(iii) \quad R(0) \geq R(\pi/n_{\text{fp}})$$

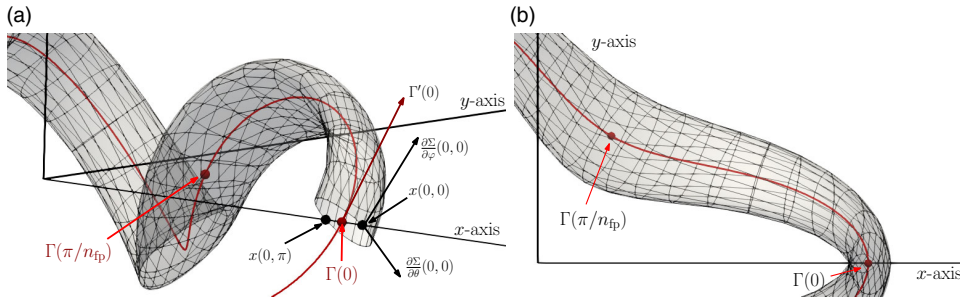


FIGURE 5. Conditions that the surface parametrization must satisfy so that the feature vector of Fourier coefficients is unique. Panel (a) illustrates how the surface coordinates satisfy $x(0, 0) \geq x(0, \pi)$, $((\partial z)/(\partial \theta))(0, 0) \geq 0$, $((\partial y)/(\partial \varphi))(0, 0) \geq 0$, and that the Z -coordinate of the magnetic axis (in red) satisfies $Z'(\phi) \geq 0$. Panel (b) is a view in the $-Z$ direction onto the XY plane, showing that the radius of the magnetic axis satisfies $R(0) \geq R(\pi/n_{fp})$. The grid lines on the magnetic surface correspond to lines of constant toroidal and poloidal Boozer angles.

and that the surface $\Sigma(\varphi, \theta) = (x(\varphi, \theta), y(\varphi, \theta), z(\varphi, \theta))$ in Boozer coordinates wraps toroidally and poloidally such that

- (iv) $x(0, 0) \geq x(0, \pi)$,
- (v) $((\partial z)/(\partial \theta))(0, 0) \leq 0$,
- (vi) $((\partial y)/(\partial \varphi))(0, 0) \geq 0$.

These conditions are illustrated in figure 5. The second and third conditions break the freedom about the stellarator symmetric points. The final three conditions ensure that the origin of the surface ($\varphi = 0, \theta = 0$) lies on the outboard side of the stellarator, that the parametrization wraps poloidally clockwise if looking in the $+Y$ direction at the curve given by $\Sigma(\theta, 0)$, and that increasing φ means going anti-clockwise about the Z -axis if looking in the $-Z$ direction. If these conditions are not satisfied, we reflect the device about the XY plane, rotate it by a half-period about the Z -axis, or reparametrize the surface so that it wraps in the correct direction. The direction of the inequalities in conditions (ii)–(vi) were arbitrarily chosen.

A more compact feature vector could have been generated by converting the representation in (5.5)–(5.9) to a cylindrical one, thereby reducing the number of Fourier series from three to two. We do not do so here because the coil design workflow relies on the Cartesian representation of the surfaces. Although Cartesian coordinates are used in the coil design workflow, this does not mean any field shape is possible. Note that in Phase I (the near-axis phase of the coil design) it is assumed that the magnetic axis can be represented in cylindrical coordinates. As a result, we will not find devices with more general magnetic axis geometries, e.g. devices where the axis is anywhere vertical, or knotted. Generalizing our implementation is the subject of future work.

In addition, this conversion is not lossless when restricted to a finite dimensional Fourier series, as explained in § 5.2. Other forms of parametrization could also be possible. It is tempting, for example, to represent each configuration by its second-order near-axis model, which constitutes a significantly reduced space. Although this could be interesting, it does not correspond to a one-to-one map, and thus should not be used as a unique representation. This is because, *a priori*, one could devise an

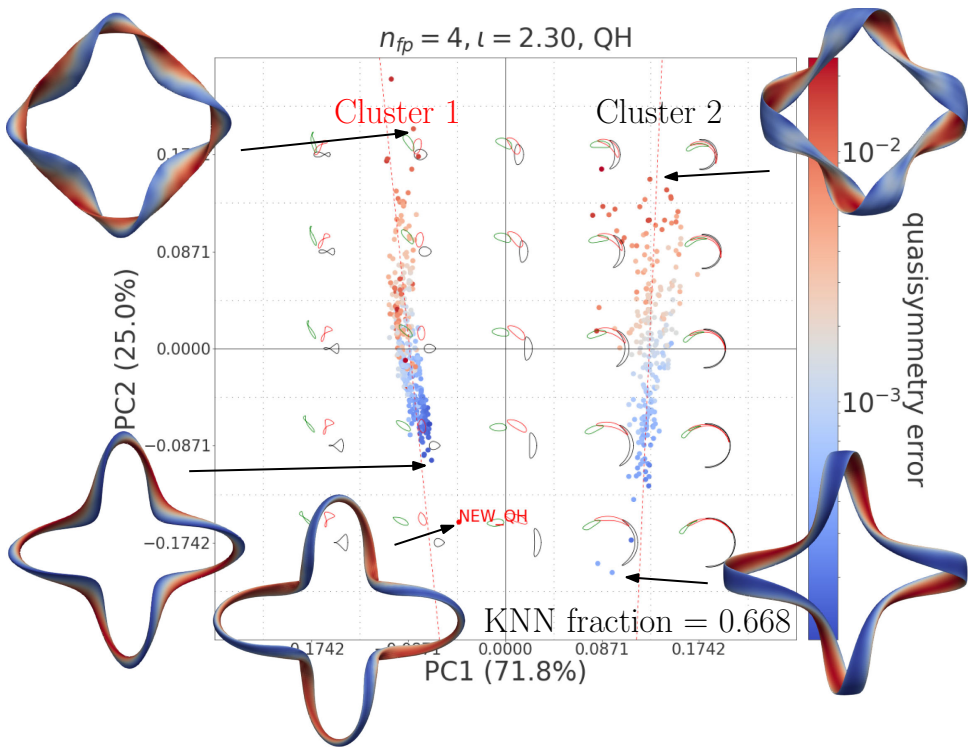


FIGURE 6. The PCA of QH symmetric devices in QUASR, along with the new QH device indicated with red text. At the centre of each dashed square, we take the associated surface DOFs and generate cross-sections to illustrate the geometric diversity of devices. This type of cross-section diagram is common in morphometrics literature (Bonhomme *et al.* 2014). The red dashed lines correspond to projections of one-dimensional PCAs of the left- and right-hand clusters onto the two-dimensional manifold. The colour on the magnetic surfaces corresponds to the local field strength. The quality of QS is calculated by taking the square root of the average of (2.1) on a number of surfaces of the magnetic field.

infinity of distinct fields, all of which correspond to the same second-order near-axis construction. Another alternative could be to use a vector of Fourier harmonics associated with the coils and their currents. We did not do this due to the inherent non-uniqueness of coil sets, i.e. multiple, vastly different coil sets can produce similar magnetic fields. That being said, we have not fully explored these directions, and it might be interesting to do this in the future.

When presenting QUASR devices in what follows, we will do so alongside quasisymmetric devices available in the literature (see details in Appendix C).

5.2. Quasihelical symmetry ($n_{fp} = 4$, $\iota = 2.3$)

To start, we look at $n_{fp} = 4$ devices with a mean rotational transform of 2.3 in figure 6. We observe that there are two distinct clusters of devices. Within these clusters, the devices form a continuum with varying qualities of QS.

One may question how faithful this two-dimensional plot is to the true high-dimensional data set. A hint that this plot is an accurate representation of the relationships in the data are the sum of the percentages in the x and y axis labels is large. This is evidence that including a third dimension is likely unnecessary to




FIGURE 7. Cross-sections of a device from cluster 2 in figure 6. The section is computed in the RZ plane (red) and the plane orthogonal the magnetic axis (blue). The magnetic axis is drawn in black. Both planes pass through the same point on the magnetic axis.

understand how devices relate to each other, which can be confirmed via a three-dimensional PCA. Using 10 nearest neighbours in the KNN calculation, the average KNN fraction over all data points is 0.689. This indicates that on average, just under 70 % of each node's 10 nearest neighbours in the high-dimensional representation are preserved on the lower-dimensional manifold. These projection error measures are comparable to the dimensionality reduction example of data sampled from simple Gaussians, in the introduction of Appendix B.

Each point on the landscape plot corresponds to a specific vacuum magnetic field, with associated magnetic surfaces. There are box-shaped regions delineated by dashed lines on the plot. Taking the point at the centre of each box, we compute cross-sections of the associated magnetic surface. The cross-sections are plotted in black, red and green, associated with the cylindrical angle $\phi = 0, \pi/(4n_{fp}), \pi/(2n_{fp})$, respectively. This figure makes it clear how devices differ from each other on the plane, and importantly how moving along the principal components affects the geometry of the device. Moving in the positive direction associated with \mathbf{d}_1^* , the elongation appears to increase, while going in the negative direction the elongation decreases. In the positive \mathbf{d}_2^* direction, the relative positions of the cross-sections move, and in particular the axis excursion increases. These cross-sections do not necessarily correspond to physically relevant devices, as illustrated by the leftmost column, where the cross-sections are self-intersecting. This is not surprising, as the devices lie on a nonlinear manifold, and the PCA slices through it with a linear one. If one strays too far away from the origin of the plot (the sample mean), the approximation errors become noticeable and unphysical devices appear.

Within the space populated by the QUASR devices, we find that converting the highly shaped surfaces, e.g. the device in figure 7, to a cylindrical representation, as is used in the more standard equilibrium solvers VMEC or DESC, requires more than $m_{pol}, n_{tor} = 20$ Fourier modes. This is significantly larger than the number of modes used to design the stellarator ($m_{pol}, n_{tor} = 10$) with the representation in (5.5)–(5.9). Figure 7 illustrates the geometrical origin of this difficulty, where due to the large helical excursion of the magnetic axis, cross-sections in the RZ plane become highly elongated, leading to a misleading interpretation of the elongation of the device. This occurrence is particularly true for the devices on the right-hand branch, in the bottom quadrant, where standard equilibrium solvers such as VMEC experience difficulties. These problems are probably shared by DESC (Panici *et al.* 2023) which uses the

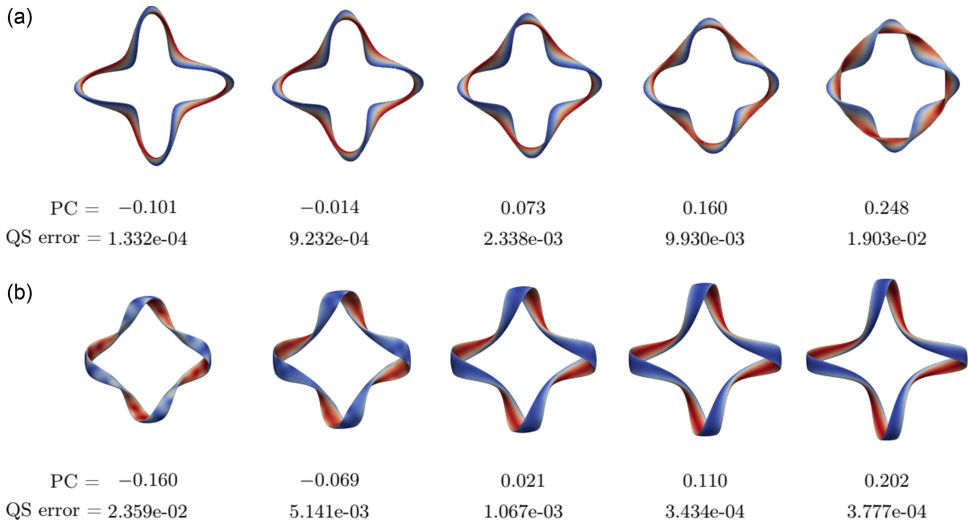


FIGURE 8. Continuum of devices in QUASR. Panels (a) and (b) correspond to the left- and right-hand cluster in figure 6, respectively. The numbers below these devices are their corresponding principal component values, and QS errors. The colour on the magnetic surfaces corresponds to the local field strength.

same cylindrical representation, but could perhaps be alleviated in the Frenet version of the GVEC code (Hindenlang, Plunk & Maj 2024).

As a reference for this set, we also provide the location of the NEW_QH device from § 4 and observe again that there are no devices in this subset of QUASR that are close to it. The NEW_QH device appears closest to the low elongation branch, in the vicinity of devices with excellent QS, but is not contained in QUASR.

5.2.1. Visualizing the continuum of devices in each cluster

We observe that the two clusters appear to lie on separate one-dimensional linear manifolds. Subselecting all the devices in the first cluster where $PC1 < 0$, we do a second PCA on the full feature vector with $D = 663$ entries, and retaining only a single principal component. It appears that the devices are well approximated by a one-dimensional linear manifold, because $\lambda_1/(\lambda_1 + \dots + \lambda_{663}) = 0.916$. The same thing can be done for devices in the second cluster, i.e. where $PC1 > 0$. Again, the devices here are well approximated by a one-dimensional linear manifold because $\lambda_1/(\lambda_1 + \dots + \lambda_{663}) = 0.876$. After these separate one-dimensional PCAs, each device from the first cluster now has a single coordinate called PC, which lies on the interval $[\min(PC), \max(PC)]$. We can easily visualize the progression of the device geometries along these one-dimensional manifold by sampling devices uniformly according to the principal component label, i.e. devices with a principal component label that is approximately $\min(PC1) + i(\max(PC) - \min(PC))/N$ for $i = 0, \dots, N$. In figure 6, we have projected these one-dimensional manifolds onto the plane, indicated by two red dashed lines. The progression of devices along these lines in both clusters is shown in figure 8, where the transition between the endpoint geometries appears smooth. This is evidence that the devices in each cluster in the full $D = 663$ dimensional feature space lie on simple one-dimensional continua.

The difference between the two clusters can in this way also be made clearer. We said before that d_1^* appeared to control the elongation of cross-sections, with cluster 2 characterized by an increased value for it. However, from [figure 8](#), there is a key geometric difference between the clusters that explains this observation and which has to do with the difference in twisting of the devices. At the sharpest corners in the RZ projection, cluster 1 devices have a magnetic axis mainly lying on the plane, while cluster 2 devices have their axes coming off the plane. This difference in bending and twisting is key in controlling rotational transform and other properties, and is worth further exploration in the future.

5.3. Quasihelical symmetry ($n_{fp} = 4, \iota = 1.1$)

We look at another region of the QH design space that has been explored in the literature ($n_{fp} = 4$ and $\iota = 1.1$) in [figure 9](#), where there are two clusters of devices once again, in this case with the largest centred at the origin, the second lying above it. The HSX, WISTELL-A and NIES_AR = 9 devices are given within the main cluster for reference. There is in this case a large disparity between clusters, with the top one being scarcely populated.

To try to make sense of this particular clustering of configurations, we consider constructing and characterizing vacuum magnetic fields for the whole PC plane. Each pair of coordinates in such a plane corresponds to one particular stellarator geometry, whose equilibrium we may then solve using VMEC. This allows us to understand the appearance of the cluster, as well as opening the door to analysing other configuration properties within such a space. We present a measure of QS quality over this plane in [figure 9\(b\)](#) (in particular a plot of the two-term QS residual from (Helander & Simakov 2008; Landreman & Paul 2022), also f_c in Rodriguez *et al.* 2022a), and overlay as black dots the QUASR devices in the clusters above with two dominant PC components. As expected, two agglomerations of devices in QUASR lie in local QS minima, providing external validation that the devices in QUASR were properly optimized for QS. However, it fails to properly describe the separation into the two clusters, for which one must consider the additional design constraint that the average rotational transform is $\iota = 1.1$. [Figure 9\(c\)](#) shows the relative violation of this constraint of rotational transform on the plane; i.e. $|\iota - 1.1|/1.1$. It appears that there are two separate manifolds where the ι design constraint is accurately satisfied. Since the devices in QUASR were designed to satisfy these constraints to within 0.1 % accuracy, the configuration clusters lie in the intersection of these delineated regions and [figure 9\(b\)](#). This is not exact, though, especially because in reality devices have more than only two principal components, and are not located precisely on the PC1–PC2 plane.

The topology of the data appears thus to be governed by the interplay of the rotational transform design constraint and the valleys of the QS objective. This perspective on the origin of the clusters also explains why the top cluster is more scarcely populated, as it has less accurate QS than the bottom one and hence corresponds to a shallower optimization well. This phenomenon could also explain the clusters in § 5.2, but VMEC failing in a large portion of the PC1–PC2 plane prevents any conclusion.

5.4. Quasihelical symmetry ($n_{fp} = 3, \iota = 1.1$)

In [figure 10\(a\)](#), we plot the landscape of devices for a lower rotational transform $\iota = 1.1$ and $n_{fp} = 3$, where it appears that little information was lost in the projection. We find the devices in QUASR here are close to an unpublished device from

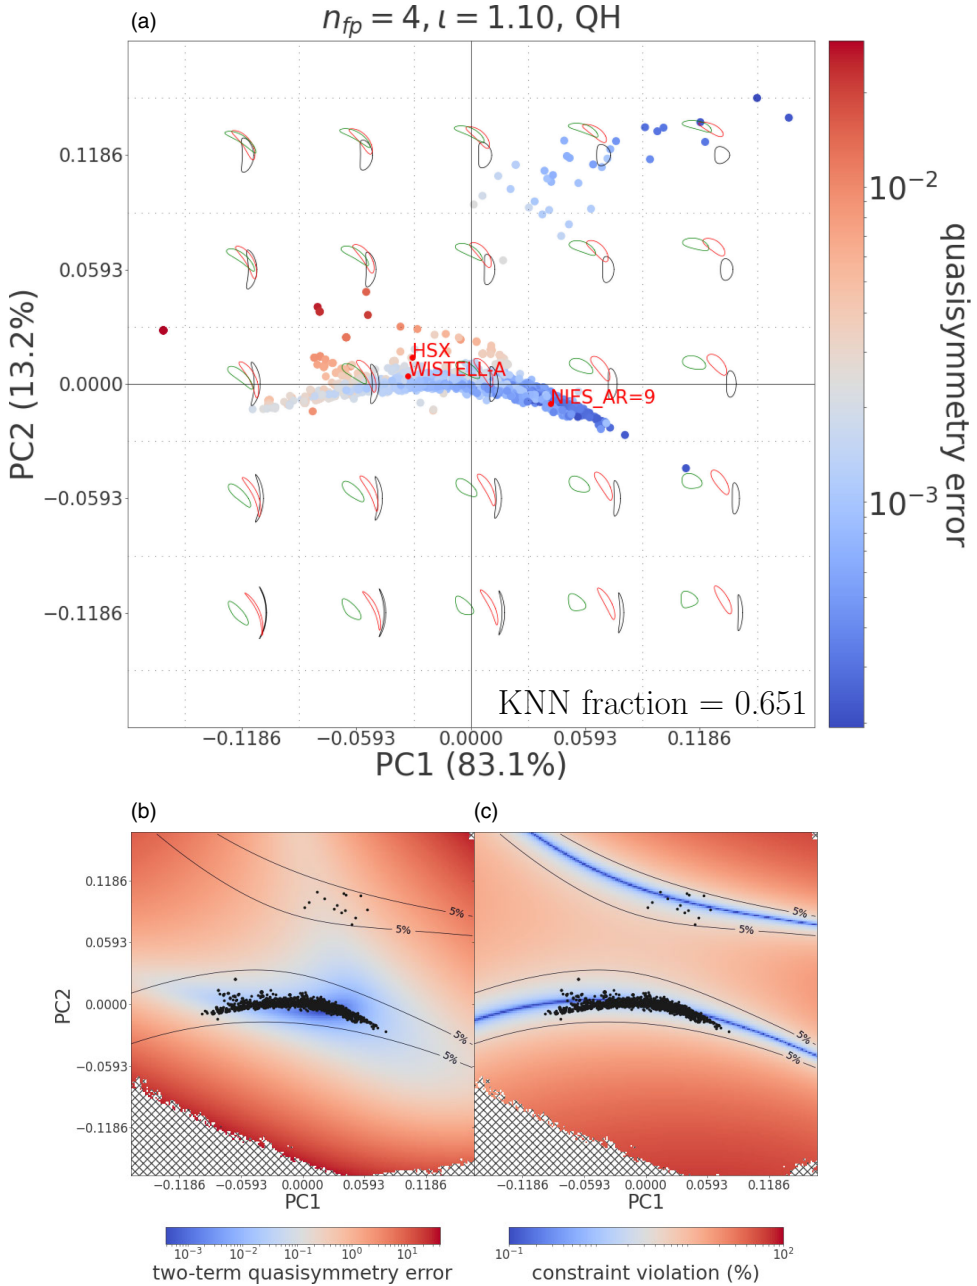


FIGURE 9. Panel (a) contains a PCA of QH symmetric devices in QUASR when $\iota = 1.1$ and $n_{fp} = 4$, where devices in QUASR with two dominant PC components are plotted. Panels (b) and (c) show, respectively, the two-term QS error from Helander & Simakov (2008) and Landreman & Paul (2022) computed using VMEC and relative ι constraint violation ($|\iota - 1.1|/1.1$) on the PC subspace. The black isolines correspond to coordinates where the constraint violation is precisely 5% and the black dots are devices in panel (a) that have only two dominant principal components. The region on the PC plane where the two-term QS error could not be evaluated or VMEC did not converge, notably on the bottom left-hand side of the panels, is hatched.

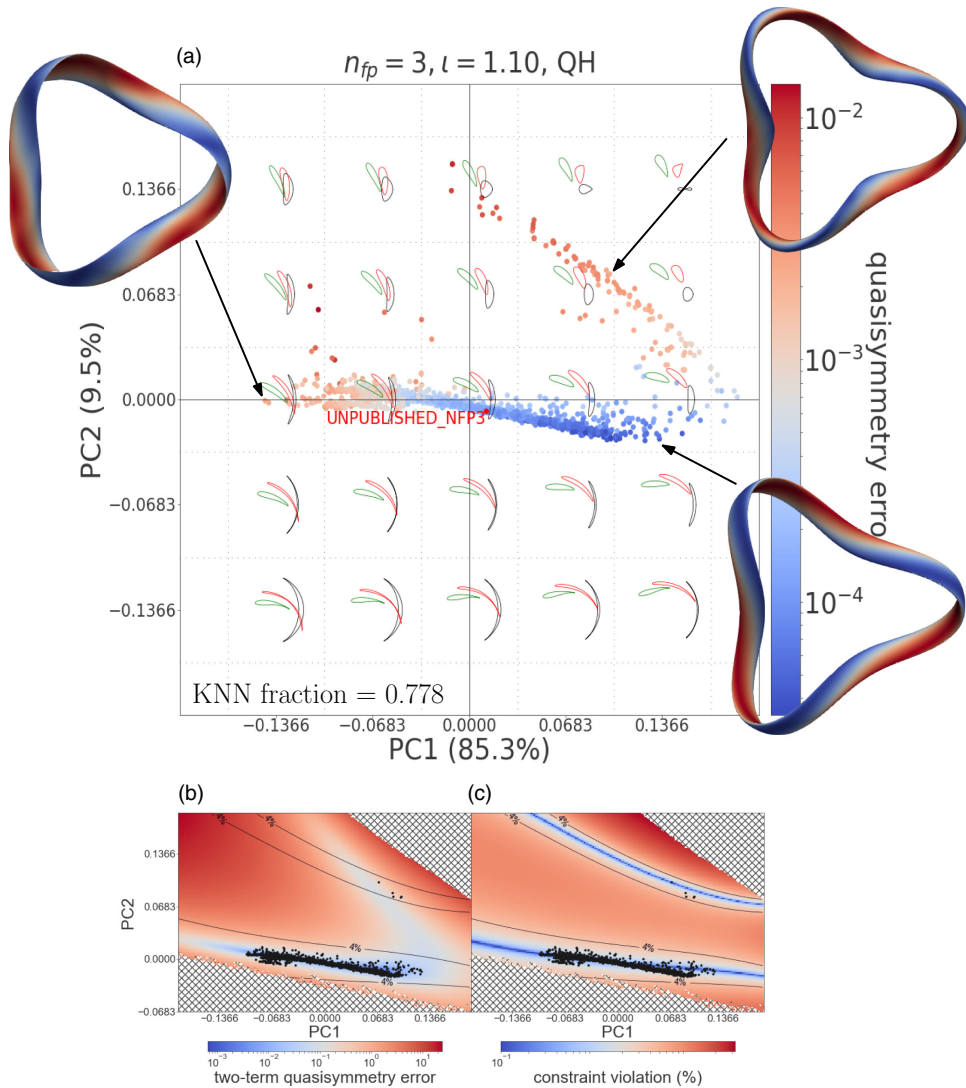


FIGURE 10. Panel (a) contains a PCA of QH symmetric devices in QUASR, along with an unpublished device from (Kappel *et al.* 2024), indicated with red text. The colour on the magnetic surfaces corresponds to the local field strength. Panels (b) and (c) show, respectively, the two-term QS error from Helander & Simakov (2008) and Landreman & Paul (2022) computed using VMEC and relative ι constraint violation ($|\iota - 1.1|/1.1$) on the PC subspace. The black isolines correspond to coordinates where the constraint violation is precisely 4% and the black dots are devices in panel (a) that have only two dominant principal components. The region on the PC plane where the two-term QS error could not be evaluated or VMEC did not converge is hatched.

(Kappel *et al.* 2024). The KNN fraction is quite high, 0.778. In contrast to the PCA in § 5.2, both principal components appear to modify the devices' elongation and the positioning of the cross-sections. We note that here, there does not appear to be two separate clusters of devices and the devices form a V-shaped continuum. Figure 10(b, c) contain the two-term QS error and ι constraint violation. The QUASR devices




FIGURE 11. The nonlinear manifold in figure 10 can be parametrized with a single coordinate using the isomap algorithm. The row of devices are uniformly sampled from the isomap embedding, and the numbers below these devices corresponds to the isomap coordinate and the quality of QS. The colour on the magnetic surfaces corresponds to the local field strength.

appear to follow the QS well throughout the plane (even suggesting the V-shape), but this clearly clashes with the rotational transform constraint. It does not in the lower portion of the cluster, which explains the large number of field cases found there. But it does in the middle portion.

Behind this apparent contradiction lies an important observation we emphasized before. Figure 10(b, c) correspond to configurations with only two principal components. But the QUASR devices truly live in a higher dimensional space. In particular, the configurations in the elbow of the V-shape cluster do, which are in fact the devices that do not fit with the interpretation of figure 10(b, c).

We may nevertheless attempt to describe the cluster in this scenario as a one-dimensional continuum. However, a PCA is inappropriate here as clearly the manifold is nonlinear. Instead, we use the isomap algorithm (Tenenbaum *et al.* 2000) to visualize the progression of the device geometries along the manifold. The isomap algorithm works by generating a connectivity graph between data points using the k -nearest neighbours of each node. Then, pairwise distances between points are computed by traversing the connectivity graph. The multidimensional scaling algorithm (Kruskal 1964) attempts to find a lower dimensional representation that preserves these pairwise distances. Our aim is to find a one-dimensional labelling of the devices so that the devices can be ordered, similar to § 5.2. Applying the isomap algorithm on the full $D = 663$ dimensional data points, we obtain the embedding shown in figure 11, i.e. each device is given a one-dimensional coordinate. Uniformly sampling this coordinate, the progression of devices in the bottom row of the figure is obtained.

5.5. Quasihelical symmetry ($n_{fp} = 5$, $\iota = 2.5$)

In this section, we show the two-dimensional landscape of devices with a rotational transform $\iota = 2.50$ and $n_{fp} = 5$ in figure 12. The first principal component appears to shift the relative positions of the cross-sections, while the second principal component appears to act on the elongation of the device (similar to the $n_{fp} = 4$ before). There appears to be two separate clusters of devices, where the second cluster contains the most quasisymmetric devices (again there seems to be a parallel to the $n_{fp} = 4$ scenario, even on the geometric behaviour within and between clusters).

However, a non-negligible amount of information was lost in this projection, as the sum of the PC ratios is just over 80 %. Similarly, the KNN fraction is just over 0.45. To provide a more accurate representation of the data, we must augment

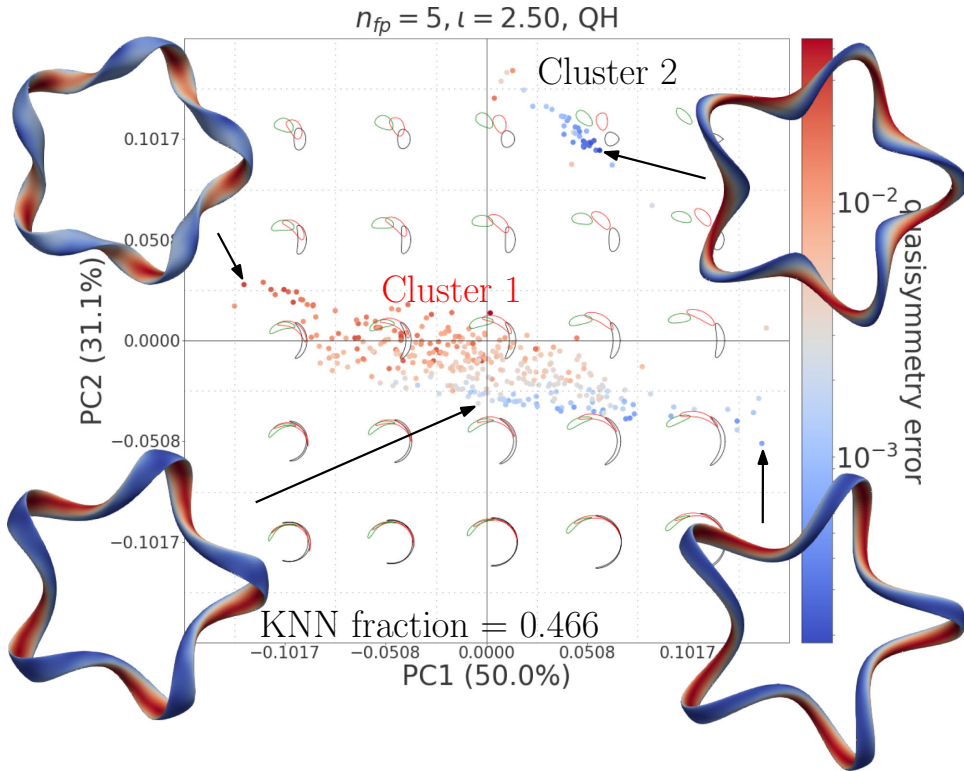


FIGURE 12. The PCA of QH symmetric devices in QUASR when $\iota = 2.5$ and $n_{fp} = 5$. Two clusters are labelled and devices from them are plotted. The colour on the magnetic surfaces corresponds to the local field strength.

figure 12 with a third principal component in figure 13. The two-dimensional projections of the data on the PC1–PC2, PC1–PC3 and PC2–PC3 planes are also provided to make the three-dimensional nature of the data more easily understood. Our projection error measures immediately improve as the sum of the PC ratios is approximately 90 %, and the KNN fraction increases to 0.678.

The third PC component evidences that the clusters are rather different: the first cluster is a three-dimensional cloud, while the second cluster appears to have very little variation in the third PC. The higher dimensionality of cluster 1 can be quantified by observing the cumulative PC ratio as the number of principal components used to represent cluster 1 and 2 in figure 13(a). It takes three PCs for the cumulative PC ratio of cluster 1 to pass 0.9, while it takes only two for cluster 2.

5.6. Quasisymmetry

For completeness, in figure 14(a), we complete a PCA on QA devices in QUASR with $n_{fp} = 2$ and $\iota = 0.4$. The cumulative PC ratio is just over 80 %, revealing that a non-negligible amount of information is lost in this projection, and it might be useful to include a third principal component in the visualization. Given the two principal axes, we project devices outside QUASR onto the figure and label them. This area of the parameter space is well explored. In figure 14, we do a PCA for QA devices, subselected based on their rotational transform ($\iota = 0.6$) and number of

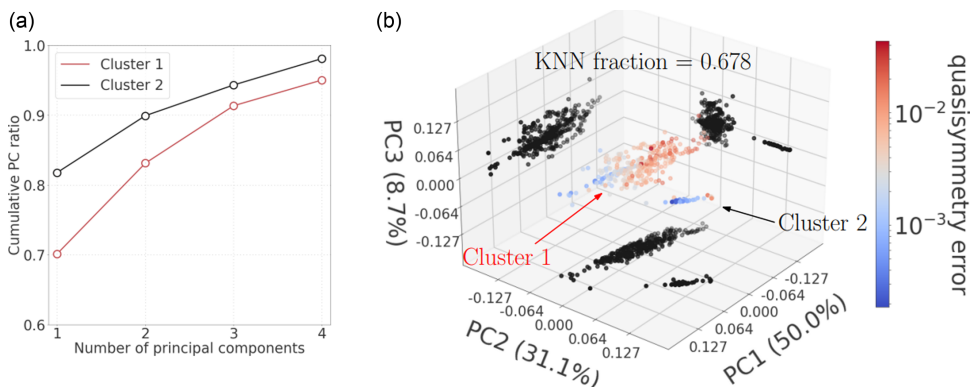


FIGURE 13. In panel (a), the cumulative PC ratio with increasing number of principal components used to represent the two clusters from figure 12. In panel (b), PCA of QH symmetric devices in QUASR when $\iota = 2.5$ and $n_{fp} = 5$, similar to figure 12 but with a third principal component. The two-dimensional projections of the data on the PC1–PC2, PC1–PC3 and PC2–PC3 planes are also provided in black to make the three-dimensional nature of the data more easily understood.

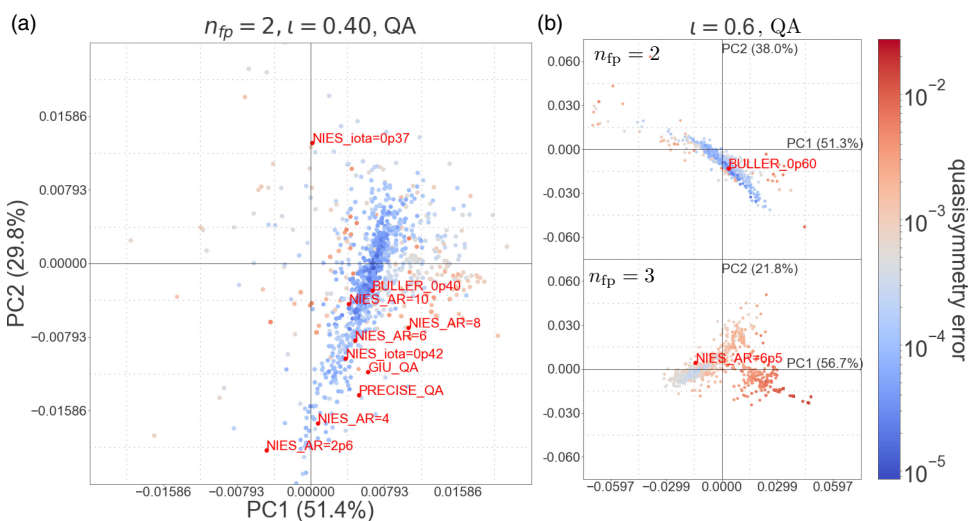


FIGURE 14. The PCA of QA devices in QUASR along with devices in the literature, indicated with red text.

field periods ($n_{fp} = 2, 3$). Labelled devices from the literature appear similar to the devices in QUASR. We do not provide cross-sections of hypothetical devices on the linear manifold as we did in previous sections because, visually, the cross-sections do not differ significantly from each other.

The analysis from § 4 reveals that the region of QA devices with favourable QS is small, and, moreover, it shrinks with increasing shaping of the device. This bears out in the PCA analysis here too, as we find that there is not much geometric diversity of the devices at the aspect ratio studied. Notably, the Landreman–Paul QA device (PRECISE_QA) is close on the plane to the device called GIU_QA. This device was

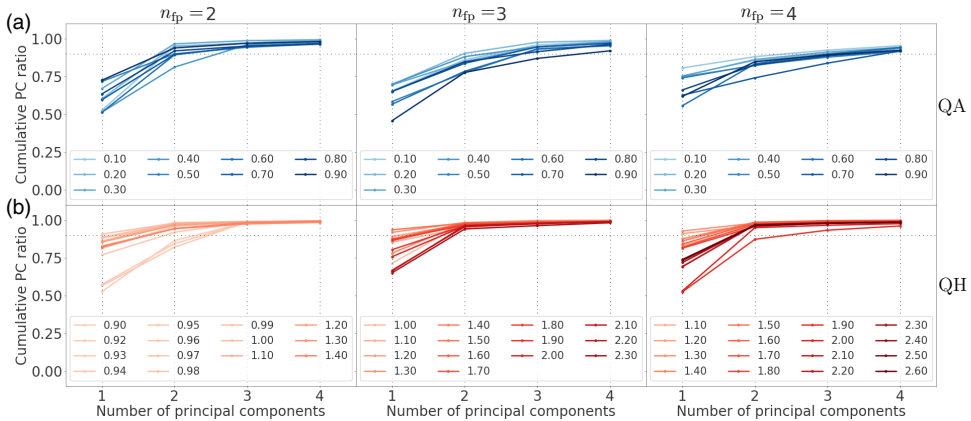


FIGURE 15. The cumulative PC ratio of subsets of the data for various combinations of n_{fp} , l and helicity. Panels (a) and (b) correspond to QA and QH devices, respectively. Subpanels (i), (ii) and (iii) correspond to particular values of n_{fp} , and the curves correspond to different values of l . The horizontal dashed line corresponds to a cumulative PC ratio of 0.9.

obtained by providing the PRECISE_QA device as an initial guess to a direct coil optimization algorithm. The quality of QS in these two devices differs by almost an order of magnitude, despite the geometrical proximity on the plot.

6. Discussion and conclusion

The main goal of this paper was to introduce a new data set of optimized stellarators, and provide tools to navigate the contents of it. We start by highlighting a few devices in QUASR first, and then move to analysing the data-set. First, we frame QUASR within the near-axis framework and its predicted landscape of quasisymmetric configurations. We find close correspondence of devices in QUASR to regions of favourable near-axis QS. However, this approach does not apply to the full set of devices in QUASR given the difference in assumptions that went into their construction, most notably constraints on elongation and simplicity of the shaping. A more model agnostic procedure to analyse the data are then presented in the form of a PCA. On the subsets of the data set that we have explored in this work, we find that the first and second principal components act predominantly on the elongation (and different field twisting), and relative positions of cross-sections of the magnetic surfaces. These two dimensions are able to capture the devices in the data set to a large degree. For some subsets of the data, we observe clear clusters. Within the clusters, we find evidence that the devices form a continuum whose shape follows that of the design constraint manifold.

For the sets presented, two to three principal components appear sufficient for visualization. This is true of even more subsets throughout the QUASR database, as illustrated in figure 15, where we compute a PCA for various combinations of n_{fp} and l for QA and QH devices. It is possible to attain a cumulative PC ratio of 0.9 consistently. This illustrates that much of the variation between the devices can be characterized on a low-dimensional, linear manifold. This points to the lower dimensional nature of the subset of quasisymmetric configurations in the larger space of stellarators. This echoes the analysis in Sengupta *et al.* (2023), where it was revealed that QS imposes a lower-dimensionality on the field strength. We would like to highlight though that this lower-dimensionality does not necessarily extend to the

coil design space, which is known to be complex with multiple minima (Wechsung *et al.* 2022a; Giuliani 2024). This is because there are many, vastly different coil sets that produce approximately the same magnetic field (Landreman 2017). This highlights that some sort of global search or a problem reformulation is indeed necessary for coil design.

Whether the information lost in the dimensionality reduction is significant depends on the application. For the purposes of visualization in this work, a cumulative PC ratio of 90 % appeared satisfactory. But in other applications, this might not be the case. Some useful extensions of this include using the very simple, linear formulae learned during the PCA to generate warm starts for optimization. Alternatively, one might attempt optimizing in the PC subspace. These directions are not explored in this paper.

Acknowledgements

The authors would like to thank Flatiron's Scientific Computing Core (SCC). We would also like to thank M. Berger, R. Jorge, M. Landreman, E. Paul, M. Rachh and L. Saul for helpful discussions.

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

Funding

E.R. was supported by a grant of the Alexander-von-Humboldt-Stiftung, Bonn, Germany, through a postdoctoral research fellowship.

Declaration of interests

The authors report no conflict of interest.

Data availability

The latest version of the QUASR data set can be found at <https://zenodo.org/doi/10.5281/zenodo.10050655>. It can be visually explored at quasr.flatironinstitute.org.

Appendix A. Fitting the second-order near-axis model to the devices

In this appendix we describe how a near-axis model to second order is constructed for a given field from the QUASR database.

The geometry of the field's magnetic axis can be directly constructed from the given coils, but an additional three parameters are needed that describe the second-order near-axis model in vacuum: B_0 , $\bar{\eta}$ and B_{2c} . In the neighbourhood of the magnetic axis, the field strength $B = \|\mathbf{B}\|$ of the near-axis model satisfies

$$B(r, \varphi, \theta) = B_0 + r B_0 \bar{\eta} \cos(\theta - N\varphi) + r^2 [B_{20}(\varphi) + B_{2c} \cos(2(\theta - N\varphi))] + \dots, \quad (\text{A.1})$$

where stellarator symmetry is assumed, φ, θ are Boozer angles and r is the radial distance off-axis. There are a multitude of methods to estimate these near-axis parameters in the field generated by the coils. For instance, one could evaluate the field strength on a three-dimensional tensor-product grid $(r_i, \varphi_j, \theta_k)$ in the neighbourhood of the magnetic axis and solve for the values of $B_0, \bar{\eta}, B_{2c}, B_{20}(\varphi_j)$ in a least-squares sense. Another possibility is to fit the predicted cross-sections of the near-axis magnetic surfaces to the true ones (Landreman 2019). These techniques both have the

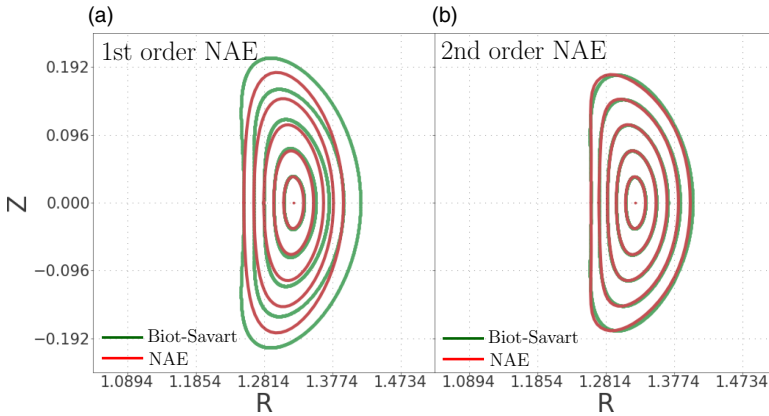


FIGURE 16. Poincaré plot (green) and cross-sections of surfaces generated by the near-axis expansion (red) fit to a device in QUASR. Panel (a) shows the first-order near-axis expansion (NAE) surfaces, while (b) shows the second-order NAE surfaces. The outermost surfaces here have aspect ratio ~ 9 . The field lines are generated from starting points along the inboard side with $Z = 0$. The Biot–Savart Poincaré sections are different in (a) and (b) because the field lines are initialized from different starting points.

difficulty that one must use information off axis, and the quality of the fit might be sensitive to the choice of how far off-axis one goes. In order to make the fitting procedure independent of this user-defined choice, we solve three minimization problems sequentially,

$$\begin{aligned}
 B_0^* &= \operatorname{argmin}_{B_0} \int_0^{2\pi} \|\mathbf{B}_{\text{coils}} - \mathbf{B}_{\text{NAE}}(B_0)\|^2 d\phi, \\
 \bar{\eta}^* &= \operatorname{argmin}_{\bar{\eta}} \int_0^{2\pi} \|\nabla \mathbf{B}_{\text{coils}} - \nabla \mathbf{B}_{\text{NAE}}(B_0^*, \bar{\eta})\|^2 d\phi, \\
 B_{2c}^* &= \operatorname{argmin}_{B_{2c}} \int_0^{2\pi} \|\nabla \nabla \mathbf{B}_{\text{coils}} - \nabla \nabla \mathbf{B}_{\text{NAE}}(B_0^*, \bar{\eta}^*, B_{2c})\|^2 d\phi, \quad (\text{A.2})
 \end{aligned}$$

where ϕ is the standard cylindrical angle. Since the magnetic field is generated by a set of coils, we can compute the magnetic field as well as its first and second derivatives using the Biot–Savart law, which we have direct access to in SIMSOPT. As a result, we can find the near-axis parameters that most closely reproduce the coils' magnetic field and its derivatives on axis. Note that the near-axis model assumes that there is perfect QS to first order. While a good approximation of QS is attainable in the volume, the devices do not present perfect near-axis QS, and we may not always be able to obtain a good fit of the coils' magnetic field.

In figure 16, we provide a Poincaré plot alongside cross-sections of surfaces generated by a first- and second-order near-axis expansion. The NAE model was fit using the method described in this section, revealing that the second-order NAE is capable of describing the magnetic surfaces in the neighbourhood of the magnetic axis much more faithfully than the first-order one.

	# PC components	1	2	3
Cumulative PC ratios	Cluster 1	0.828	0.916	1.000
	Cluster 2	0.501	0.953	1.000
	Cluster 3	0.361	0.712	1.000
KNN fraction	Cluster 1	0.216	0.468	1.000
	Cluster 2	0.140	0.614	1.000
	Cluster 3	0.095	0.352	1.000

TABLE 2. The top three rows are the cumulative PC ratio and the bottom three rows are the KNN fraction with respect to the number of PCs for the clusters in [figure 4](#).

Appendix B. The PCA example for Gaussian processes

To gain some intuition on the quality measures of cumulative PC ratios and KNN fraction introduced in § 5, we apply PCA to data generated from three Gaussians in this appendix. These Gaussians are defined to have mean and covariance matrices,

$$\begin{aligned}
 5\mu_1 &= \begin{pmatrix} -6 \\ 0 \\ 0 \end{pmatrix}, \Sigma_1 = Q_1^T \begin{pmatrix} 1 & 0 & 0 \\ 0 & 10^{-1} & 0 \\ 0 & 0 & 10^{-1} \end{pmatrix} Q_1 \\
 \mu_2 &= \begin{pmatrix} 6 \\ -6 \\ -6 \end{pmatrix}, \Sigma_2 = Q_2^T \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 10^{-1} \end{pmatrix} Q_2, \\
 \mu_3 &= \begin{pmatrix} 6 \\ 0 \\ 0 \end{pmatrix}, \Sigma_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},
 \end{aligned} \tag{B.1}$$

where Q_1, Q_2 are randomly generated rotation matrices. Individually, the Gaussians are well approximated by linear manifolds of dimension one, two and three, respectively. Samples from these distributions are shown in [figure 4\(a\)](#) in the main text, while a two-dimensional projection of the data is shown in [figure 4\(b\)](#).

In two-dimensional, the cumulative PC ratio is quite high 0.986, revealing that the majority of variation in the data are captured by two principal components. Despite this favourable quality measure and such a simple data set, the KNN fraction is approximately 0.4, i.e. $\approx 40\%$ of the 10 nearest neighbours in three dimensions remain so on the two-dimensional projection on average.

In [table 2](#), we provide the cumulative PC ratio as a function of the number of PCs. In order for the PC ratio to attain 0.9, each cluster requires a different number of PCs. This hints that cluster 1 is well approximated on a line, cluster 2 on a plane and cluster 3 on a volume. In fact, since cluster 3 is genuinely three-dimensional, dimensionality reduction is not possible without significant loss of information. The KNN fractions increase with the number of PCs, with a value greater than 0.4 appearing to be a lower bound on a reasonable lower-dimensional representation.

Label	n_{fp}	Aspect ratio	ι	QS type	Figure #	Reference
BULLER_0p40	2	6.00	0.40 (m)	QA	14	Buller <i>et al.</i> (2024)
GIU_QA	2	6.00	0.42 (m)	QA	14	Giuliani <i>et al.</i> (2022b)
NIES_iota = 0p37	2	6.00	0.37 (e)	QA	14	Nies <i>et al.</i> (2024)
NIES_iota = 0p42	2	6.00	0.42 (e)	QA	14	Nies <i>et al.</i> (2024)
NIES_AR = 10	2	10.0	0.42 (e)	QA	14	Nies <i>et al.</i> (2024)
NIES_AR = 8	2	8.00	0.42 (e)	QA	14	Nies <i>et al.</i> (2024)
NIES_AR = 6	2	6.00	0.42 (e)	QA	14	Nies <i>et al.</i> (2024)
NIES_AR = 4	2	4.00	0.42 (e)	QA	14	Nies <i>et al.</i> (2024)
NIES_AR = 2p6	2	2.60	0.42 (e)	QA	14	Nies <i>et al.</i> (2024)
PRECISE_QA	2	6.00	0.42 (m)	QA	14	Landreman & Paul (2022)
BULLER_0p60	2	6.00	0.60 (m)	QA	14	Buller <i>et al.</i> (2024)
NIES_AR = 6p5	3	6.50	0.42 (e)	QA	14	Nies <i>et al.</i> (2024)
UNPUBLISHED_NFP3	3	6.00	1.06 (m)	QH	12	Kappel <i>et al.</i> (2024)
NEW_QH	4	13.5	2.29 (m)	QH	6	Rodríguez <i>et al.</i> (2023)
HSX (vacuum)	4	10.0	1.06 (m)	QH	9	Anderson <i>et al.</i> (1995)
WISTELL-A	4	6.72	1.1 (m)	QH	9	Bader <i>et al.</i> (2020)
NIES_AR = 9	4	9.00	1.3 (e)	QH	9	Nies <i>et al.</i> (2024)

TABLE 3. Labelling of devices from the literature, the figure number on which they appear in this manuscript, along with the value of n_{fp} , type of QS targeted, aspect ratio and rotational transform. We quote either the mean rotational transform, or the edge rotational transform. These separate cases are highlighted with an ‘m’ or ‘e’ in parentheses.

Appendix C. Reference QS configurations from the literature

The QS configurations from the literature referenced throughout the paper are summarized in table 3. These devices were all optimized for QS at various aspect ratios, and target rotational transforms. Moreover, some devices were designed to have a target mean rotational transform (Landreman & Paul 2022; Giuliani *et al.* 2022b), while others had a target edge rotational transform (Nies *et al.* 2024).

To make a direct comparison of these to the devices in the QUASR dataset, we must construct their corresponding feature vectors using the following procedure. First, we take the geometry of the device’s outermost surface and use VMEC (Hirshman & Whitson 1983) to determine the vacuum magnetic field on the entire toroidal volume. We take the surface in the device that has aspect ratio closest to 10 for QA and 12 for QH. The Boozer coordinates on this surface are computed using BOOZ_XFORM (Sanchez *et al.* 2000). Finally, we compute the Fourier harmonics of this surface parametrized in Boozer coordinates. We may then treat them as we did

with QUASR objects, and when necessary project them to the appropriate principal component plots. We directly compare these configurations to the relevant ones in QUASR when they share the value of n_{fp} , helicity and mean rotational transform (out to aspect ratio 10 (for QA) and 12 (for QH) to an accuracy of ± 0.05).

REFERENCES

- ANDERSON, F.S.B., ALMAGRI, A.F., ANDERSON, D.T., MATTHEWS, P.G., TALMADGE, J.N. & SHOHET, J.L. 1995 The helically symmetric experiment, (HSX) goals, design and status. *Fusion Technol.* **27**, 273–277.
- BADER, A. *et al.* 2020 Advancing the physics basis for quasi-helically symmetric stellarators. *J. Plasma Phys.* **86**, 905860506.
- BISHOP, C. 2006 *Pattern Recognition and Machine Learning*. Springer.
- DE BLANK, H.J. 2004 Guiding center motion. *Fusion Sci. Technol.* **45**, 47–54.
- BONHOMME, V., PICQ, S., GAUCHEREL, C. & CLAUDE, J. 2014 Momocs: outline analysis using R. *J. Stat. Softw.* **56**, 1–24.
- BOOZER, A.H. 1981a Plasma equilibrium with rational magnetic surfaces. *Tech. Rep.*, Princeton Univ.
- BOOZER, A.H. 1981b Transport and isomorphic equilibria. *Tech. Rep.*, Princeton Plasma Physics Lab.(PPPL).
- BOURLARD, H. & KAMP, Y. 1988 Auto-association by multilayer perceptrons and singular value decomposition. *Biol. Cybern.* **59**, 291–294.
- BULLER, S., LANDREMAN, M., KAPPEL, J. & GAUR, R. 2024 A family of quasi-axisymmetric stellarators with varied rotational transform. arXiv: [2401.09021](https://arxiv.org/abs/2401.09021).
- BURBY, J.W., KALLINIKOS, N. & MACKAY, R.S. 2020 Some mathematics for quasi-symmetry. *J. Math. Phys.* **61** (9), 093503.
- ERIKSSON, D., PEARCE, M., GARDNER, J., TURNER, R.D. & POLOCZEK, M. 2019 Scalable global optimization via local Bayesian optimization. In *Advances in Neural Information Processing Systems*, (ed. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox & R. Garnett) vol. 32, pp. 5496–5507. Curran Associates, Inc. Available at: https://proceedings.neurips.cc/paper_files/paper/2019/file/6c990b7aca7bc7058f5e98ea909e924b-Paper.pdf.
- FULLER JR, E.J. 1999 *The Geometric and Topological Structure of Holonomic Knots*. University of Georgia.
- GARREN, D.A. & BOOZER, A.H. 1991 Existence of quasihelically symmetric stellarators. *Phys. Fluids B: Plasma Phys.* **3**, 2822–2834.
- GIULIANI, A. 2024 Direct stellarator coil design using global optimization: application to a comprehensive exploration of quasi-axisymmetric devices. *J. Plasma Phys.* **90**, 905900303.
- GIULIANI, A., WECHSUNG, F., CERFON, A., LANDREMAN, M. & STADLER, G. 2023 Direct stellarator coil optimization for nested magnetic surfaces with precise quasi-symmetry. *Phys. Plasmas* **30**, 042511.
- GIULIANI, A., WECHSUNG, F., CERFON, A., STADLER, G. & LANDREMAN, M. 2022a Single-stage gradient-based stellarator coil design: optimization for near-axis quasi-symmetry. *J. Comput. Phys.* **459**, 111147.
- GIULIANI, A., WECHSUNG, F., STADLER, G., CERFON, A. & LANDREMAN, M. 2022b Direct computation of magnetic surfaces in Boozer coordinates and coil optimization for quasisymmetry. *J. Plasma Phys.* **88**, 905880401.
- HELANDER, P. 2014 Theory of plasma confinement in non-axisymmetric magnetic fields. *Rep. Prog. Phys.* **77**, 087001.
- HELANDER, P. & SIMAKOV, A.N. 2008 Intrinsic ambipolarity and rotation in stellarators. *Phys. Rev. Lett.* **101**, 145003.
- HINDENLANG, F.J., PLUNK, G.G. & MAJ, O. 2024 A generalized Frenet frame for computing MHD equilibria in stellarators. arXiv: [2410.17595](https://arxiv.org/abs/2410.17595).
- HIRSHMAN, S.P. & WHITSON, J.C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Phys. Fluids* **26**, 3553–3568.

- JOLLIFFE, I.T. 1990 Principal component analysis: a beginner's guide—I. Introduction and application. *Weather* **45**, 375–382.
- JORGE, R., GIULIANI, A. & LOIZU, J. 2024 Simplified and flexible coils for stellarators using single-stage optimization. *Phys. Plasmas* **31**, 112501.
- JORGE, R., GOODMAN, A., LANDREMAN, M., RODRIGUES, J. & WECHSUNG, F. 2023 Single-stage stellarator optimization: combining coils with fixed boundary equilibria. *Plasma Phys. Control. Fusion* **65**, 074003.
- KAPPEL, J., LANDREMAN, M. & MALHOTRA, D. 2024 The magnetic gradient scale length explains why certain plasmas require close external magnetic coils. *Plasma Phys. Control. Fusion* **66**, 025018.
- KAPTANOGLU, A.A., LANGLOIS, G.P. & LANDREMAN, M. 2024 Topology optimization for inverse magnetostatics as sparse regression: application to electromagnetic coils for stellarators. *Comput. Meth. Appl. Mech. Engng* **418**, 116504.
- KAPTANOGLU, A.A., WIEDMAN, A., HALPERN, J., HURWITZ, S., PAUL, E.J. & LANDREMAN, M. 2025 Reactor-scale stellarators with force and torque minimized dipole coils. *Nucl. Fusion* **65**, 046029.
- KOBAK, D. & BERENS, P. 2019 The art of using t-SNE for single-cell transcriptomics. *Nat. Commun.* **10**, 5416.
- KRUSKAL, J.B. 1964 Nonmetric multidimensional scaling: a numerical method. *Psychometrika* **29**, 115–129.
- LANDREMAN, M. 2017 An improved current potential method for fast computation of stellarator coil shapes. *Nucl. Fusion* **57**, 046003.
- LANDREMAN, M. 2019 Optimized quasisymmetric stellarators are consistent with the Garren–Boozor construction. *Plasma Phys. Control. Fusion* **61**, 075001.
- LANDREMAN, M. 2022 Mapping the space of quasisymmetric stellarators using optimized near-axis expansion. *J. Plasma Phys.* **88**, 905880616.
- LANDREMAN, M., JORGE, R., RODRIGUEZ, E. & DUDT, D. 2024 *landreman/pyQSC*. GitHub repository.
- LANDREMAN, M., MEDASANI, B., WECHSUNG, F., GIULIANI, A., JORGE, R. & ZHU, C. 2021 SIMSOPT: a flexible framework for stellarator optimization. *J. Open Source Softw.* **6**, 3525.
- LANDREMAN, M. & PAUL, E. 2022 Magnetic fields with precise quasisymmetry for plasma confinement. *Phys. Rev. Lett.* **128**, 035001.
- LANDREMAN, M. & SENGUPTA, W. 2019 Constructing stellarators with quasisymmetry to high order. *J. Plasma Phys.* **85**, 815850601.
- LANDREMAN, M., SENGUPTA, W. & PLUNK, G.G. 2019 Direct construction of optimized stellarator shapes. Part 2. Numerical quasisymmetric solutions. *J. Plasma Phys.* **85**, 905850103.
- LEE, J.A. & VERLEYSSEN, M. 2009 Quality assessment of dimensionality reduction: rank-based criteria. *Neurocomputing* **72**, 1431–1443.
- VAN DER MAATEN, L. & HINTON, G. 2008 Visualizing data using t-SNE. *J. Machine Learning Res.* **9** (Nov), 2579–2605.
- MOFFATT, H.K. & RICCA, R.L. 1992 Helicity and the $\text{c\acute{a}lug\text{ă}reanu}$ invariant. *Proc. R. Soc. Lond. A: Math. Phys. Sci.* **439**, 411–429.
- NAJMABADI, F. *et al.* 2008 The ARIES-CS compact stellarator fusion power plant. *Fusion Sci. Technol.* **54**, 655–672.
- NIES, R., PAUL, E.J., PANICI, D., HUDSON, S.R. & BHATTACHARJEE, A. 2024 Exploration of the parameter space of quasisymmetric stellarator vacuum fields through adjoint optimisation. *J. Plasma Phys.* **90** (6), 905900620.
- NÜHRENBURG, J. & ZILLE, R. 1988 Quasi-helically symmetric toroidal stellarators. *Phys. Lett. A* **129**, 113–117.
- OBERTI, C. & RICCA, R.L. 2016 On torus knots and unknots. *J. Knot Theory Ramific.* **25**, 1650036.
- PANICI, D., CONLIN, R., DUDT, D.W., UNALMIS, K. & KOLEMEN, E. 2023 The DESC stellarator code suite. Part 1. Quick and accurate equilibria computations. *J. Plasma Phys.* **89**, 955890303.
- PAUL, E.J., BHATTACHARJEE, A., LANDREMAN, M., ALEX, D., VELASCO, J.L. & NIES, R. 2022 Energetic particle loss mechanisms in reactor-scale equilibria close to quasisymmetry. *Nucl. Fusion* **62**, 126054.
- RODRIGUEZ, E., HELANDER, P. & BHATTACHARJEE, A. 2020 Necessary and sufficient conditions for quasisymmetry. *Phys. Plasmas* **27** (6), 062501.

- RODRIGUEZ, E., PAUL, E.J. & BHATTACHARJEE, A. 2022a Measures of quasisymmetry for stellarators. *J. Plasma Phys.* **88**, 905880109.
- RODRIGUEZ, E., SENGUPTA, W. & BHATTACHARJEE, A. 2022b Phases and phase-transitions in quasisymmetric configuration space. *Plasma Phys. Control. Fusion* **64**, 105006.
- RODRÍGUEZ, E., SENGUPTA, W. & BHATTACHARJEE, A. 2023 Constructing the space of quasisymmetric stellarators through near-axis expansion. *Plasma Phys. Control. Fusion* **65**, 095004.
- SANCHEZ, R., HIRSHMAN, S.P., WARE, A.S., BERRY, L.A. & SPONG, D.A. 2000 Ballooning stability optimization of low-aspect-ratio stellarators. *Plasma Phys. Control. Fusion* **42**, 641–653.
- SENGUPTA, W., NIKULSIN, N., BULLER, S., MADAN, R., PAUL, E.J., NIES, R., KAPTANOGLU, A.A., HUDSON, S.R. & BHATTACHARJEE, A. 2023 Periodic Korteweg-de Vries soliton potentials generate magnetic field strength with excellent quasisymmetry. arXiv: [2302.13924](https://arxiv.org/abs/2302.13924).
- SPITZER, L. 1958 The stellarator concept. *Phys. Fluids* **1**, 253–264.
- TENENBAUM, J.B., SILVA, V.DE & LANGFORD, J.C. 2000 A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319–2323.
- TREFETHEN, L.N. & WEIDEMAN, J.A.C. 2014 The exponentially convergent trapezoidal rule. *SIAM Rev.* **56**, 385–458.
- WECHSUNG, F., GIULIANI, A., LANDREMAN, M., CERFON, A. & STADLER, G. 2022a Single-stage gradient-based stellarator coil design: stochastic optimization. *Nucl. Fusion* **62**, 076034.
- WECHSUNG, F., LANDREMAN, M., GIULIANI, A., CERFON, A. & STADLER, G. 2022b Precise stellarator quasi-symmetry can be achieved with electromagnetic coils. *Proc. Natl Acad. Sci.* **119**, e2202084119.
- WIEDMAN, A., BULLER, S. & LANDREMAN, M. 2024 Coil optimization for quasi-helically symmetric stellarator configurations. *J. Plasma Phys.* **90**, 905900307.
- ZHU, C., HUDSON, S.R., SONG, Y. & WAN, Y. 2017 New method to design stellarator coils without the winding surface. *Nucl. Fusion* **58**, 016008.