How does ion temperature gradient turbulence depend on magnetic geometry? Insights from data and machine learning

Matt Landreman[1] , **Jong Youl Choi[2], Caio Alves[2], Prasanna Balaprakash[2],**
Michael Churchill[3], Rory Conlin[1] and **Gareth Roberg-Clark[4]**

[1]Institute for Research in Electronics & Applied Physics, University of Maryland, College Park, MD 20742, USA
[2]Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA
[3]Princeton Plasma Physics Laboratory, Princeton, NJ 08540, USA
[4]Max Planck Institute for Plasma Physics, Wendelsteinstraße 1, Greifswald 17491, Germany
Corresponding author: Matt Landreman, mattland@umd.edu

Magnetic geometry has a significant effect on the level of turbulent transport in fusion plasmas. Here, we model and analyse this dependence using multiple machine learning methods and a dataset of $>200\,000$ nonlinear gyrokinetic simulations of ion-temperature-gradient turbulence in diverse non-axisymmetric geometries. The dataset is generated using a large collection of both optimised and randomly generated stellarator equilibria. At fixed gradients and other input parameters, the turbulent heat flux varies between geometries by several orders of magnitude. Trends are apparent among the configurations with particularly high or particularly low heat flux. Regression and classification techniques from machine learning are then applied to extract patterns in the dataset. Due to a symmetry of the gyrokinetic equation, the heat flux and regressions thereof should be invariant to translations of the raw features in the parallel coordinate, similar to translation invariance in computer vision applications. Multiple regression models including convolutional neural networks (CNNs) and decision trees can achieve reasonable predictive power for the heat flux in held-out test configurations, with highest accuracy for the CNNs. Using Spearman correlation, sequential feature selection and Shapley values to measure feature importance, it is consistently found that the most important geometric lever on the heat flux is the flux surface compression in regions of bad curvature. The second most important geometric feature relates to the magnitude of geodesic curvature. These two features align remarkably with surrogates that have been proposed based on theory, while the methods here allow a natural extension to more features for increased accuracy. The dataset, released with this publication, may also be used to test other proposed surrogates, and we find that many previously published proxies do correlate well with both the heat flux and stability boundary.

Key words: fusion plasma, plasma confinement, plasma nonlinear phenomena

1. Introduction

The level of turbulent transport in magnetised plasmas depends on the magnetic field geometry in complicated ways. In tokamaks, geometric factors include the aspect ratio, elongation and triangularity, while in stellarators, there is vast additional freedom in the non-axisymmetric shaping. While direct numerical simulation of the turbulence can be used to evaluate the transport for specific geometries, numerical calculations do not immediately give insight into general parametric dependencies. Numerical turbulence simulations also require non-negligible computation time and are noisy due to the chaotic dynamics, making it challenging to include them directly in optimisation or parametric studies. Therefore, a variety of physics-based surrogates for the geometry-dependence of turbulent transport have been proposed (Mynick, Pomphrey & Xanthopoulos 2010; Mynick *et al*. 2014; Xanthopoulos *et al*. 2014; Proll *et al*. 2015; Hegna, Terry & Faber 2018; Mackenbach, Proll & Helander 2022; Nakata & Matsuoka 2022; Nakayama *et al*. 2023; Roberg-Clark *et al*. 2023; Goodman *et al*. 2024). However, owing to advances in computing hardware and simulation software, it is now possible to assemble large sets of turbulence simulation data that span a wide range of possible geometries, even in the high-dimensional parameter space of stellarators. Moreover, progress in machine learning (ML) methods and software provide many new opportunities for understanding and exploiting these data. The purpose of this article is to show how ML methods can discover patterns in the geometry-dependence of plasma turbulence. Specifically, we present a new method for generating training data using random three-dimensional (3-D) geometries, we compare the accuracy of several ML methods at predicting the turbulent heat flux, and we show several interpretable ML techniques that can identify which geometric factors determine the turbulent heat flux. This last aspect shows that ML can be more than a black-box interpolation method to accelerate computations – it can in fact also feed back into more traditional physics analysis. For example, the patterns in the data identified here can provide evidence for or against physics-inspired approximate models and motivate future theoretical studies.

The methods here extend prior work in several ways. ML methods, specifically neural networks, have been applied previously as surrogates for transport in tokamaks (Meneghini *et al*. 2014; Citrin *et al*. 2015; Meneghini *et al*. 2017; Narita *et al*. 2019; Honda & Narita 2019; van de Plassche *et al*. 2020; Abbate, Conlin & Kolemen 2021; Boyer & Chadwick 2021; Li *et al*. 2024). In contrast to this previous work, here we allow for non-axisymmetric shaping, in which case, the geometric parameter space is much higher dimensional. In the context of stellarators, neural networks have been applied to neoclassical transport coefficients (Wakasa *et al*. 2007) and magneohydrodynamic (MHD) equilibria (Merlo *et al*. 2021; Curvo, Ferreira & Jorge 2025). For stellarator turbulent transport, regressions have been performed using theory-based analytic models with tuning parameters (Nakayama *et al*. 2023). Theory-based surrogates for transport in stellarators have previously been compared with gyrokinetic simulations for a small number of geometries N, e.g. $N = 10$ by Proll *et al*. (2015), $N = 4$ by Mackenbach *et al*. (2022) or $N = 9$ by Roberg-Clark *et al*. (2023). In the present work, the number of geometries considered is increased by four orders of magnitude to $N > 10^5$. Compared with these earlier studies, the work here is also unique in the new method of data generation using random novel 3-D geometries, in the application of both neural-network- and non-neural-network-based ML methods, and in the use of interpretable ML methods to identify important features in the geometry. Finally, in the supplemental material available from Landreman (2025), we are making the training dataset publicly

available, so other researchers can test the accuracy of different proposed turbulence surrogates.

To obtain ML models for this problem that are interpretable, we demonstrate a novel approach that combines a large library of candidate features with forward sequential feature selection (FSFS) and traditional regression and classification methods. The new approach here is used to ensure that the models respect a translation-invariance in the gyrokinetic equation: the heat flux should be invariant to periodic translation of all geometric quantities in the direction along the magnetic field (see §2). The invariance is guaranteed in our models by using a library of features that respect the symmetry. The use of a combinatorial library of candidate features in our method is reminiscent of SINDy (sparse identification of nonlinear dynamics) (Brunton, Proctor & Kutz 2016), except that we are not interested in time-dependence and we obtain parsimony through FSFS rather than via sparsity-promoting optimisation. Our method is also reminiscent of symbolic regression (Koza 1994) in that we seek symbolic expressions for interpretability. However, unlike SINDy and symbolic regression, the method here combines the symbolic feature library with additional regression and classification models (decision trees and nearest-neighbours) to efficiently allow extra nonlinearity.

We find a remarkable alignment between the results of interpretable ML analysis here and recently proposed physics-inspired surrogates for turbulence. Across several ML regression methods and several ways to measure feature importance, the most important two features are consistent. The most important geometric factor is found to be the flux surface compression $|\nabla \psi|$ in regions of bad curvature, where $2\pi\psi$ is the toroidal flux, reflecting the gradient drive in real space in regions of linear instability. This factor has been used as an optimisation objective function for ion temperature gradient (ITG) turbulence by Mynick *et al*. (2014), Xanthopoulos *et al*. (2014), Stroteich *et al*. (2022) and Goodman *et al*. (2024), and a more complicated objective with these elements was used previously by Mynick *et al*. (2010). The second most important geometric factor identified in our analysis is the average magnitude of the geodesic curvature, equivalent to the radial magnetic drift $\propto \boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla\psi$, where $\boldsymbol{\kappa}$ is the curvature. This quantity has been proposed as a correlate of turbulence by Xanthopoulos *et al*. (2011) and Nakata & Matsuoka (2022). The theoretical motivation for this quantity is related to zonal flows: larger geodesic curvature leads to stronger linear damping of zonal flows, resulting in higher levels of saturated heat flux. For both of the two most important features, our analysis matches the theoretical predictions for the sign of the correlation: increased $|\nabla\psi|$ and increased absolute geodesic curvature correlate with heat flux increase.

Since ML methods are most effective when large amounts of data are available that cover the parameter space of interest, we make several simplifying assumptions in this work. First, we consider the electrons to be adiabatic. This choice makes the direct numerical simulations faster by a factor of $\sim \sqrt{m_i/m_e}$, where m_i and m_e are the ion and electron masses. Hence, the maximum stable time step is not limited by the electron parallel speed, increasing the number of simulations that can be run for a given computational budget. Compared with simulations with adiabatic electrons, simulations with kinetic electrons have higher heat fluxes (Chen *et al*. 2003) and the difference depends on magnetic geometry (McKinney *et al*. 2019; Goodman *et al*. 2024), so some conclusions of the analysis here may be modified if revisited with kinetic electrons. By the use of adiabatic electrons, our analysis is necessarily electrostatic and focuses on ion temperature gradient turbulence.

When attempting to model the turbulent heat flux, there are two natural options for the independent geometric variables: one could either use shape parameters of the plasma or the geometric quantities appearing directly in the gyrokinetic equation. Examples of the former include elongation and triangularity, or Fourier amplitudes of the boundary surface, while examples of the latter include the field magnitude $|B|$ and the guiding centre drift component $B^{-3}\boldsymbol{B} \times \boldsymbol{\nabla} B \boldsymbol{\cdot} \boldsymbol{\nabla}\psi$. By taking the independent variables to be the plasma shape parameters, an ML regression effectively models physics of both MHD equilibrium and turbulence. If the independent variables are instead taken to be the geometry factors in the gyrokinetic equation, the ML regression becomes a model only for the turbulence, without including MHD equilibrium physics. Here, we take the latter approach, because by focusing on a narrower aspect of the physics, there is greater hope for interpretability of the model. At the same time, we nonetheless generate the geometric inputs to the gyrokinetic equation from actual global MHD equilibria so correlations and constraints among the geometric features are respected.

To introduce the methods here in detail, we begin in the following section by reviewing the gyrokinetic equation used in the direct numerical simulations, highlighting the geometric inputs and its translation invariance. This translation invariance is important because we will want ML models to respect this symmetry. The procedure for generating data, including both MHD equilibria and turbulence simulations, is then detailed in §3. Neural network fits to the data are presented in §4. Section 5 presents alternative ML methods that do not use neural networks, allowing for greater interpretability at the expense of reduced accuracy. These alternative methods require manual feature engineering and feature selection. In §6, it is demonstrated how the same dataset can be used to assess other proposed proxies for turbulence, by evaluating the accuracy of several ITG objective functions from earlier papers. Finally, we conclude and discuss future directions in §7.

2. Relevant properties of the gyrokinetic equation

Here we review the electrostatic gyrokinetic turbulence model for a flux tube to identify the geometric features that appear. A similar discussion can be found from Jorge & Landreman (2020). The magnetic field can be written $\boldsymbol{B} = \boldsymbol{\nabla}\psi \times \boldsymbol{\nabla}\alpha$, where $2\pi\psi$ is the toroidal flux and α is a field line label. Coordinates (x, y, z) are introduced where z is the arclength along the field line, satisfying $\boldsymbol{B} \boldsymbol{\cdot} \boldsymbol{\nabla} z = B$, and $x = x(\psi)$ and $y = y(\alpha)$ are functions of (ψ, α) that each resemble a distance. (The results of this section are independent of the exact choice of these two functions.) We consider a flux tube, for which the domain's small extent in x and y is of the order of a few gyroradii, much smaller than typical equilibrium scales like the minor radius. Therefore, x and y (or ψ and α) can be considered fixed for all equilibrium quantities. However, we do care about gyroradius-scale fluctuations in the distribution function and electrostatic potential, so the x- and y-dependence of these perturbations is retained. Fluctuating quantities are taken to vary with z on the same scale length as the equilibrium, and the extent of the flux tube domain in z is of a comparable scale.

The fluctuating electrostatic potential Φ is Fourier expanded as

$$\Phi(x, y, z, t) = \sum_{\boldsymbol{k}} \hat{\Phi}_{\boldsymbol{k}}(z, t) \exp(ik_x x + ik_y y), \qquad (2.1)$$

where $\boldsymbol{k}$ denotes (k_x, k_y). The distribution function of species s is expanded as

$$f_s = F_{Ms} - \frac{e_s \Phi}{T_s} F_{Ms} + h_s, \tag{2.2}$$

where $F_{Ms} = n_s (2\pi)^{-3/2} v_s^{-3} \exp(-v^2/(2v_s^2))$ is the leading-order Maxwellian with density n_s and temperature T_s, v is the speed, $v_s = \sqrt{T_s/m_s}$ is the thermal speed, and e_s is the species charge. The non-adiabatic part of the distribution, h_s, has a Fourier expansion

$$h_s(X, Y, z, v_{||}, \mu, t) = \sum_{\boldsymbol{k}} \hat{h}_{s,\boldsymbol{k}}(z, v_{||}, \mu, t) \exp(i k_x X + i k_y Y), \tag{2.3}$$

where (X, Y) are the values of (x, y) at the guiding centre position, $\mu = v_\perp^2/(2B)$, and $v_{||}$ and $v_\perp$ are the velocity components along $\boldsymbol{B}$ or perpendicular to it. The non-adiabatic distribution is computed by evolving the gyrokinetic equation (Frieman & Chen 1982),

$$\frac{\partial \hat{h}_{s,\boldsymbol{k}}}{\partial t} + v_{||} \frac{\partial \hat{h}_{s,\boldsymbol{k}}}{\partial z} - \mu \frac{\partial B}{\partial z} \frac{\partial \hat{h}_{s,\boldsymbol{k}}}{\partial v_{||}} + \boldsymbol{v}_d \cdot \nabla \hat{h}_{s,\boldsymbol{k}} + \mathcal{N}_{s,\boldsymbol{k}}$$
$$= \frac{e_s J_{0,\boldsymbol{k}} F_{Ms}}{T_s} \left(\frac{\partial \hat{\Phi}_{\boldsymbol{k}}}{\partial t} + i \omega_{*s}^T \hat{\Phi}_{\boldsymbol{k}} \right) + C_{s,\boldsymbol{k}}. \tag{2.4}$$

Here, $\boldsymbol{b} = B^{-1}\boldsymbol{B}$ is the unit vector along the magnetic field, $\omega_{*s}^T = \omega_{*s}[1 + \eta_s(m_s v^2/2T_s - 3/2)]$ and $\eta_s = \mathrm{d} \ln T_s/\mathrm{d} \ln n_s$. Also, $\omega_{*s} = [\sigma k_y T_s/(e_s B_{ref})]\mathrm{d} \ln n_s/\mathrm{d}x$ and $\sigma = (B_{ref}/B^2)\boldsymbol{B} \cdot \nabla x \times \nabla y = B_{ref}(\mathrm{d}x/\mathrm{d}\psi)(\mathrm{d}y/\mathrm{d}\alpha)$ is constant over the flux tube domain to leading order. A constant reference field strength is denoted B_{ref} and $C_{s,\boldsymbol{k}}$ is the gyroaveraged collision operator. The factor $J_{0,\boldsymbol{k}}$ is shorthand for the Bessel function $J_0(k_\perp v_\perp/\Omega_s)$, where $k_\perp = |k_x \nabla x + k_y \nabla y|$ and $\Omega_s = e_s B/m_s$ is the gyrofrequency. Also, the magnetic drift is

$$\boldsymbol{v}_d = \frac{m_s v_\perp^2}{2 e_s B^3} \boldsymbol{B} \times \nabla B + \frac{m_s v_{||}^2}{e_s B^2} \boldsymbol{B} \times \boldsymbol{\kappa}, \tag{2.5}$$

where $\boldsymbol{\kappa} = \boldsymbol{b} \cdot \nabla \boldsymbol{b}$ is the curvature and $\mathcal{N}_{s,\boldsymbol{k}}$ denotes the nonlinear term:

$$\mathcal{N}_{s,\boldsymbol{k}} = \sum_{\boldsymbol{k}'} \frac{\sigma J_{0,\boldsymbol{k}'} \hat{\Phi}_{\boldsymbol{k}'} \hat{h}_{s,\boldsymbol{k}''}}{B_{ref}} (k_y' k_x'' - k_x' k_y''), \tag{2.6}$$

where $\boldsymbol{k}'' = \boldsymbol{k} - \boldsymbol{k}'$. Note that where v^2 is needed in F_{Ms} and ω_{*s}^T, it can be computed from the independent variables via $v^2 = v_{||}^2 + 2\mu B$, so the magnetic geometry enters via B.

The system is closed with the quasineutrality equation:

$$\hat{\Phi}_{\boldsymbol{k}} \sum_s \frac{e_s^2 n_s}{T_s} = \sum_s e_s \int d^3 v \, J_{0,\boldsymbol{k}} \hat{h}_{s,\boldsymbol{k}}. \tag{2.7}$$

The integral over velocity space for these velocity coordinates $\int d^3 v = 2\pi \int_{-\infty}^{\infty} \mathrm{d}v_{||} \int_0^{\infty} d\mu \, B$ depends on the magnetic field via B.

Once the potential and distribution function are computed, we are principally interested in the turbulent heat flux for the ions, $s = i$:

$$Q = \sum_{\boldsymbol{k}} \left\langle \int d^3 v \, \frac{m_i v^2}{2} \hat{h}_{i,k} \frac{i k_y \sigma J_{0,k} \hat{\Phi}_{-k}}{B_{ref}} \right\rangle, \tag{2.8}$$

where $\langle \ldots \rangle$ denotes a flux surface average: $\langle u \rangle = \left(\int d\ell / B \right)^{-1} \int d\ell \, u / B$ for any quantity u.

Now, let us examine the places where the magnetic geometry enters the abovementioned model. From the $\boldsymbol{v}_d$ term in (2.4), which can be expanded as

$$(\boldsymbol{v}_d \cdot \nabla x) \, i k_x \hat{h}_{s,k} + (\boldsymbol{v}_d \cdot \nabla y) \, i k_y \hat{h}_{s,k}, \tag{2.9}$$

we see that the ∇x and ∇y components of the magnetic drift (2.5) appear. Note that $\boldsymbol{b} \times \nabla B \cdot \nabla \psi = \boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla \psi$ exactly for an MHD equilibrium, even if the ratio β of thermal to magnetic pressure is not small, so the two terms in $\boldsymbol{v}_d \cdot \nabla x$ can be combined as a multiple of $\boldsymbol{b} \times \nabla B \cdot \nabla x$. Next, in the argument of the Bessel functions,

$$k_\perp = \sqrt{k_x^2 |\nabla x|^2 + 2 k_x k_y \nabla x \cdot \nabla y + k_y^2 |\nabla y|^2}, \tag{2.10}$$

the quantities $|\nabla x|^2$, $\nabla x \cdot \nabla y$ and $|\nabla y|^2$ appear. Finally, B appears in numerous places: through Ω_s in the Bessel functions, in the flux surface average in the heat flux, in $v(v_\parallel, \mu B)$, in $\int d^3 v$ and via $\partial B / \partial z$ in (2.4). Thus, we find that the raw geometric features entering the gyrokinetic model are the following seven functions of z:

$$B, \quad \frac{\boldsymbol{B} \times \nabla B \cdot \nabla x}{B^3}, \quad \frac{\boldsymbol{B} \times \nabla B \cdot \nabla y}{B^3}, \quad \frac{\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y}{B^2}, \quad |\nabla x|^2, \quad \nabla x \cdot \nabla y, \quad |\nabla y|^2. \tag{2.11}$$

As stated previously, the x and y variation of these equilibrium quantities over the flux tube simulation domain is negligible due to the small extent of the domain in these coordinates, so we only need the variation of these quantities in z along a field line. If a parallel coordinate z other than arclength is used, there would be an additional geometric input $\boldsymbol{b} \cdot \nabla z$ required in the parallel streaming term of (2.4).

We can now understand a translation-invariance property of the abovementioned gyrokinetic model, which ML models should preserve. The key idea is that z does not appear explicitly anywhere in the model – z-dependence in the equations enters only through the functions in (2.11). To precisely state the translation-invariance property, suppose the gyrokinetic-quasineutrality system is solved with periodic boundary conditions in z. Then, the heat flux is exactly unchanged if a periodic shift $f(z) \to f(z + \Delta)$ is applied, where f indicates each of the seven raw features (2.11) along with Φ and h_i, where the shift Δ is the same for all quantities. If the seven raw features are shifted in this way but the initial conditions for Φ and h_i are not, the detailed dynamics will change, but under the usual assumption that the mean heat flux is independent of the initial condition, the mean heat flux would still be invariant. We confirmed that this translation-invariance indeed holds for nonlinear GX simulations, as shown in figure 1. If a twist-and-shift boundary condition (Beer, Cowley & Hammett 1995; Martin *et al.* 2018) is used in z instead of periodic boundary conditions, and a finite number of k_x and k_y modes are included, the translation-invariance will be slightly violated, because the outgoing distribution
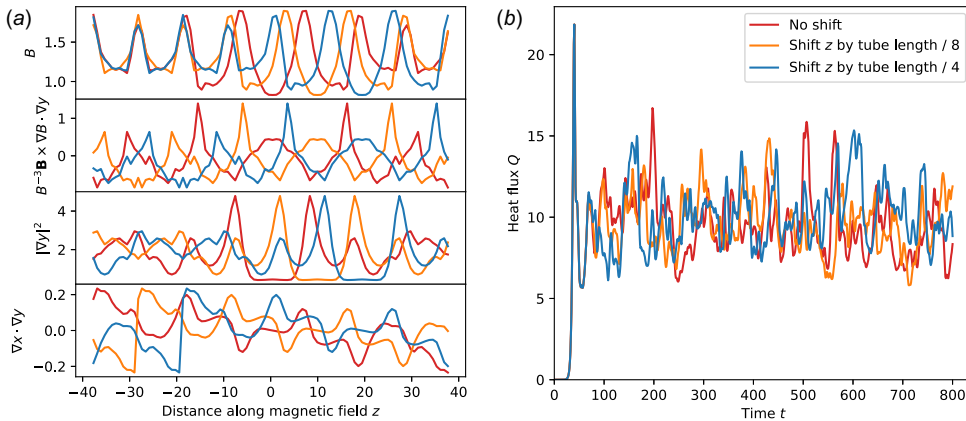
FIGURE 1. Translation-invariance of the gyrokinetic-quasineutrality system. (*a*) A periodic translation in z is applied to all of the z-dependent inputs to the gyrokinetic system, (2.11). Only four of the seven are shown for simplicity, but all are translated. (*b*) Average heat flux is unchanged by the translation.

function is set to zero at the unlinked flux tube ends. However, if the simulation is well resolved with respect to the number of k_x modes, the heat flux is evidently approximately independent of the number of times a tube segment links to itself and hence approximately independent of where the outgoing distribution function is set to zero. Thus, the breaking of translation-invariance should be small.

For all results that follow, the raw features (2.11) are normalised by a reference length L_{ref} and reference field strength B_{ref}. The effective minor radius a is adopted for L_{ref} and B_{ref} is defined by equating $\pi a^2 B_{ref} = 2\pi |\psi_a|$, where ψ_a is the value of ψ at the equilibrium boundary. We choose the perpendicular coordinates to be $x = a\sqrt{\psi/\psi_a}$ and $y = -\alpha x \, \text{sign}(\psi_a)$, so $\sigma = -1$.

3. Dataset generation

In this section, we first present the method used to generate stellarator MHD equilibria with diverse geometries. Next, details of the nonlinear gyrokinetic simulations in these equilibria are given. After describing some general properties of the dataset, a few configurations are highlighted that are interesting due to their extreme values of the heat flux.

3.1. *Magnetohydrodynamic equilibria*

An MHD equilibrium is determined by the boundary shape together with two functions of the toroidal flux, typically the pressure and enclosed toroidal current (Kruskal & Kulsrud 1958). It is not obvious how to best sample this space, particularly the space of boundary shapes. Boundaries with self-intersections or other such pathologies should be avoided. Also, a compromise must be struck between preserving similarity of the shapes to 'real' devices (built experiments or theoretical configurations designed through serious optimisation) while also allowing for new possible geometries.

To balance these considerations, we assemble a set of 23577 equilibria drawn from three groups. The first group has heliotron-like rotating ellipse shapes, in which parameters of the shape (number of field periods, aspect ratio, elongation,
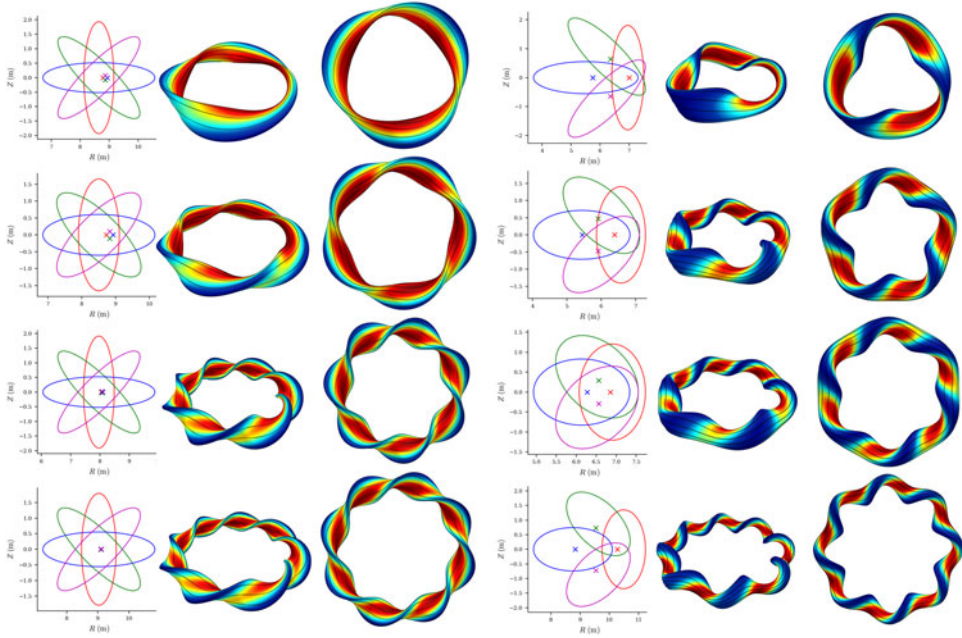
FIGURE 2. Examples of the rotating-ellipse equilibria included in the dataset. Two-dimensional plots show the cross-sections at which the toroidal angle is 0, 1/4, 1/2 and 3/4 of a field period. Three-dimensional images show each configuration from two angles, with colour indicating $|B|$ (red = high, blue = low) and field lines in black. Left columns, configurations in which the boundaries are centred on a circle. Right columns, configurations in which the boundaries are centred on a curve with torsion.

axis torsion, beta) are sampled randomly. Equilibria in the second group are taken from the QUASR database of quasi-axisymmetric (QA) and quasi-helically symmetric (QH) configurations (Giuliani 2024; Giuliani, Rodríguez & Spivak 2024). We use both the original QUASR configurations, which are all vacuum fields, and also generate new configurations by adding pressure while keeping the plasma boundary shape fixed. For the third group, random boundary shapes are generated by sampling Fourier modes that have been fit to a dataset of previous stellarator shapes. The combined set of configurations includes values of aspect ratio ranging from 2.9 to 10, volume-averaged beta from 0 % to 5 %, and number of field periods from 2 to 8. Thus, the set of equilibria is diverse, and includes both omnigenous and non-omnigenous geometries. All equilibria have the same minor radius and same boundary toroidal flux, resulting in the same normalising length and normalising field strength, so the gyro-Bohm normalisations are identical. Examples of equilibria from the three classes are shown in figures 2–4. More details of the procedures for generating equilibria are given in Appendix A.

3.2. *Gyrokinetic simulations*

In each equilibrium, we extract at least four flux tubes, more in the QUASR vacuum configurations so as to include more omnigenous geometries. The radial location $\rho = \sqrt{s}$ for each flux tube is sampled randomly from $[0, 1]$, where $s = \psi/\psi_a$ is the normalised toroidal flux. All flux tubes are stellarator-symmetric, with at least
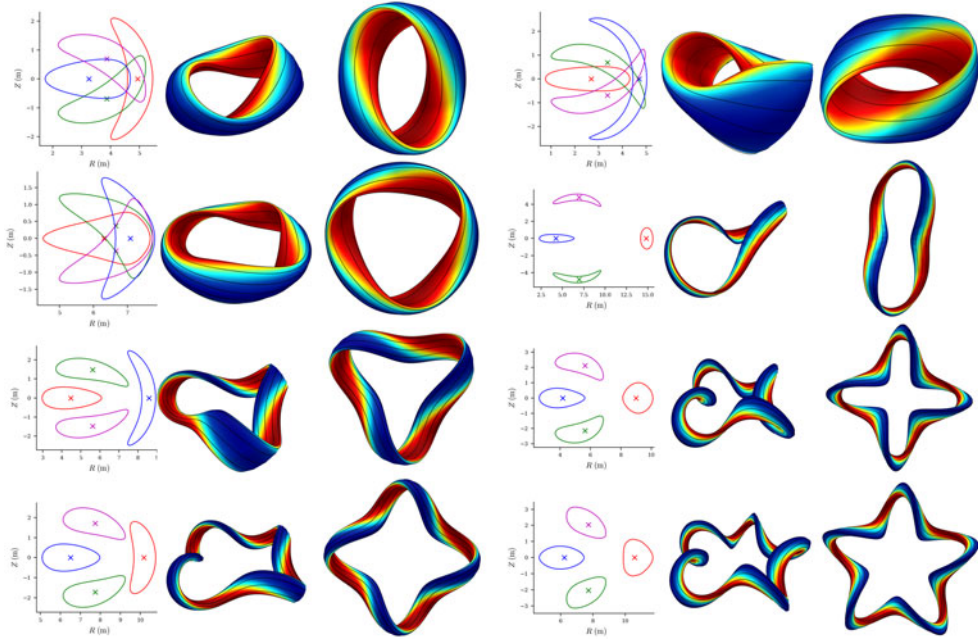
FIGURE 3. Examples of the QUASR quasi-axisymmetric and quasi-helically symmetric equilibria included in the dataset. Two-dimensional plots show the cross-sections at which the toroidal angle is 0, 1/4, 1/2 and 3/4 of a field period. Three-dimensional images show each configuration from two angles, with colour indicating $|B|$ (red = high, blue = low) and field lines in black.

one flux tube in every configuration centred on each of the points $(\theta, \phi) = (0, 0)$, $(\pi, 0)$, $(0, \pi/n_{\text{fp}})$ and $(\pi, \pi/n_{\text{fp}})$. Nothing in the procedure limits the method to stellarator-symmetric equilibria or flux tubes. The total number of flux tube geometries in the dataset is 100 705. The length of the flux tubes in real space was chosen to be the same for every configuration, 75 times the minor radius, which for aspect ratio 6 corresponds to roughly two full toroidal transits. Due to the variation in ι, the number of poloidal transits varied between configurations. By making all flux tubes have the same physical length, distances in the z coordinate can be compared meaningfully between flux tubes, and for each integer j, the jth Fourier series coefficients in all flux tubes correspond to the same physical $k_\parallel$. These properties may be advantageous for extracting patterns in the dataset.

Nonlinear electrostatic simulations with adiabatic electrons were then carried out with the GX code (Mandell *et al.* 2024). We run GX twice for each flux tube, resulting in two datasets each of size $N = 100\,705$. In the first dataset, the temperature and density gradients were fixed to the same values for all flux tubes. In the second dataset, these gradients were chosen randomly for each flux tube. Using the latter dataset, models can potentially find interactions between geometry and gradients in their effect on the heat flux; for example, the dependence on geometry may differ close to the critical gradient compared with far above the critical gradient. However, we also assembled the dataset with fixed gradients since it enables several of the analyses in § 5 and allows us to focus cleanly on the effects of geometry. For the fixed-gradient dataset, we take $a/L_{Ti} = 3$ and $a/L_n = 0.9$, where $a/L_{Ti} = -(a/T_i)\,dT_i/dx$ and $a/L_n = -(a/n)\,dn/dx$ are the normalised gradients of
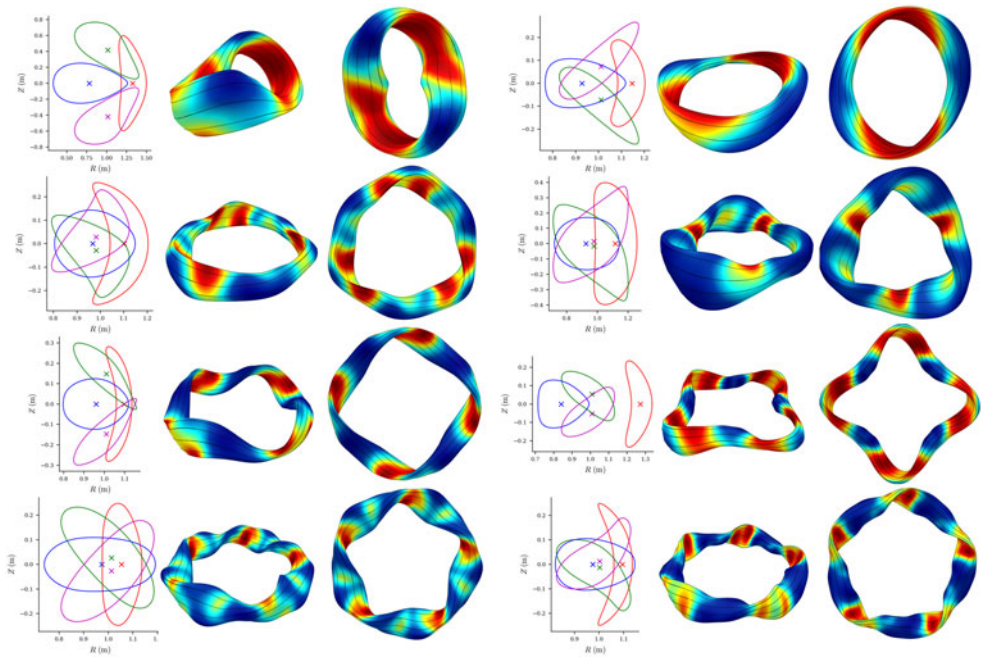
FIGURE 4. Examples of the equilibria generated with random boundary Fourier modes. Two-dimensional plots show the cross-sections at which the toroidal angle is 0, 1/4, 1/2 and 3/4 of a field period. Three-dimensional images show each configuration from two angles, with colour indicating $|B|$ (red = high, blue = low) and field lines in black.

ion temperature and density, reflecting typical measurements at the $s = 0.5$ surface of W7-X (Beurskens *et al.* 2021; Lunsford *et al.* 2021; Zhang *et al.* 2023). For the varying-gradient dataset, each simulation has a/L_{Ti} and a/L_n sampled randomly within plausible experimental ranges.

Additional details of the turbulence simulations are given in Appendix B.

3.3. *Configurations with very low or high heat flux*

The distribution of heat fluxes for the dataset is shown in figure 5. It is striking that even for the dataset with fixed gradients, Q varies by over four orders of magnitude over the data. This variation is evidently due purely to the geometry.

It is interesting to examine configurations with particularly low or particularly high heat flux. Several such flux tubes are displayed in figure 6. The columns represent six flux tubes from the dataset. In the fixed gradient dataset with $a/L_T = 3$ and $a/L_n = 0.9$, the first three tubes are stable, while the other three have $Q > 500$, even though the gradients are identical. These flux tubes are all taken from $n_{fp} = 3$ equilibria from the group with random boundary Fourier modes. The first six of the seven rows show the geometric inputs to the gyrokinetic-quasineutrality system. To simplify the figure, $B^{-3}\boldsymbol{B} \times \nabla B \cdot \nabla y$ is not shown since it is similar to $B^{-2}\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y$. The bottom row shows the contribution to the heat flux Q as a function of z. For each of the seven rows, the vertical scales are the same for each column to allow comparison.

Several patterns are apparent. The configurations with highest heat flux have regions with very large $|\nabla x|^2$. The high-heat-flux tubes also have larger magnitudes
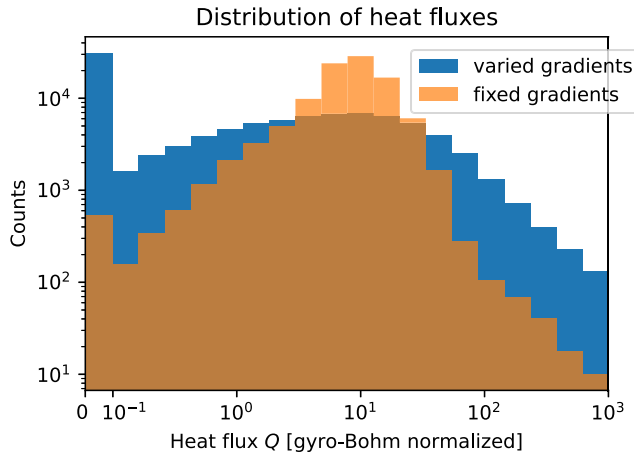
FIGURE 5. Distribution of heat fluxes for the fixed-gradient and varied-gradient datasets. In the latter, 30 % of the simulations were stable, with $Q \approx 0$. Simulations with $Q < 0.1$, considered stable, are included in the leftmost bar.
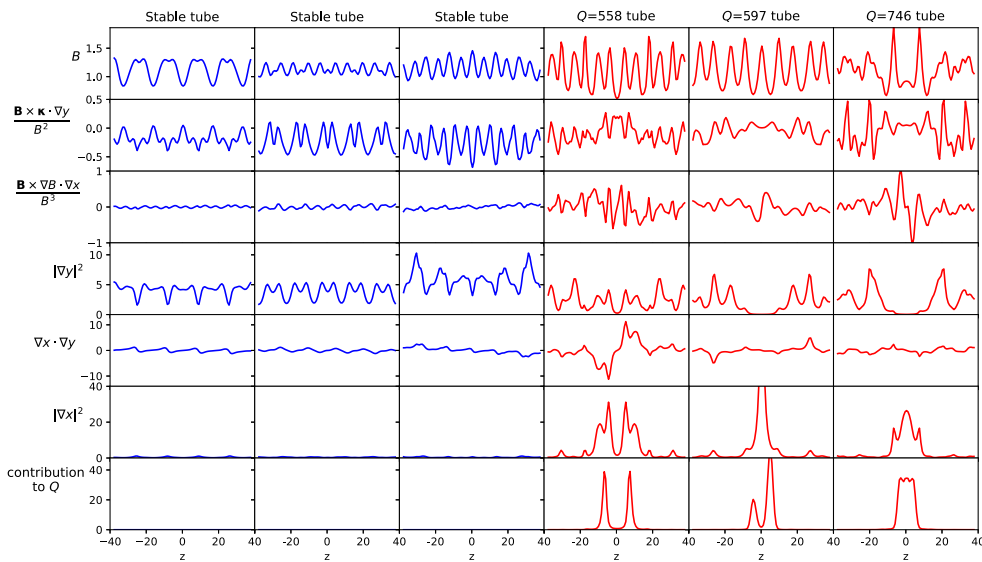

FIGURE 6. Some trends are apparent between flux tubes with very low or high heat flux. The columns show six flux tubes from the $n_{fp} = 3$ equilibria with random boundary Fourier modes, the first three stable, the last three with very high Q at the same gradients. The top six rows are the inputs to the gyrokinetic-quasineutrality system ($B^{-3}\boldsymbol{B} \times \nabla B \cdot \nabla y$ is omitted for simplicity since it is similar to $B^{-2}\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y$), while the bottom row shows the contribution to the total heat flux versus z.

of $B^{-3}\boldsymbol{B} \times \nabla B \cdot \nabla x$. The stable tubes have more negative values of $B^{-2}\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y$, meaning mostly good curvature. All of these patterns foreshadow findings in § 5.

For the geometry-gradient pairings with very large heat flux, where the gyro-Bohm-normalised Q is not small compared with $1/\rho_*$, with ρ_* the ratio of typical

gyroradius to macroscopic scale, the ordering used to derive the gyrokinetic model may no longer be accurate. Therefore, these simulation runs may not accurately describe real physical scenarios. Nonetheless, it is still reasonable to consider the δf-gyrokinetic model as an abstract mathematical function, a map from seven functions (the geometry) and two real numbers (the gradients) to a real number (Q), and to develop approximations of this function. Moreover, since surrogates will be more accurate at interpolation than extrapolation, it is desirable to include at least a few of these extreme data points in the dataset.

3.4. *Weighting of errors for regression*

When fitting regression models to the heat flux data and assessing their accuracy, some care is required to measure and weigh errors appropriately. The first issue is that the heat flux can span several orders of magnitude. If a model makes an absolute prediction error of 1 in gyro-Bohm units, this error is significant if the true heat flux is 0.5, but it is a relatively minor error if the true heat flux is 100. Moreover, the standard deviation (over time) of the heat flux within one simulation tends to scale approximately linearly with the mean heat flux, so high-Q cases have higher uncertainty than low-Q cases. For these reasons, it is natural to measure errors in a relative rather than absolute sense. Weighing errors in this relative sense can be achieved by performing regression on $\ln Q$ rather than Q directly. In particular, models should be scored based on the mean squared error or coefficient of determination R^2 evaluated using $\ln Q$. Thus, in the standard definition, $R^2 = 1 - \left[\sum_j (y_j - f_j)^2 \right] / \left[\sum_j (y_j - \bar{y})^2 \right]$, where y_j are the true target values, $\bar{y}$ is their mean and f_j are the predicted values, one would take y_j and f_j to be the true and predicted values of $\ln Q$ rather than of Q itself.

However, an additional important consideration is that some simulations are stable. In a stable simulation, the heat flux is typically slightly positive but $\ll 1$ at the end of the computation due to decay of the initial perturbation and the exact value of Q is not very meaningful. There is not a steady state to average over and the final Q is influenced by the magnitude of the initial condition. If the heat flux from a simulation is $Q = 10^{-12}$ and a regression model predicts it to be 10^{-4}, the model is accurately predicting that the system is stable (or nearly so), but scoring the model based on the error in $\ln Q$ would treat the prediction error to be large. In some stable simulations where $|Q|$ has very small magnitude, Q may even be slightly negative, in which case, $\ln Q$ cannot even be evaluated.

There are several possible approaches to account for these considerations. One approach is to consider a classification problem for stability versus instability separately from the regression problem. In this case, the regression only needs to be evaluated if the classifier has first predicted instability, i.e. in the unstable region of parameter space, so the regression can be fit using only the subset of the data in which Q exceeds some small threshold. Here, we adopt this approach with a threshold of 0.1 and our experience is that results are insensitive to the exact choice of threshold. With this approach, the training set does not include any points with $Q < 0.1$ or negative Q, so it is straightforward to use $\ln Q$ as the regression target. The full dataset is still used for training the classifier. If the classifier is interpretable, it directly provides insight into the factors that determine the critical gradient.

Several other approaches are possible. As reported by van de Plassche *et al.* ([2020](#)), a regression could be fit by minimising a custom objective that penalises

errors in Q (or $\ln Q$) only in the unstable region, while in the stable region, only penalising positive values of Q, to prevent spurious prediction of instability. However, not all regression methods allow for a custom objective, e.g. nearest-neighbours. Or, regression could be performed on the full dataset with a standard objective after all values of Q below a threshold are replaced by the threshold. While straightforward to implement, this method introduces non-smoothness in the predicted Q; this is not a problem for decision tree models, but it may introduce inaccuracies with smooth models like neural networks. A fourth option is to perform regression on the full dataset using the target quantity

$$\hat{Q} = \ln(1 + Q). \tag{3.1}$$

For $Q \gg 1$, the quantity $\hat{Q}$ behaves like $\ln Q$, while for $Q \ll 1$, $\hat{Q}$ varies approximately linearly with Q. Therefore, for a regression model that is fit by minimising the mean squared error in $\hat{Q}$, or R^2 computed from $\hat{Q}$, prediction errors will effectively be based on relative error in Q for large Q, but based on absolute error for small Q, as desired. We have tried the first, third and fourth of these methods, and our experience is that they give very similar results. For the rest of the discussion, we use the first method.

4. Convolutional neural network models

We develop neural network surrogates to predict heat flux values from the simulations described previously. Our approach involves two key strategies: (i) designing a neural network that remains invariant to cyclic permutations and (ii) ensuring reliable predictions through an ensemble of the top-k models selected via hyperparameter optimisation.

4.1. *Cyclic invariant neural network*

Our goal is to develop a surrogate model based on a neural network that maintains invariance to cyclic permutations inherent to the gyrokinetic system. An analogy exists with computer vision research, where translation invariance is important as well: if a cat is present in an image, it should be recognised as a cat even if its location within a bitmap is shifted. A standard approach to achieving approximate translation-invariance in computer vision is through convolutional (as opposed to fully connected) neural networks. This type of neural network contains convolutional layers, in which the spatial data are convolved with a kernel, which is a translation-equivariant operation. Between convolutional layers, pooling layers then reduce the spatial resolution, typically by a factor of two, and a sufficient number of alternating convolutional and pooling layers converts the spatial data to approximately translation-invariant features. While computer vision applications use two-dimensional convolution and pooling operations, for the turbulence application here, we will use one-dimensional convolution and pooling.

During training, we further reinforce the translation-invariance by augmenting the dataset with a sufficient number of sequences that are randomly permuted (cyclically) in z, allowing the network to learn patterns that are independent of any particular alignment with respect to z. This strategy prevents the model from developing biases towards specific alignments, ultimately improving generalisation and stability.

Another key aspect of our surrogate model development is the use of ensemble learning, a powerful machine learning technique that aggregates predictions from multiple base models to improve overall performance. Ensemble methods have proven to be highly effective in reducing prediction variance, mitigating model biases and improving generalisation, making them well suited for applications involving complex, high-dimensional data such as ours. By leveraging an ensemble of multiple models, we aim to construct a predictive framework that is not only more accurate, but also more stable when encountering unseen data. The core advantage of ensemble learning is its ability to balance trade-offs between different models – some models may be more sensitive to specific features, while others may capture different patterns. By combining them, we create a more comprehensive representation of the underlying physical system.

To implement this strategy, we design our base model to be highly flexible, allowing variations in neuron count, layer depth, kernel shape, activation functions, etc. also known as hyperparameters. This flexibility is essential for tuning the model to achieve optimal predictive performance across different configurations. However, selecting the best architecture requires careful exploration of the hyperparameter space. To efficiently search for the best hyperparameters, we employ DeepHyper (Balaprakash *et al*. 2018; Egelé *et al*. 2023), an advanced hyperparameter optimisation tool designed for deep learning. DeepHyper allows us to automate the search for high-performing configurations, reducing manual tuning efforts and accelerating model development. We will discuss the details of this optimisation process in subsequent sections.

With these considerations in mind, we design our surrogate neural network using a structured approach that consists of three main components, shown in figure 7.

(i) **Feature extraction with convolutional layers.** The first component comprises n consecutive layers of neural network blocks, each containing one-dimensional convolutional layers (Conv1D). These layers are responsible for extracting spatial features from the input sequences while preserving cyclic invariance. Each convolutional layer processes x data channels and outputs y number of transformed channels, which are subsequently normalised using batch normalisation to stabilise training and improve convergence. A max-pooling layer follows each block to downsample the extracted features, reducing computational complexity while retaining the most relevant information.

(ii) **Global average pooling (GAP) for dimensionality reduction.** To further condense the extracted feature maps, we integrate GAP. Unlike traditional pooling methods that operate on small local regions, GAP computes the average activation value for each feature map across all spatial dimensions. This not only reduces the number of parameters, but also ensures that the model focuses on high-level feature representations rather than localised details, making it particularly effective for cyclic data.

(iii) **Fully connected layers for prediction.** The final stage consists of m fully connected linear layers that process the compact feature representation and map it to the target variable: predicted heat flux averages. These layers refine the extracted features and produce the final output by capturing complex, nonlinear relationships in the data. Since the temperature and density gradient are not functions of z, these two inputs bypass the convolutional layers and feed directly to the first fully connected layer.

Hyperparameter	Count	Range	Description
kernel_size	5	(2, 16)	Conv1D's convolving kernel size
conv_channel	5	(4, 32)	Conv1D's number of output channels
fc_dimension	2	(4, 32)	Fully connected layer's output feature size
batch_size	1	(16, 64)	Batch size
learning_rate	1	(1e-5, 1e-1)	Learning rate (sampled in log-uniform)
scheduler_patience	1	(5, 20)	ReduceLROnPlateau scheduler's patience parameter
scheduler_factor	1	(0.1, 1.0)	ReduceLROnPlateau scheduler's factor parameter

TABLE 1. List of hyperparameters and their search ranges explored using DeepHyper.
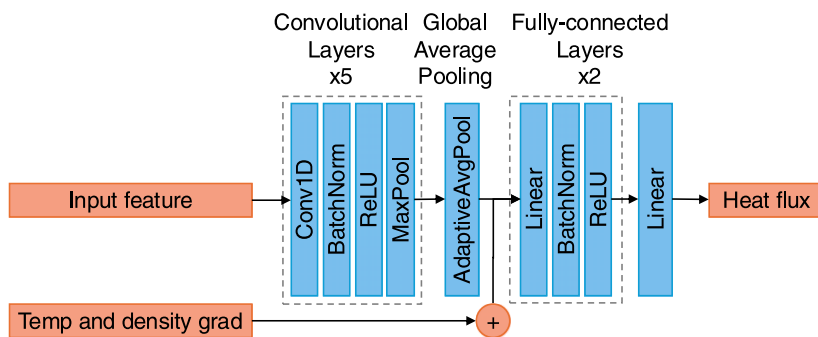

FIGURE 7. Surrogate model architecture to learn heat flux averages using a structured neural network. It consists of three main components: (*a*) feature extraction with Conv1D; (*b*) global average pooling; and (*c*) fully connected layers.

4.2. *Hyper parameter optimisations and ensemble learning*

To enhance the predictive performance of our surrogate model, we employ ensemble learning, which aggregates predictions from multiple base models. By combining different models, we create a more robust and comprehensive representation of the underlying physical system, improving accuracy and generalisation.

In our model, we carefully parametrise several key hyperparameters that influence performance. These include the kernel size of the 1-D convolution layers, the number of output channels in Conv1D layers, the number of neurons in fully connected layers, batch size, learning rate and scheduler patience for dynamically adjusting the learning rate. A summary of these hyperparameters is provided in table 1. We fix the number of Conv1D layers at five, so the associated pooling layers reduce the number of spatial points from 96 to $96/2^5 = 3$, and include two fully connected layers. While

these constraints simplify the search space, the total number of possible hyperparameter configurations remains extremely large, exceeding one hundred quadrillion (10^{17}). Exploring this vast space through brute-force methods is computationally infeasible.

To efficiently navigate this large hyperparameter space, we leverage DeepHyper (Balaprakash *et al*. 2018; Egelé *et al*. 2023), an advanced hyperparameter optimisation tool. DeepHyper systematically finds the best set of hyperparameters that maximise model accuracy using different search strategies including Bayesian optimisation, evolutionary strategies and random search. Bayesian optimisation uses a probabilistic approach to refining hyperparameters, while evolutionary strategies optimise parameters through genetic algorithms. Random search serves as a baseline method. DeepHyper is designed for parallel execution in distributed computing environments such as clusters and supercomputers, making it effective for hyperparameter optimisation tasks at scale. In DeepHyper, Bayesian optimisation (BO) that we adopt follows a single-manager, multiple-worker scheme to explore and exploit the search space. The manager maintains a probabilistic surrogate model (e.g. Random Forest) to predict the performance of configurations and select promising candidates based on an acquisition function (e.g. expected improvement or upper confidence bound). The workers asynchronously evaluate these candidates in parallel, returning results to update the surrogate model iteratively. To improve robustness and generalisation, DeepHyper generates an ensemble of models by aggregating the top-performing configurations identified during the search.

In this study, we employ DeepHyper to explore the hyperparameter space of our base model, aiming to identify and select the top-N best models for ensemble learning. DeepHyper systematically searches the hyperparameter space to optimise model performance. To guide this optimisation, we use the coefficient of determination R^2, as the objective function, ensuring the selected models achieve the highest predictive accuracy. The base neural network model is implemented using PyTorch.

Figure 8 presents the results of our hyperparameter search. Each point in the figure represents a model explored during the optimisation process. The x-axis denotes the time of completion, while the y-axis shows the model's performance measured by the R^2 score. Over a period of approximately nine hours, we used 64 GPUs to conduct this search. DeepHyper evaluated 443 different models with varying hyperparameter configurations, systematically assessing their performance. After completing the search, we selected the top 100 models with the highest R^2 scores to form an ensemble, ensuring improved predictive accuracy and generalisation. The bottom plot presents a histogram of model sizes, represented by the number of parameters on the x-axis, for all 443 models explored by DeepHyper and the top 100 high-performing selected models, illustrating the distribution of model sizes before and after selection.

Figure 9 illustrates the prediction performance of this ensemble, which consists of the 100 highest-performing models from the DeepHyper optimisation, on the varied-gradient data. Each red dot in the figure represents the mean prediction of the ensemble at each target value in the test dataset, while the vertical bars indicate the range of $\pm 1\sigma$ (standard deviation) of predictions across the 100 models. We evaluated the ensemble using a total of 9785 held-out test samples, achieving an overall R^2 score of 0.989, computed by treating the ensemble mean as the prediction. Additionally, the neural networks exhibited efficient inference speed, taking a total 1031.0 s on a single NVIDIA A100 GPU to process 9785 test samples across the 100 models. To make a single prediction of the heat flux, this time corresponds to
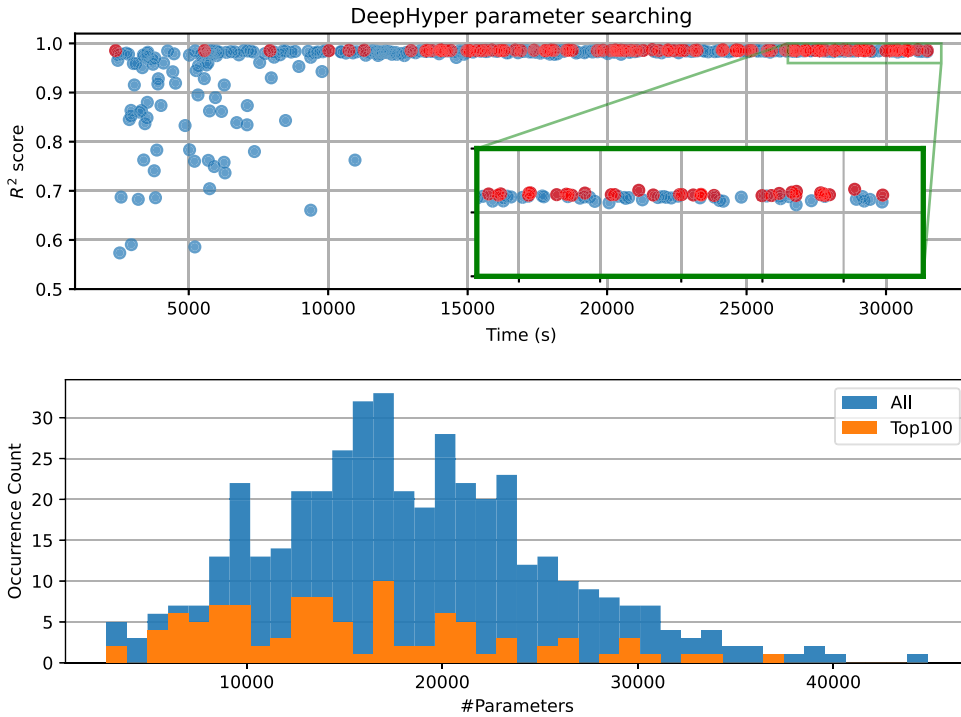
FIGURE 8. Results of the DeepHyper hyperparameter search. Each point represents a model with a unique hyperparameter configuration explored by DeepHyper, plotted against its completion time (x-axis) and performance score (y-axis). The search was conducted using 64 GPUs over approximately nine hours, evaluating a total of 443 models. The top 100 highest-performing models, selected for the final ensemble, are highlighted in red. The bottom plot shows the histogram of model sizes (number of parameters on the x-axis) for all 443 models explored by DeepHyper and the top 100 selected models.

an average of 0.001 s for a single CNN and 0.1 s for the ensemble. These times represent a factor of $\sim$4000 speed-up for evaluating the ensemble compared with one of the gyrokinetic simulations used to generate the training data; for evaluating a single CNN, the speed-up factor over a gyrokinetic simulation is 4×10^5. These results demonstrate the effectiveness of our ensemble-based approach in providing accurate and stable predictions across the dataset while maintaining competitive inference performance.

5. Manual feature engineering

Complementary to the neural network approach discussed in the previous section, we next present a method to identify patterns in the data using different machine learning algorithms together with manual feature extraction. While these methods are not yet able to achieve as close a fit to the data as the neural networks, they naturally provide some level of interpretability. We proceed by first defining a combinatorial procedure to compute a large number of derived features from the raw data that respect the spatial translation invariance. The Spearman correlation between each feature and the heat flux provides a first measure of each feature's importance. Then we apply FSFS (Efroymson 1960; Whitney 1971), where the order in which
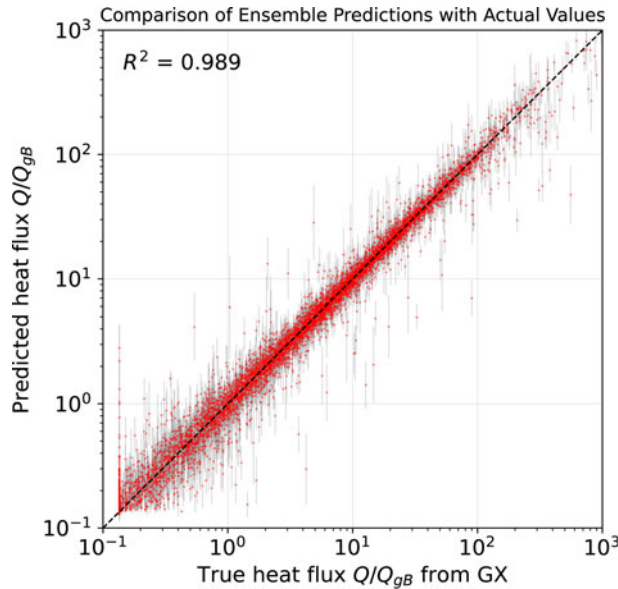
FIGURE 9. Prediction performance of an ensemble of the top 100 models selected from DeepHyper for the varied-gradient dataset. After exploring 443 models by DeepHyper, the top 100 were chosen based on their performance and evaluated against a test dataset of 9785 samples. Each dot represents the mean prediction of the ensemble, while the vertical bars indicate ± 1 standard deviation. The ensemble achieved an overall R^2 score of 0.989, demonstrating strong predictive accuracy and stability.

features are selected provides a second measure of feature importance. FSFS can be applied using any regression method, and several methods will be compared for the fixed-gradient and varied-gradient data, with the conclusion that there is general agreement. Finally, a third measure of feature importance is provided by computing Shapley values for the models that result from FSFS. We will find broad consistency among the three measures of importance.

5.1. *Feature extraction*

Since the heat flux from the gyrokinetic model is invariant under periodic translation in z, as discussed in § 2, a regression or classification model should preserve this invariance. However, this invariance would generally be broken if we use the raw features on grid points as inputs to regression or classification methods (other than CNNs) directly. For example, a regression could, in principle, find that the heat flux depends specifically on $|\nabla x|^2$ at $z = 0.3$, which is not a translation-invariant quantity. Instead, the inputs to the regression or classification should be translation-invariant features that are derived from the raw z-dependent geometric features. To ensure this invariance in the extracted features, we consider features that are the composition of equivariant operations with translation-invariant reductions. To explain this method, let P denote the set of periodic real-valued functions on the interval $[0, 1)$, representing the z-domain. Each of the seven original raw features can be considered as an element of the set P. An equivariant operation E is a function from $P \to P$, such that translating the input and then applying E gives the same result as first applying E and then translating. More formally, let f be any element of P and

let $g : P \to P$ indicate a periodic shift by some distance Δ, so $g(f)(x) = f(x + \Delta$ mod 1). Then the equivariance of E means that $E(g(f)) = g(E(f))$ for all f and Δ. A reduction is a map from $P \to \mathbb{R}$. A reduction A is translation-invariant if $A(f(x + \Delta \mod 1)) = A(f(x))$ for all x and any shift Δ.

For the results here, we consider 23 possible translation-invariant reductions: max; min; range (max−min); mean; median; mean of the square; variance; skewness; L_1 norm; quantiles 0.1, 0.25, 0.75 and 0.9; count above -2, -1, 0, 1 and 2; the absolute amplitudes of the three longest-wavelength Fourier modes; the expected $k_{||}$; and the largest amplitude $k_{||}$. The last two quantities are defined as follows: letting $\sum_{j=-N}^{N} \hat{f}_j \exp(ik_j z)$, denote the Fourier transform of a raw feature $f(z)$, the expected $k_{||}$ is $\sum_{j=1}^{N} k_j |\hat{f}_j| / \sum_{j=1}^{N} |\hat{f}_j|$, while the largest amplitude $k_{||}$ is k_j, where j is the solution of $\arg\max_j |\hat{f}_j|$. These last two reductions were included due to the importance of parallel length scales in critical-balance models of turbulence (Barnes, Parra & Schekochihin 2011). We indicate the set of these 23 reductions by R.

For the equivariant operations, we consider combinations of unary operations and products. The unary operations, each an equivariant map $P \to P$, include 11 operations on an input function f: the identity f, $|f|$, df/dz, Heaviside(f), Heaviside($-f$), ReLU(f) = max(f, 0), ReLU($-f$), $1/f$, f^2, fB and f/B. The last two operations, multiplying or dividing by the field strength B, were included since the coordinate Jacobian is $1/B$.

We can now state the full set of features that are considered. Let F denote a set of eight features, given by the original seven raw features together with the local shear $S = (d/dz)(\nabla x \cdot \nabla y / \nabla x \cdot \nabla x)$. The local shear (Greene & Johnson 1968) is added here since it is plausible that it could play a role in determining the turbulence intensity. Then, the set of features $U(F)$ is obtained by applying each possible unary operation to each of the eight elements of F. We let $C(U(F))$ denote the set of all possible pairwise products of functions in $U(F)$, supplemented with $U(F)$. Then, $U(C(U(F)))$ indicates all possible unary operations applied to all elements of $C(U(F))$. The full set of features is finally $R(U(C(U(F))))$, obtained by applying all of the reductions, for a total of just over 1 million extracted features. This set includes a few features that are duplicated, improper due to division by 0, or that include unnecessary operations (e.g. ReLU of a positive-definite function), but a suitably large fraction are distinct and finite. Of course, larger feature sets could be considered by including more operations and combinations, but the set here is suitably rich to find accurate regressions of the data.

5.2. *Spearman correlation*

One approach to measuring the potential importance of an isolated feature in a regression problem is to compute its Spearman correlation with the target quantity. Spearman correlation is defined by the Pearson correlation between the sorted rank of the target with the sorted rank of the feature. The absolute magnitude of the Spearman correlation has the appealing properties of being fast to compute and invariant to any monotonic nonlinear function, e.g. for any sequence of points $\{x_j\}$, the Spearman correlation of $\{x_j\}$ with $\{\exp(x_j)\}$ is 1. We can evaluate the Spearman correlation between each feature from the previous subsection and the heat flux. Note that no regression or classification model is used. The most highly correlated features for the fixed-gradient dataset are listed in table 2. All these top features include the factor $\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y)$, where Θ is the Heaviside function. In our

Feature	Correlation
Variance$(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^2/B)$	0.775
Mean$(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^8/B^2)$	0.774
Mean$(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^4/B)$	0.772
Variance$(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^4/B)$	0.769
Mean$(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^4/B^2)$	0.769

TABLE 2. Geometric features from § 5.1 with highest absolute magnitude of Spearman correlation to the nonlinear heat flux at fixed temperature and density gradient. Here, Θ denotes the Heaviside function.

sign convention, this quantity is 1 in regions of bad curvature. The most correlated features all have a similar form, indicating the heat flux is highest when the flux surface compression $|\boldsymbol{\nabla} x|$ is large in regions of bad curvature. For the features listed, the sign of the correlations is positive, so greater flux surface compression is associated with a higher heat flux, as expected physically. The features also include an inverse weighting with B, which perhaps reflects the Jacobian $\propto 1/B$. The exact powers of $|\boldsymbol{\nabla} x|$ and B vary among the top features, but a consistent pattern is evident.

As mentioned in § 1, these features with highest Spearman correlation to the heat flux are consistent with ideas by Mynick *et al*. (2010), Xanthopoulos *et al*. (2014), Stroteich *et al*. (2022) and Goodman *et al*. (2024). In these earlier works, stellarator configurations were sought with smaller $|\boldsymbol{\nabla} x|$ in regions of bad curvature, motivated by the following physical intuition. At fixed $\mathrm{d}T/\mathrm{d}x$, reducing $|\boldsymbol{\nabla} x|$ reduces the real-space temperature gradient $|\boldsymbol{\nabla} T| = (\mathrm{d}T/\mathrm{d}x)|\boldsymbol{\nabla} x|$. Reducing $|\boldsymbol{\nabla} T|$ reduces the source of free energy for instabilities and the associated turbulence. This is particularly true if done in the regions where instabilities and turbulence are localised due to bad curvature. Positive correlation between $|\boldsymbol{\nabla} x|$ and quasilinear heat flux estimates was also noted by Jorge *et al*. (2024).

Kendall's τ is another correlation coefficient with many similarities to Spearman's correlation. For all results in this paper involving Spearman correlation, essentially the same findings are obtained in comparable computational time if Kendall correlation is used instead, though numerical values of Kendall correlation are smaller.

Spearman and Kendall correlation cannot account for dependence of the target on multiple features that may interact, so we now proceed to more sophisticated methods.

5.3. *Sequential feature selection*

In FSFS (Efroymson 1960; Whitney 1971), regression or classification is first performed using one feature at a time. In other words, if there are n features, we fit n distinct models, each with one feature. Any regression or classification method can be used. The feature that yields the best model fit to the data is selected to progress to the next step. Then, the data are fit using $n-1$ independent models that each use two features, the one selected in the first step plus all possible second features. Of these $n-1$ models, the one with closest fit is selected to progress to the next step and so forth. Thus, FSFS results in a parsimonious set of features, which
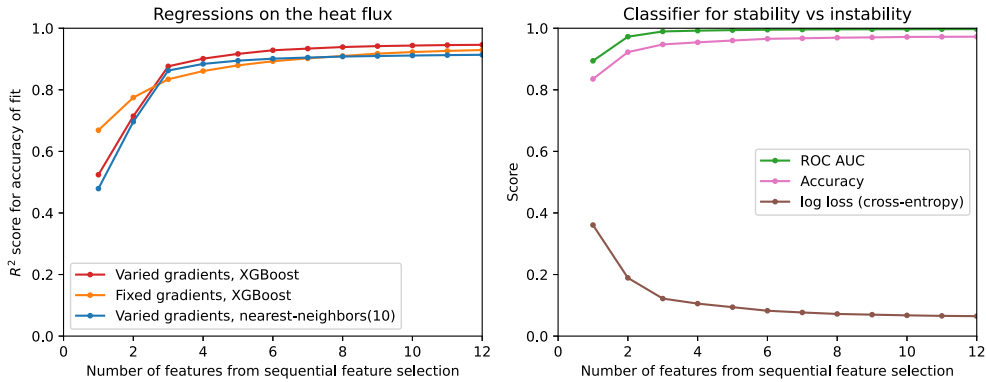
FIGURE 10. Scores for regression (left) and classification (right) in forward sequential feature selection, showing improvement as the first few features are added. Each point shows the mean score on held-out data using five-fold cross-validation.

are effectively ranked from most to least important by the order in which they are selected.

For the regression and classification models to use in FSFS, we primarily rely on the gradient boosted decision tree package XGBoost (Chen & Guestrin 2016). This choice is due to its speed with datasets of this size. Speed is a priority since $O(10^6)$ independent models must be fit at each step. Like other decision tree methods, XGBoost fits a piecewise-constant function to the data, with the location and number of discontinuities chosen to balance accuracy of fitting the data against complexity of the surrogate. At each step of FSFS, we use an average score from five-fold cross-validation. When applying FSFS to the varied-gradient dataset, the temperature and density gradients are also included among the features that can be selected. Other scalar features such as ι, $\mathrm{d}\iota/\mathrm{d}x$, $\mathrm{d}p/\mathrm{d}x$ and n_{fp} can be included, but these features are not selected by the procedure, which makes sense because they do not appear in the gyrokinetic equation and hence have only a more indirect relation to the heat flux.

Figure 10 shows the results of FSFS for regression on both the fixed-gradient and varied-gradient datasets, as well as for stable versus unstable classification on the varied-gradient dataset. For regression, we assess the performance using the coefficient of determination R^2. For classification, we choose features in FSFS using the log-loss measure (also known as cross-entropy; lower is better), while reporting also the accuracy and ROC-AUC scores (receiver operating characteristic area under the curve; higher is better) in figure 10(*b*). It can be seen that the prediction accuracy rapidly improves with the first few features, then levels off. With three features, the heat flux for the varied-gradient dataset can be predicted with $R^2 = 0.88$; adding nine more features results in a marginal improvement to $R^2 = 0.92$. Qualitatively similar results are obtained using other regression or classification models. As an example, figure 10 also shows FSFS results for 10-nearest-neighbours regression; the behaviour is similar to the XGBoost regression but with slightly lower R^2.

Another view on FSFS is provided by figure 11. Each panel shows, for a given number of features, how the predictions with that feature subset compare with the actual heat flux. In each panel, the regression was fit using 80 % of the data, and the figure shows the performance on the held-out 20 % of the data. A perfect prediction would correspond to all the points lying on the grey dotted line. As more features
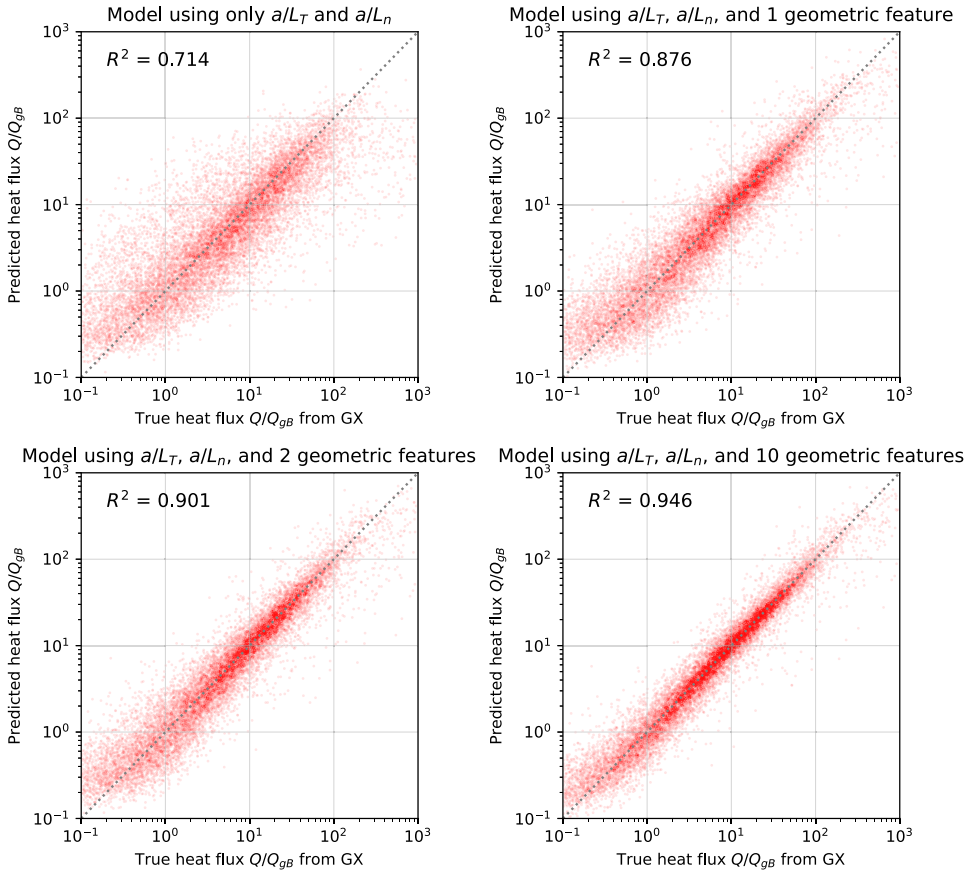
FIGURE 11. Accuracy of regression models improves to a point as more features are included from sequential feature selection. Each panel shows the performance of the XGBoost regression on 20 % held-out test data from the varied-gradient dataset.

are added, the regression models more accurately predict the heat flux. There is a trade-off in that the models with more features are more complicated and harder to interpret due to possible interactions between the features.

For each of the FSFS curves in figure 10, the first five selected features are listed in table 3. In the table, absFFTCoeff1 indicates the absolute magnitude of the longest-wavelength (but not constant) Fourier mode in z. In the varied-gradient dataset, for both regression and classification, the first selected feature is the temperature gradient and the second selected feature is the density gradient. These findings are consistent with expectations from linear theory that the temperature gradient is the primary drive for the instability and the density gradient can significantly affect stability. In both the varied-gradient and fixed-gradient cases, the first geometric feature selected again reflects the flux surface compression in regions of bad curvature. While the top geometric feature for the classifier involves unfavourable ∇B drift rather than curvature drift, the feature $\mathrm{mean}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y)|\nabla x|^2/B)$ has almost an identical score, log-loss = 0.123 as opposed to 0.122, so there is little distinction between the two drifts in this case. Hence, the results are quite consistent with the Spearman correlation analysis of §5.2, and support the physical intuition

Classification, varied gradients, XGBoost	
Feature	Log-loss
a/L_T	0.361
a/L_n	0.189
$\text{Mean}(\Theta(\boldsymbol{B} \times \nabla B \cdot \nabla y)|\nabla x|^2/B)$	0.122
$\text{Mean}(\Theta(-\boldsymbol{B} \times \nabla B \cdot \nabla x)|\nabla x|^2 B)$	0.105
$\text{Mean}((\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y)/B)$	0.094
Regression, fixed gradients, XGBoost	
Feature	R^2
$\text{Mean}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y)|\nabla x|^4/B)$	0.669
$\text{Quantile0.75}((\boldsymbol{B} \times \nabla B \cdot \nabla x)|\nabla x|^4/B^5)$	0.775
$\text{Mean}(|\nabla x|^4/B^6)$	0.834
$\text{Quantile0.9}((\boldsymbol{B} \times \nabla B \cdot \nabla y)/B^5)$	0.861
$\text{AbsFFTCoeff1}((\boldsymbol{B} \times \nabla B \cdot \nabla x)/B^3)$	0.879
Regression, varied gradients, XGBoost	
Feature	R^2
a/L_T	0.524
a/L_n	0.714
$\text{Mean}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y)|\nabla x|^4/B^2)$	0.876
$\text{Median}((\boldsymbol{B} \times \nabla B \cdot \nabla x)^2|\nabla x|^8/B^8)$	0.901
$\text{AbsFFTCoeff1}(\text{ReLU}(\boldsymbol{B} \times \nabla B \cdot \nabla x)/B^5)$	0.917
Regression, varied gradients, 10NN	
Feature	R^2
a/L_T	0.479
a/L_n	0.696
$\text{Quantile0.9}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla y)|\nabla x|^2/B)$	0.863
$\text{Mean}(\text{ReLU}(\boldsymbol{B} \times \nabla B \cdot \nabla x)|\nabla x|^2/B^4)$	0.884
$\text{Median}((\boldsymbol{B} \times \nabla B \cdot \nabla y)|\nabla y|^2/B^6)$	0.895

TABLE 3. First five features from § 5.1 selected with forward sequential feature selection. Results are shown both for classification of stability vs instability, and for regression on the logarithm of the heat flux Q. Results are also shown for both the gradient-boosted decision tree package XGBoost and for 10-nearest-neighbors (10NN). Here, Θ denotes the Heaviside function.

from Mynick *et al.* (2010), Xanthopoulos *et al.* (2014), Stroteich *et al.* (2022) and Goodman *et al.* (2024).

For each of the four tables within table 3, the second geometric feature again includes the flux surface compression $|\nabla x|$, but this time involving the radial rather than binormal component of the ∇B drift, $\boldsymbol{B} \times \nabla B \cdot \nabla x$. Note that there is no distinction between the ∇B drift and curvature drift in the radial direction: $\boldsymbol{B} \times \nabla B \cdot \nabla x = B \boldsymbol{B} \times \boldsymbol{\kappa} \cdot \nabla x$ for any MHD equilibrium at any β. So, the quantity $\boldsymbol{B} \times \nabla B \cdot \nabla x$ appearing in the second geometric features is the geodesic curvature (times $B|\boldsymbol{B} \times \nabla x|$). As mentioned in § 1, the geodesic curvature has been discussed recently as a correlate of turbulence (Xanthopoulos *et al.* 2011; Nakata & Matsuoka 2022), motivated by the fact that geodesic curvature plays a prominent role in the gyrokinetic equation for zonal flow modes (Rosenbluth & Hinton 1998). More work would be needed to confirm that the high significance of geodesic curvature in our

Sequential feature selection step 3	
Feature	R^2
$\text{Mean}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^4/B^2)$	0.876
$\text{Mean}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^4/B)$	0.874
$\text{Variance}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^2/B)$	0.871
$\text{Quantile0.9}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^2/B)$	0.870
$\text{Mean}(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^8/B^4)$	0.869
Sequential feature selection step 4	
Feature	R^2
$\text{Median}((\boldsymbol{B} \times \boldsymbol{\nabla} B \cdot \boldsymbol{\nabla} x)^2|\boldsymbol{\nabla} x|^8/B^8)$	0.901
$\text{Quantile0.75}(\text{ReLU}(-\boldsymbol{B} \times \boldsymbol{\nabla} B \cdot \boldsymbol{\nabla} x)^2|\boldsymbol{\nabla} x|^8/B^6)$	0.901
$\text{Median}((\boldsymbol{B} \times \boldsymbol{\nabla} B \cdot \boldsymbol{\nabla} x)^2|\boldsymbol{\nabla} x|^8/B^6)$	0.901
$\text{Median}(|\boldsymbol{B} \times \boldsymbol{\nabla} B \cdot \boldsymbol{\nabla} x||\boldsymbol{\nabla} x|^4/B^4)$	0.901
$\text{Quantile0.75}(\text{ReLU}(-\boldsymbol{B} \times \boldsymbol{\nabla} B \cdot \boldsymbol{\nabla} x)|\boldsymbol{\nabla} x|^4/B^3)$	0.901

TABLE 4. Top-scoring features from steps 3–4 of forward sequential feature selection, for regression on the heat flux using the varied-gradient dataset with XGBoost. At each step, there are many features which are variations on a theme that have nearly identical R^2 score. Here, Θ denotes the Heaviside function.

analysis is due to the effect of zonal flows on the turbulence, but the connection is at least suggestive.

At each step of FSFS, there are typically many features which are variations on a theme and which would give nearly the same score, similar to table 2. For instance, table 4 shows the top five features in steps three and four of FSFS for regressions on the varied-gradient dataset. At step 3, the top features are all based on the flux surface compression $|\boldsymbol{\nabla} x|$ in regions of bad curvature, weighted by various powers of the Jacobian. At step 4, the top features all involve the radial magnetic drift, weighted by the flux surface compression and powers of the Jacobian. Since the R^2 score is so similar among the various options at each step, we should not ascribe too much importance to the details that vary, such as the power of B at step 3. However, aspects that are consistent among the top features, such as the consistent appearance of the radial magnetic drift (i.e. geodesic curvature) at step 4, are more likely to be physically meaningful.

Other regression and classification models besides XGBoost can achieve a reasonable fit to the data. Figure 12 shows a comparison of several regression methods for the heat flux, all using the same set of 12 features with the varied-gradient dataset (a/L_T, a/L_n and the top 10 geometric features from FSFS). In addition to XGBoost and 10-nearest-neighbours (10NN), performance is also shown for three other models. One is the gradient-boosted decision tree package LightGBM (Ke *et al.* 2017). Another is random forest regression. Lastly is linear regression after the features are transformed with the Yeo–Johnson power transform (Yeo & Johnson 2000). Models other than XGBoost and LightGBM use the implementations from scikit-learn (Pedregosa *et al.* 2011). All models use the default hyperparameters from the relevant python package and the performance shown is the average score from five-fold cross-validation. While the prediction accuracy is highest for the decision-tree methods, it is meaningfully high for the other methods as well.
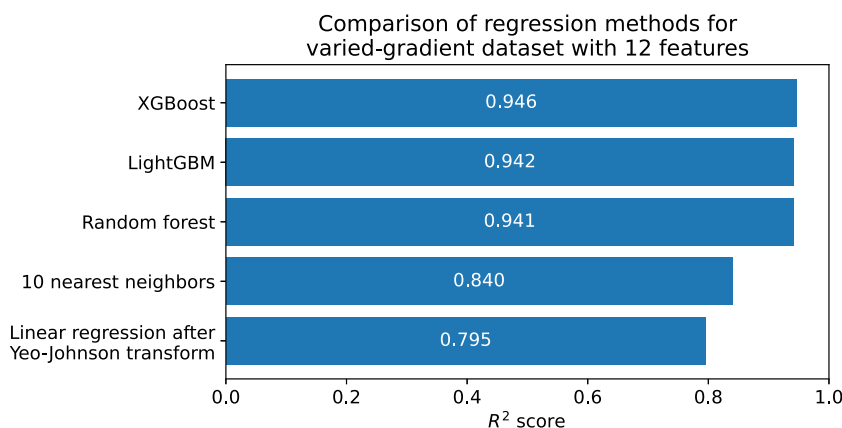
Comparison of regression methods for
varied-gradient dataset with 12 features

FIGURE 12. Comparison of regression methods for the heat flux in the varied-gradient dataset. In all cases, the feature set is a/L_T, a/L_n and the top 10 geometric features from FSFS with XGBoost.

5.4. *Shapley values*

Another way to measure the importance of the features in regression and classification models is via Shapley values. A concept originating from game theory, Shapley values are a fair way to divide value among the input features, based on how the model performance degrades when that feature is removed, and considering all subsets of features (Shapley 1953; Lipovetsky & Conklin 2001). Here, we specifically use the Shapley Additive exPlanation (SHAP) formulation by Lundberg & Lee (2017). Shapley values naturally provide the average sign of correlation between each feature and the target. For calculations here, we use the SHAP python package (Lundberg & Lee 2017; Lundberg *et al*. 2020), which allows for efficient computations with decision tree models.

Figure 13 shows the distribution of Shapley values for regression on the varied-gradient dataset, using the first 12 features from FSFS. Within each row, corresponding to one of the features, a histogram of the Shapley values is shown. The horizontal dimension gives the contribution to $\ln Q$ from that feature. Wider distributions therefore indicate more important features. The features are sorted in decreasing order of the average absolute Shapley value, i.e. decreasing importance. Each point is coloured by the value of the feature for that flux tube, scaled to the range of that feature over the dataset. So, if the distribution is purple/dark on the left and yellow/light on the right, larger values of that feature increase the predicted Q. Conversely if the distribution is yellow/light on the left and purple/dark on the right, increasing values of the feature decrease Q.

The order of feature importance as measured by Shapley values in figure 13 is not exactly the same as the order in which features were selected in FSFS, but the first three features are the same, and the pattern is broadly similar. Unsurprisingly, the most important feature is a/L_T, which is found to increase Q. The next most important feature is a/L_n, which decreases Q, as expected from linear theory. Consistent with the findings from Spearman correlation and FSFS, the most important geometric feature is an average of $|\nabla x|$ in regions of bad curvature. Here, the next most important feature, mean($|\nabla x|^4/B^6$), also reflects the average flux surface compression, but now independent of bad curvature. As the distributions for these
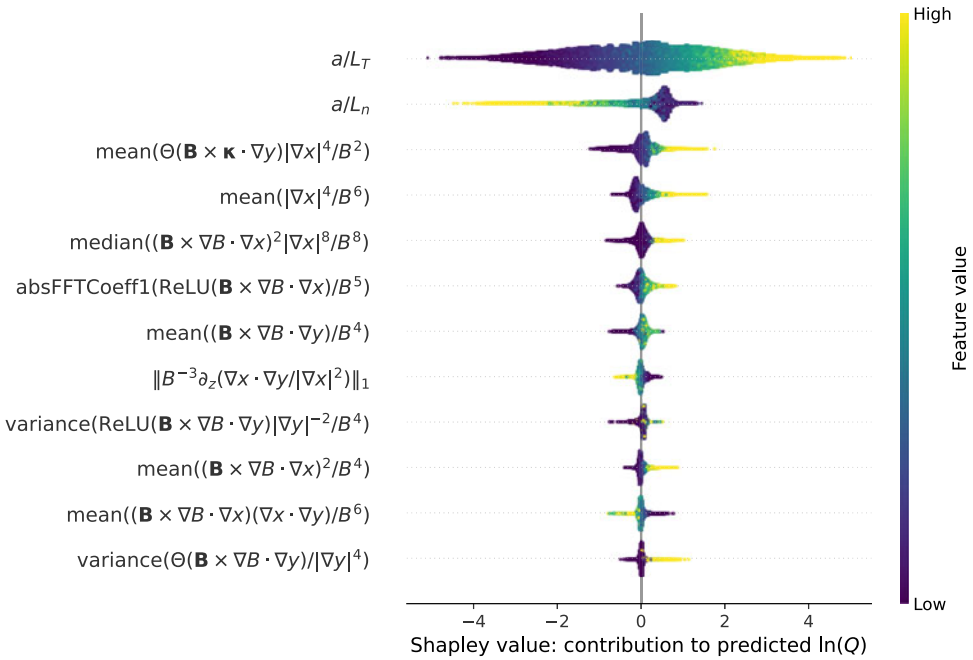
FIGURE 13. Distribution of Shapley values for regression with the varied-gradient dataset, using an XGBoost fit with top 12 features from FSFS. Features are listed in decreasing importance as measured by mean magnitude of the Shapley values.

two features are blue on the left and yellow on the right, both of these features increase Q, consistent with physical intuition. The next most important features, median$((\boldsymbol{B} \times \nabla B \cdot \nabla x)^2 |\nabla x|^8 / B^8)$ and absFFTCoeff1(ReLU$(\boldsymbol{B} \times \nabla B \cdot \nabla x)/B^5)$, both involve the geodesic curvature, consistent with the pattern in table 3. The colours indicate that larger average magnitude of geodesic curvature increases Q, coinciding with the theoretical prediction of Xanthopoulos *et al.* (2011) and Nakata & Matsuoka (2022). Physically, larger geodesic curvature means greater neoclassical damping of zonal flows, hence smaller average zonal flow magnitude, and therefore higher turbulence intensity.

The second-most-important geometric feature according to Shapely values, mean$(|\boldsymbol{\nabla} x|^4 / B^6)$, was the fifth most important geometric feature according to FSFS order. Other differences appear in the ranking of feature importance between Shapley values and FSFS after the two geodesic curvature features. Since Shapely values and FSFS order are different measures of importance, there is no guarantee that they will rank the features in the same order, particularly farther down in the list where the effect sizes become smaller.

5.5. *Fine-tuning the top feature*

From the previous subsections, there is a consistent and robust finding that the most important geometric feature is the flux surface compression in regions of bad curvature. With this in mind, we can repeat the analyses of the previous subsections with a new set of extracted features that focuses on this general combination of quantities, giving more variations on this theme. Optimisation can also be used
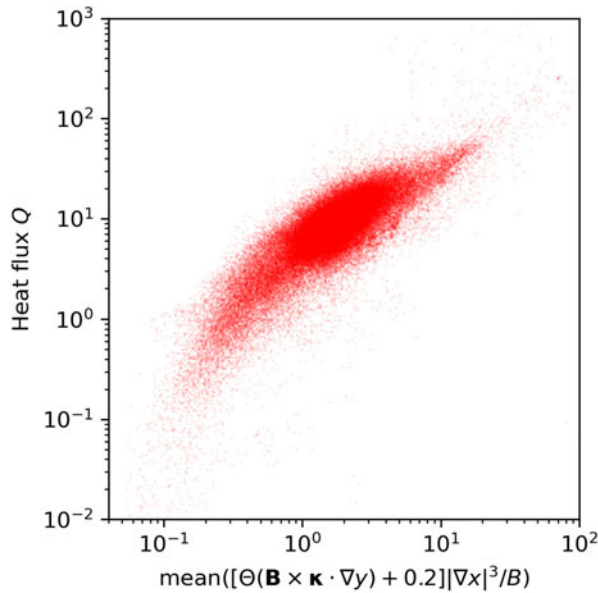
FIGURE 14. Comparing the true heat flux and the single geometric feature f_Q in (5.1) for fixed gradients, it is clear that there is significant correlation. No regression model is used here.

to tune constants and exponents appearing within a proposed functional form. For example, considering features of the form $\text{mean}([\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y) + \alpha]|\boldsymbol{\nabla} x|^\beta / B^\gamma)$, the parameters $\{\alpha, \beta, \gamma\}$ can be optimised to maximise R^2. Both approaches – trying a set of variations on a theme and optimisation – result in broadly similar conclusions: slightly improved prediction of Q is found if a shift is added to the Heaviside function and the exponent on $|\boldsymbol{\nabla} x|$ is modified, yielding the feature

$$f_Q = \text{mean}([\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y) + 0.2]|\boldsymbol{\nabla} x|^3 / B). \tag{5.1}$$

The shift of 0.2 to the Heaviside function effectively combines the top two geometric features according to the Shapley values in figure 13 (though with different powers of B.) For this single feature with the fixed-gradient dataset, the Spearman correlation is 0.788 and XGBoost regression on $\ln Q$ yields $R^2 = 0.737$. These values are slightly increased from 0.775 and 0.669, respectively, without the shift to the Heaviside function. For the varied-gradient dataset, using models with the three features $\{a/L_T, a/L_n, f_Q\}$, regression on $\ln Q$ gives $R^2 = 0.887$, slightly improved from $R^2 = 0.876$ from table 3.

Figure 14 shows a comparison of the single feature f_Q and the GX heat flux for the fixed-gradient dataset. Note that there is no regression model applied here. While (5.1) does not explain the full variation in the heat flux, it clearly does explain some. It is possible that with additional work, a single geometric feature of comparable complexity could be found with even higher predictive accuracy.

The set of three features $\{a/L_T, a/L_n, f_Q\}$ does reasonably well in the classifier for predicting stability: log-loss $= 0.128$, accuracy $= 0.944$, ROC-AUC $= 0.989$. However, the geometric feature that best predicts the stiffness of the heat flux above marginal stability is not necessarily the best feature for predicting the stability boundary. Using the same methods described earlier in this subsection, a single geometric feature can be fine-tuned to marginally improve the stability classifier's accuracy.
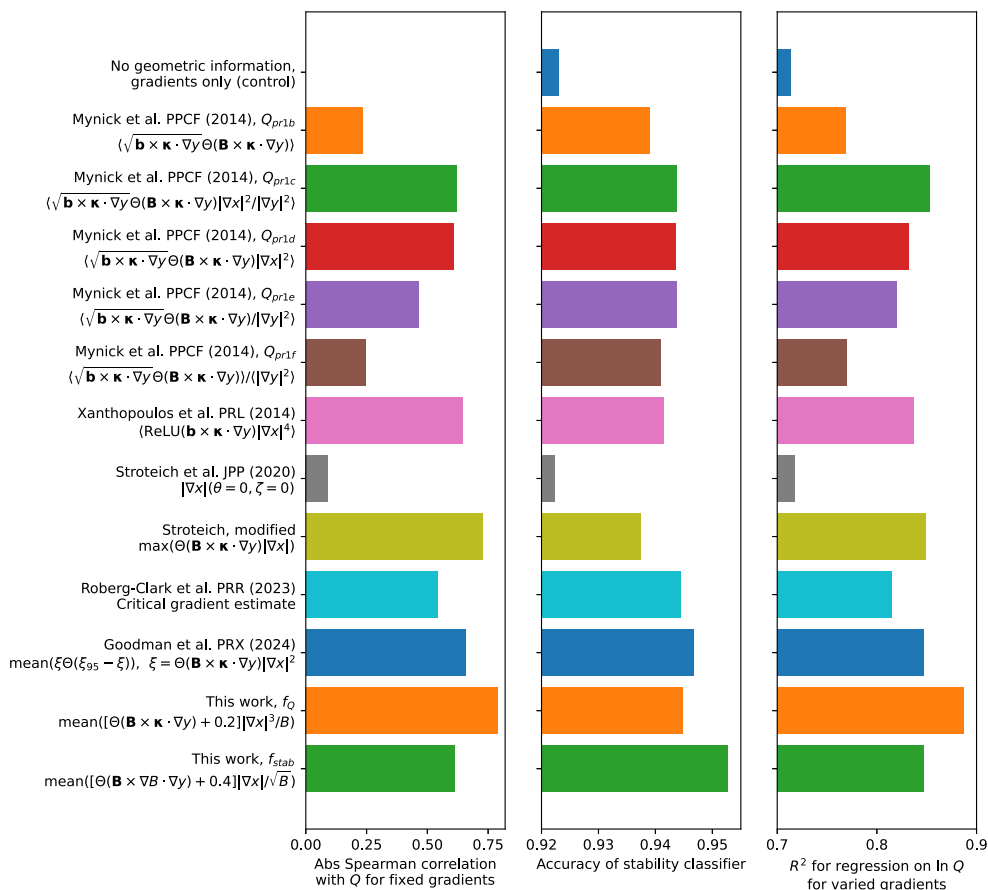
FIGURE 15. Comparing several proposed ITG objectives using three scores. The classification and regression scores are computed using XGBoost with three features: the single geometric feature, a/L_T and a/L_n.

Again, adding a shift to the Heaviside function and modifying the exponents, we arrive at

$$f_{\text{stab}} = \text{mean}([\Theta(\boldsymbol{B} \times \boldsymbol{\nabla} B \cdot \boldsymbol{\nabla} y) + 0.4]|\boldsymbol{\nabla} x|/\sqrt{B}). \tag{5.2}$$

Using XGBoost with the three features $\{a/L_T, a/L_n, f_{\text{stab}}\}$, the classifier achieves log-loss = 0.111, accuracy = 0.953 and ROC-AUC = 0.991.

6. Testing other proposed surrogates

A natural application of this dataset is to test objective functions for reduced ITG transport that have been proposed previously. Several such objectives are compared in figure 15. Each quantity is rated by three scores: Spearman correlation with Q for the fixed-gradient dataset, accuracy for classification with the varied-gradient dataset and R^2 for regression on the varied-gradient dataset. For the latter two scores, an XGBoost model is fit using three features: the geometric feature in question along with a/L_T and a/L_n. The classification and regression scores are averages from five-fold cross-validation. A control is also shown, in which XGBoost is fit using only

the two features $\{a/L_T, a/L_n\}$ with no geometric features, since the accuracy using only these gradients is already significant.

One set of proposed objectives in figure 15 comes from Mynick *et al*. (2014): Q_{pr1b}-Q_{pr1f}. The definitions can be found in figure 15, in terms of the flux tube average $\langle \ldots \rangle = [\int \ldots \mathrm{d}z/B]/\int \mathrm{d}z/B$. In the earlier publication (Mynick *et al*. 2014), the average was not defined, and it was stated that the component of curvature used was the radial component, but we confirmed with the first two authors that the expressions in figure 15 are what was actually used. Another similar proposed objective in figure 15 is that from Xanthopoulos *et al*. (2014): $\langle \mathrm{ReLU}(\boldsymbol{b} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^4 \rangle$. In that paper, the average was not written and again the component of curvature used was described as radial, but we confirmed with the author that the expression in figure 15 is what was actually used. Our sign convention for y is reversed compared with these earlier papers, so the arguments of Heaviside and ReLU functions have opposite sign.

One more previously proposed objective in figure 15 is that from Stroteich *et al*. (2022), namely $|\boldsymbol{\nabla} x|$ at a specific point $\boldsymbol{p}$: the outboard midplane of the taller stellarator-symmetric cross-section. This quantity is not well correlated with Q in our dataset because 3/4 of the flux tubes in our data do not include this point. Therefore, we also include a modified objective which is usually similar for flux tubes that include $\boldsymbol{p}$ but more meaningful in those that do not: the maximum $|\boldsymbol{\nabla} x|$ in regions of bad curvature: $\max(\Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|)$.

Another feature included in figure 15 is the estimated critical gradient from equation (6) of Roberg-Clark *et al*. (2023). We use the curvature drift rather than ∇B drift to compute this quantity as it results in slightly higher scores, though the differences are small. A typical flux tube may have multiple intervals of bad curvature, in which case, the estimated critical gradient is computed for each interval and the minimum over all intervals is used as the feature. We note here that the critical gradient estimate is not intended to correlate inversely with heat fluxes, but rather to predict the gradient where heat fluxes become significant. Any inverse correlation of the feature with heat fluxes above the critical gradient, perhaps through the physical connection between $|\boldsymbol{\nabla} \alpha|$ and finite Larmor radius damping of ITG modes, is incidental.

Finally, the comparison in figure 15 includes the ITG objective from Goodman *et al*. (2024). In that work, an integral over a full flux surface was used; here, we modify the expression to integrate only over a flux tube. The specific expression we consider is $\mathrm{mean}(\xi \Theta(\xi_{95} - \xi))$, where $\xi = \Theta(\boldsymbol{B} \times \boldsymbol{\kappa} \cdot \boldsymbol{\nabla} y)|\boldsymbol{\nabla} x|^2$ and ξ_{95} is the 0.95 quantile of ξ over the flux tube.

Figure 15 shows that with the exception of $|\boldsymbol{\nabla} x|(\theta = 0, \zeta = 0)$, all of these proposed objectives have significant predictive power for both the stability boundary and heat flux. This is perhaps not surprising as these objectives are similar to each other, all measures of the flux surface compression, bad curvature or both. Their performance is nearly as good as for our optimised features f_Q and f_{stab}, and is significantly better than the control with no geometric information.

7. Discussion and future work

In this work, we have presented a large new dataset of nonlinear gyrokinetic turbulence simulations covering a wide and diverse range of magnetic geometries. Applied to the dataset, a variety of machine learning methods can accurately predict the nonlinear heat flux and classify stable versus unstable conditions. It was important that these classification and regression methods be applied in such a way as to respect the

translation-invariance of the gyrokinetic system, as can be done using convolutional neural networks or translation-invariant reductions of transation-equivariant operations. Beyond providing these fast surrogates, machine learning methods can also extract insights that can stimulate theory. Thus, machine learning can be more than a black-box interpolation: it can provide understanding and feed back into more traditional physics calculations. To extract insights in this case, we have demonstrated a variety of methods – Spearman correlation, sequential feature selection and Shapley values – which can measure the importance of geometric features in the data. While some details differ between the conclusions of these methods, such as the exponent of B in the most important features, certain patterns are quite robust.

Multiple lines of evidence point to the flux surface compression in regions of bad curvature as the most important geometric feature, with more surface compression yielding higher heat flux Q. This feature is identified from its high Spearman correlation, from its early choice in sequential feature selection and by its large Shapley values. It arises in FSFS using multiple regression algorithms (decision trees and nearest-neighbours), using both regression and classification, and both with fixed or varied temperature and density gradients. These findings provide evidence supporting the physical arguments regarding this feature by Mynick *et al.* (2010), Xanthopoulos *et al.* (2014), Stroteich *et al.* (2022) and Goodman *et al.* (2024). It is not obvious that the importance of this feature can be explained purely in terms of linear growth rates, as $|\nabla x|$ does not appear in the linear gyrokinetic equation for $k_x = 0$ modes, which are typically the most unstable.

Another robust finding is that the next most important feature is the geodesic curvature, with larger magnitudes giving higher Q. These conclusions are supported by both FSFS and Shapley values. Our findings support the discussion of this feature's importance by Xanthopoulos *et al.* (2011) and Nakata & Matsuoka (2022), and are suggestive of the effect of zonal flow dynamics. Further analysis of the existing dataset could be done to elucidate the role of zonal flows.

Another theoretical framework for turbulence that has been discussed recently is critical balance (Barnes *et al.* 2011), in which the parallel wavenumber $k_\parallel$ plays an important role. Although thousands of features involving $k_\parallel$ were included in the set of features that could have been selected, these features were not selected in FSFS, indicating lower importance than the surface compression or geodesic curvature. It would be valuable to better understand these findings theoretically in the future.

There are many directions in which this research can be extended in the areas of data generation, fitting the data with surrogate models, physics understanding and applications. Regarding data generation, the set of equilibria could be expanded with more quasi-isodynamic configurations. New datasets of nonlinear gyrokinetic simulations could be generated with kinetic electrons and electromagnetic effects. The physics model would then include additional instabilities such as trapped electron modes and kinetic ballooning modes. Similar analysis methods to those used here could then be applied to such data, including also regression on the particle flux. For both future data and the present data, the feature engineering methods here could be applied with larger sets of possible features beyond the set from § 5.1, and alternative regression methods could be tried. Of particular interest would be symbolic regression and Kolmogorov Arnold networks (Liu *et al.* 2024) due to their advantages for interpretability. It would also be valuable to use saliency maps and related methods to understand the features learned by the neural networks. If these saliency maps can be understood, the results may suggest new features that could

be checked directly for correlation with the true heat flux and to include in the FSFS.

In the area of physics understanding, researchers can aim to derive relationships between the top geometric features here and the nonlinear heat flux using traditional analytic methods. This should be done first for the most important feature (and variations thereof) related to flux surface compression, but could also be attempted for the next most important features. Other quantities inspired by plasma theory could be checked for correlations against the heat flux and added to the menu of possible features for sequential feature selection. It would be natural for instance to check quantities from linear zonal flow dynamics in this way.

The surrogate models developed in this paper could be applied in multiple ways. One application is for predicting the radial profiles of temperature, using the surrogate as a fast model for the gradient-flux relationship inside a solver of the transport equations. (Prediction of the density profile would require a dataset of nonlinear simulations with kinetic electrons.) Such profile prediction using surrogate models is already available for tokamaks (Citrin *et al*. 2015; Meneghini *et al*. 2017), but this could now be extended to stellarators. The other evident application would be for turbulence optimisation of stellarators. While optimisation of a geometric feature similar to the top feature here is already being done (Goodman *et al*. 2024), the results of this paper allow several improvements. First, this turbulence objective can now be justified through validation on this data and modified slightly as in § 5.5 for better correlation to nonlinear simulations. Second, the surrogates here provide better correlation to the true heat flux by incorporating more geometric information, either via multiple features in the FSFS method or through holistic use of all the geometric information in the neural networks. Lastly, since the surrogates here could allow information from all flux surfaces and the gradients to be combined in a transport solver to rapidly predict the fusion power, this enables the fusion power to be used directly as an optimisation objective.

Declaration of interests

M.L. is a consultant for Type One Energy Group.

Data availability statement

Appendix A. Details of equilibrium generation

As discussed in §3.1, the equilibria in this study consist of three classes. Here, we give additional details of how the equilibria in each class were generated. In all cases, the equilibria were computed using DESC (Dudt & Kolemen 2020) release v0.12.0.

For the group of rotating-ellipse equilibria, the number of field periods was sampled randomly from the interval [2, 8], the aspect ratio was chosen randomly from [6, 10], and the elongation (ratio of major to minor axis of the cross-section in a constant-ϕ plane) was chosen randomly from [1, 4]. For half of the rotating-ellipse configurations, the ellipses were centred on a constant-major-radius circle. For the other half, the ellipses were centred on a curve with torsion, by setting the $m = 0$, $n = n_{fp}$ Fourier mode of R and Z on the boundary to a random number between 0 and the minor radius, with sign chosen to increase iota. A pressure profile with fixed shape $p(s) = 1 - 1.8s + 0.8s^2$ is chosen, where s is the normalised toroidal flux, reflecting a plausible level of peaking, with random magnitude chosen for a uniform distribution of volume-averaged $\beta \in [0, 0.05]$. Equilibrium calculations were then run, assuming no toroidal current for simplicity. Configurations with $|\iota| < 0.2$ were dropped.

Another group of equilibria were derived from the QUASR database (Giuliani 2024; Giuliani *et al.* 2024). This database includes coils in addition to plasma shapes, and many entries in QUASR differ primarily in coil geometry while having similar plasma parameters. Also, some configurations in QUASR have much better quasisymmetry than others and we wish to focus on those with better quasisymmetry (since the other equilibrium groups in our study contain many non-quasisymmetric geometries.) Therefore, a subset of QUASR was selected as follows. First, all configurations with high quasisymmetry error, $\iota < 0.2$, or aspect ratio >10 were excluded. Then, for each symmetry class (QA versus QH) and n_{fp}, for each interval in aspect ratio (1–2, 2–3, 3–4, .., 9–10) and ι, the two configurations in QUASR with lowest QS error were chosen. DESC was run for these boundary shapes assuming a vacuum field and also with the aforementioned pressure profile for multiple pressure magnitudes spanning $\langle \beta \rangle \in [0, 5\%]$. For some strongly shaped geometries or high $\langle \beta \rangle$ values, the resulting force residual was high, indicating the equilibrium was not well converged, so these cases were dropped.

The third group of equilibria were generated by randomly sampling Fourier modes from distributions that have been fit to a dataset of previous stellarator shapes. Consider the common representation of toroidal boundary shapes as a double Fourier series:

$$R(\theta, \phi) = \sum_{m,n} R_{m,n} \cos(m\theta - n_{fp}n\phi), \tag{A1}$$

$$Z(\theta, \phi) = \sum_{m,n} Z_{m,n} \sin(m\theta - n_{fp}n\phi), \tag{A2}$$

where (R, ϕ, Z) are cylindrical coordinates, n_{fp} is the number of field periods, θ is some poloidal angle and stellarator symmetry has been assumed. The sums are

considered to include only non-negative m, with both positive and negative n for $m > 0$, but only non-negative n for $m = 0$. To randomly generate boundary shapes, we sample the $R_{m,n}$ and $Z_{m,n}$ coefficients from independent normal distributions for each (m, n) (and independent for R and Z). The mean and standard deviation of each distribution are taken from the dataset of 44 stellarators collected by Kappel, Landreman & Malhotra (2024). This set includes both built experiments (W7-X, LHD, HSX, CFQS, TJ-II, etc.) and theoretical configurations. Before extracting the sample mean and standard deviation, all configurations were scaled to the same minor radius. To further standardise the data, $\phi \to -\phi$ reflections were applied to configurations with $\iota < 0$, so all configurations have rotational transform of the same sign, and toroidal rotations by half a field period were applied to any configurations in which the cross-section at $\phi = \pi/n_{\text{fp}}$ is not as tall as at $\phi = 0$. After this data cleaning, the sample mean and sample standard deviation are computed over the 44 configurations for each $R_{m,n}$ and $Z_{m,n}$. Configurations with all values of n_{fp} are included together in this calculation. Only modes with $m \leqslant 4$ and $|n| \leqslant 4$ are considered, since some theoretical stellarators in Kappel's dataset resulting from optimisation do not include boundary modes with higher mode number. The resulting mean and standard deviation values are used to define normal distributions which can then be sampled.

To generate a new random equilibrium, new values of $R_{m,n}$ and $Z_{m,n}$ are first sampled from the distributions determined previously, after which $R_{0,0}$ is computed by root-finding to give the desired aspect ratio. To ensure that the same gyro-Bohm normalisation is used in every turbulence simulation, each configuration is scaled slightly to the same minor radius, and the toroidal flux is set to the same value. As with the other equilibrium groups, a pressure profile with fixed shape $p(s) = 1 - 1.8s + 0.8s^2$ is chosen, with random magnitude chosen for a uniform distribution of volume-averaged $\beta \in [0, 0.05]$. The current profile is taken to be zero for simplicity. Next, a fast and low-resolution equilibrium calculation is run using the code VMEC (Hirshman & Whitson 1983). If this calculation does not converge to a threshold value of force residual in a given number of iterations, the configuration is rejected, ensuring that self-intersecting boundaries are excluded quickly.

A potential issue with the abovementioned procedure is that most resulting configurations in the third group have larger values of mirror ratio B_{max}/B_{min} than typical stellarators. For this reason, configurations with $B_{max}/B_{min} > 5$ are immediately rejected. Also, a subset of the random configurations is selected to bias the distribution towards smaller mirror ratios. For the selected configurations, higher-resolution equilibrium calculations are then run using DESC.

The final set of 23 577 equilibria included 3200 rotating-ellipse configurations centred on a circle, 3200 rotating-ellipse configurations centred on a curve with torsion, 413 QUASR vacuum configurations, 3964 QUASR finite-beta configurations and 3200 random boundaries for each of four values of n_{fp}. More flux tubes were drawn from the QUASR vacuum fields to increase the representation of this class. The final set of 100 705 flux tubes included 12 795 tubes from rotating-ellipse configurations centred on a circle, 12 791 tubes from rotating-ellipse configurations centred on a curve with torsion, 8235 tubes from QUASR vacuum configurations, 15 809 tubes from QUASR finite-beta configurations and 51 075 tubes from random-boundary equilibria.
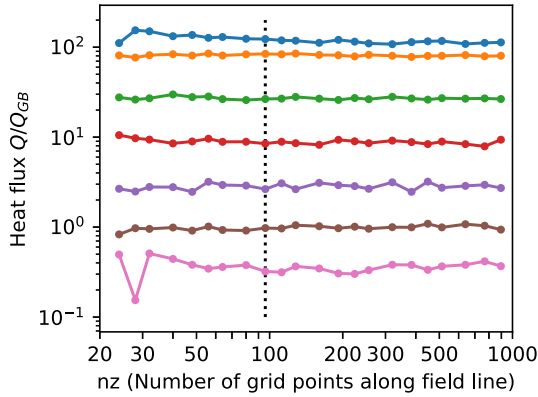
FIGURE 16. Dependence of Q on the resolution parameter nz for a selection of flux tubes with associated gradients from the varied-gradient dataset. The seven tubes, shown by different colours, all have the highest value of n_{fp} in the dataset (8) to give short scales in z. The vertical dotted line shows the value of nz used for the main dataset (96). Variation of Q with increasing resolution is small compared with the differences between geometries and gradients, indicating sufficient convergence.

Appendix B. Details of turbulence simulations

Here, more details are given of the nonlinear turbulence simulations. Calculations were performed using the version of GX from git commit b88d763.

For the dataset with varied gradients, the gradients were sampled randomly in each simulation using the following procedure. Since a/L_{Ti} typically increases with radius, we choose it to be ρ times a random number with mean 4 and standard deviation 3, resampling the latter until the result is $\geqslant 1.5$. Similarly, a/L_n is chosen to be ρ times a random number with mean 1 and standard deviation 2, resampling until the number is $\geqslant -0.5$. This procedure results in 70 % of the simulations yielding instability, 30 % stable.

In all simulations, physical collisions were included with magnitude $\nu_{ii} a_{\mathrm{minor}}/v_i = 0.01$ and we take $T_i/T_e = 1$. The transient behaviour for $t v_i/a_{\mathrm{minor}} < 150$ was ignored and a mean of the heat flux was computed over the remaining simulation time.

Resolution parameters for the simulations were as follows: box size x0=y0=10 $\times 2\pi$, perpendicular spatial resolution nx=ny=64, (so the dealiased k_x and k_y grids were 2.1, $-2.0, \ldots, -0.1, 0, 0.1, \ldots, 2.0, 2.1$), number of grid points along the field line nz=96, parallel velocity resolution nhermite=8, perpendicular velocity resolution nlaguerre=4, simulation time $t_{\mathrm{max}} v_i/a_{\mathrm{minor}} = 800$ and time step given by 0.9 times the Courant–Friedrichs–Lewy condition. To arrive at these values, first, each of these parameters was scanned individually for a selection of flux tubes and gradients. An example is shown in figure 16, in which nz is varied for several configurations with $n_{fp} = 8$, the highest value of n_{fp} included in the dataset, resulting in the shortest parallel length scales among the flux tubes in the collection. For each parameter, a value was adopted at which Q had approximately reached an asymptote. For changes to the box size in x or y, nx or ny was varied proportionally to keep the highest wavenumber fixed. Flux tube length was not considered to be a resolution parameter like the others abovementioned, since longer tubes do sample different regions of the flux surface geometry, so changes to Q are physical rather than numerical. Then, to confirm that the selected resolution parameters would be
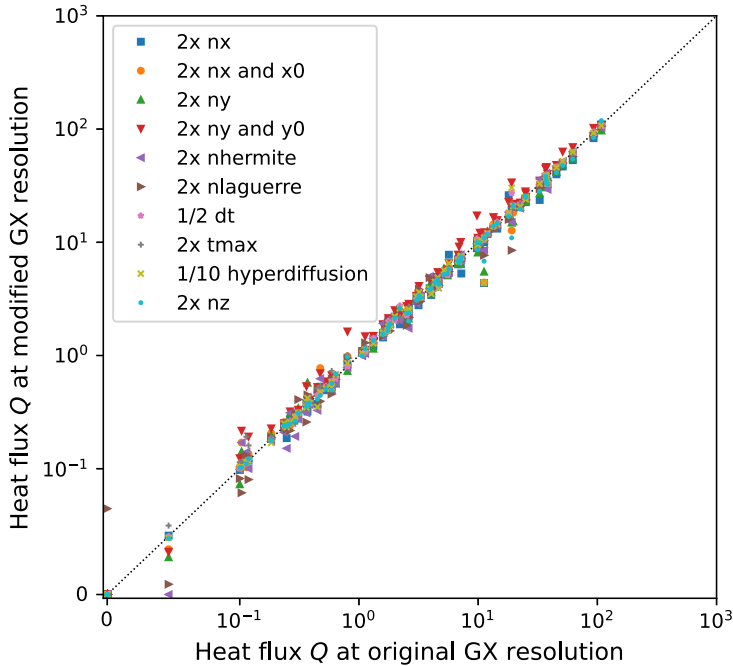
FIGURE 17. Evidence of sufficient convergence with respect to numerical resolution parameters. For 100 randomly sampled entries in the varied-gradient dataset, every resolution parameter is varied by a factor of 2 or 10. The box sizes in *x* and *y* are denoted *x*0 and *y*0 in the legend.

sufficient for most simulations in the dataset, 100 random flux tubes and gradients were selected from the varied-gradient data. Every resolution parameter was varied by a factor of two or more for each flux tube and GX was re-run. The results are displayed in figure 17. The coefficient of determination between the original-resolution data and modified-resolution data is $R^2 = 0.993$ for the unstable cases, or including both stable and unstable cases $R^2 = 0.995$ using (3.1). The accuracy score for predicting the stability at high resolution from the stability at standard resolution is 0.994. These values are greater than the values for any of the surrogate fits discussed in this paper, meaning scatter due to discretisation error is smaller than imperfection in the fits. We therefore deem the resolution parameters to be adequate.

Periodic boundary conditions were employed in all three spatial coordinates. For the *z* coordinate, the twist-and-shift boundary condition (Beer *et al.* 1995) and its generalisation to stellarators (Martin *et al.* 2018), while well-motivated physically, are inconvenient when the integrated local magnetic shear is small. In this case, the box size in *x* is required to be very large, necessitating large nx to resolve adequately high k_x, increasing computation time. Were we to use the alternative boundary conditions by Martin *et al.* (2018) that constrain the tube length to certain specific values, it would no longer be possible to use the same tube length for all configurations, complicating the analysis. Martin *et al.* (2018) found that the heat flux is insensitive to the choice of boundary conditions as long as enough Fourier modes in *x* are included to resolve sufficiently high k_x. We find the same to be true for the simulations here. Figure 18 shows a comparison of periodic versus twist-and-shift boundary conditions (using equations (24)–(25) of Martin *et al.* 2018) for a
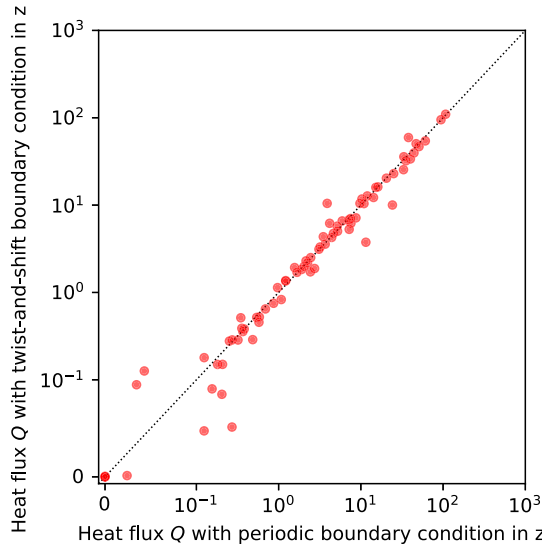
FIGURE 18. Comparison of boundary conditions in z for 100 randomly sampled flux tubes from the varied-gradient dataset. The heat flux is insensitive to the choice of boundary condition if the number of Fourier modes in x is increased in each twist-and-shift calculation to match the same maximum k_x as the periodic calculations, as is done here.

random sample of 100 flux tubes and associated gradients from the varied-gradient dataset. For the twist-and-shift calculations, the box size in x is set by the quantisation condition and nx is increased as needed for each tube to match the same maximum k_x as in the periodic case ($k_x = 2.1$). The heat fluxes for the two choices of boundary conditions are highly correlated. The R^2 for predicting $\ln(Q_{linked})$ from $\ln(Q_{periodic})$ is 0.97, and the accuracy score for predicting stability for the twist-and-shift case based on stability for the periodic case is 0.95. These values are deemed sufficiently high that the periodic boundary condition was adopted for the study, given the computational savings it allows.

At the resolution parameters given previously, the mean simulation wallclock time was ∼8 minutes on one Nvidia A100 GPU. The ∼2×10^5 nonlinear simulations took <7000 node-hours, equivalent to <28 000 gpu-hours (4 gpus/node).

In terms of normalised variables, the heat flux returned directly by GX is $Q/\langle|\nabla\rho|\rangle$, where $Q = \langle \int d^3 v \, f (v^2/2) \boldsymbol{v}_d \cdot \nabla\rho \rangle$ and $\langle \ldots \rangle = (\int dz/B)^{-1} \int dz (\ldots)/B$. For results in this paper, we multiply through by $\langle|\nabla\rho|\rangle$ to focus on Q itself.

Simulations with $Q > 10^3$ are dropped from the dataset, since longer simulation times and larger box sizes are likely needed for adequate resolution of these cases.

REFERENCES

ABBATE, J., CONLIN, R. & KOLEMEN, E. 2021 Data-driven profile prediction for DIII-D. *Nucl. Fusion* **61**, 046027.

BALAPRAKASH, P., SALIM, M., URAM, T.D., VISHWANATH, V. & WILD, S.M. 2018 Deephyper: Asynchronous hyperparameter search for deep neural networks. In *2018 IEEE 25th international conference on high performance computing (HiPC)*, pp. 42–51. IEEE.

BARNES, M., PARRA, F.I. & SCHEKOCHIHIN, A.A. 2011 Critically balanced ion temperature gradient turbulence in fusion plasmas. *Phys. Rev. Lett.* **107**, 115003.

BEER, M.A., COWLEY, S.C. & HAMMETT, G.W. 1995 Field-aligned coordinates for nonlinear simulations of tokamak turbulence. *Phys. Plasmas* **2**, 2687–2700.

BEURSKENS, M.N.A., *et al.* 2021 Ion temperature clamping in Wendelstein 7-X electron cyclotron heated plasmas. *Nucl. Fusion* **61**, 116072.

BOYER, M.D. & CHADWICK, J. 2021 Prediction of electron density and pressure profile shapes on NSTX-U using neural networks. *Nucl. Fusion* **61**, 046024.

BRUNTON, S.L., PROCTOR, J.L. & KUTZ, J.N. 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci.* **113**, 3932–3937.

CHEN, T. & GUESTRIN, C. 2016 Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794.

CHEN, Y., PARKER, S.E., COHEN, B.I., DIMITS, A.M., NEVINS, W.M., SHUMAKER, D., DECYK, V.K. & LEBOEUF, J.N. 2003 Simulations of turbulent transport with kinetic electrons and electromagnetic effects. *Nucl. Fusion* **43**, 1121–1127.

CITRIN, J., BRETON, S., FELICI, F., IMBEAUX, F., ANIEL, T., ARTAUD, J.F., BAIOCCHI, B., BOURDELLE, C., CAMENEN, Y. & GARCIA, J. 2015 Real-time capable first principle based modelling of tokamak turbulent transport. *Nucl. Fusion* **55**, 092001.

CURVO, P., FERREIRA, D.R. & JORGE, R. 2025 Using deep learning to design high aspect ratio fusion devices. *J. Plasma Phys.* **91**, E38.

DUDT, D.W. & KOLEMEN, E. 2020 Desc: a stellarator equilibrium solver. *Phys. Plasmas* **27**, 102513.

EFROYMSON, M.A. 1960 Multiple regression analysis. *Math. Meth. Digital Comput.* **3**, 191–203.

EGELÉ, R., GUYON, I., VISHWANATH, V. & BALAPRAKASH, P. 2023 Asynchronous decentralized Bayesian optimization for large scale hyperparameter optimization. In *2023 IEEE 19th International Conference on e-Science (e-Science)*, pp. 1–10. IEEE.

FRIEMAN, E.A. & CHEN, L. 1982 Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. *Phys. Fluids* **25**, 502–508.

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., RODRÍGUEZ, E. & SPIVAK, M. 2024 A comprehensive exploration of quasisymmetric stellarators and their coil sets. arXiv:2409.04826.

GOODMAN, A.G., XANTHOPOULOS, P., PLUNK, G.G., SMITH, H., NÜHRENBERG, C., BEIDLER, C.D., HENNEBERG, S.A., ROBERG-CLARK, G., DREVLAK, M. & HELANDER, P. 2024 Quasi-isodynamic stellarators with low turbulence as fusion reactor candidates. *PRX Energy* **3**, 023010.

GREENE, J.M. & JOHNSON, J.L. 1968 Interchange instabilities in ideal hydromagnetic theory. *Plasma Phys.* **10**, 729–745.

HEGNA, C.C., TERRY, P.W. & FABER, B.J. 2018 Theory of ITG turbulent saturation in stellarators: Identifying mechanisms to reduce turbulent transport. *Phys. Plasmas* **25**, 022511.

HIRSHMAN, S.P. & WHITSON, J.C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Phys. Fluids* **26**, 3553–3568.

HONDA, M. & NARITA, E. 2019 Machine-learning assisted steady-state profile predictions using global optimization techniques. *Phys. Plasmas* **26**, 102307.

JORGE, R., DORLAND, W., KIM, P., LANDREMAN, M., MANDELL, N.R., MERLO, G. & QIAN, T. 2024 Direct microstability optimization of stellarator devices. *Phys. Rev. E* **110**, 035201.

JORGE, R. & LANDREMAN, M. 2020 The use of near-axis magnetic fields for stellarator turbulence simulations. *Plasma Phys. Control. Fusion* **63**, 014001.

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.

KE, G., MENG, Q., FINLEY, T., WANG, T., CHEN, W., MA, W., YE, Q. & LIU, T.-Y. 2017 Lightgbm: a highly efficient gradient boosting decision tree. *Adv. Neur. Info. Process. Syst.* **30**. https://proceedings.neurips.cc/paper_files/paper/2017/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf.

KOZA, J.R. 1994 Genetic programming as a means for programming computers by natural selection. *Stat. Comput.* **4**, 87–112.

KRUSKAL, M.D. & KULSRUD, R.M. 1958 Equilibrium of a magnetically confined plasma in a toroid. *Phys. Fluids* **1**, 265–274.

LANDREMAN, M. 2025 https://doi.org/10.5281/zenodo.14867776.

LI, H., WANG, L., FU, Y.L., WANG, Z.X., WANG, T.B. & LI, J.Q. 2024 Surrogate model of turbulent transport in fusion plasmas using machine learning. *Nucl. Fusion* **65**, 016015.

LIPOVETSKY, S. & CONKLIN, M. 2001 Analysis of regression in game theory approach. *Appl. Stoch. Model. Bus. Indust.* **17**, 319–330.

LIU, Z., WANG, Y., VAIDYA, S., RUEHLE, F., HALVERSON, J., SOLJAČIĆ, M., HOU, T.Y. & TEGMARK, M. 2024, KAN: Kolmogorov-Arnold networks. arXiv:2404.19756.

LUNDBERG, S.M., ERION, G., CHEN, H., DEGRAVE, A., PRUTKIN, J.M., NAIR, B., KATZ, R., HIMMELFARB, J., BANSAL, N. & LEE, S.-I. 2020 From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* **2**, 2522–5839.

LUNDBERG, SCOTT M., LEE, SU-IN, (ED., IGUYON, LUXBURG, U.V., BENGIO, S., WALLACH, H., FERGUS, R., VISHWANATHAN, S. & GARNETT), R. 2017 A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, vol. 30, pp. 4765–4774. Curran Associates, Inc.

LUNSFORD, R., *et al.* 2021 Characterization of injection and confinement improvement through impurity induced profile modifications on the Wendelstein 7-X stellarator. *Phys. Plasmas* **28**, 082506.

MACKENBACH, R J J., PROLL, J H E. & HELANDER, P. 2022 Available energy of trapped electrons and its relation to turbulent transport. *Phys. Rev. Lett.* **128**, 175001.

MANDELL, N.R., DORLAND, W., ABEL, I., GAUR, R., KIM, P., MARTIN, M. & QIAN, T. 2024 GX: a GPU-native gyrokinetic turbulence code for tokamak and stellarator design. *J. Plasma Phys.* **90**, 905900402.

MARTIN, M.F., LANDREMAN, M., XANTHOPOULOS, P., MANDELL, N.R. & DORLAND, W. 2018 The parallel boundary condition for turbulence simulations in low magnetic shear devices. *Plasma Phys. Control. Fusion* **60**, 095008.

MCKINNEY, I.J., PUESCHEL, M.J., FABER, B.J., HEGNA, C.C., TALMADGE, J.N., ANDERSON, D.T., MYNICK, H.E. & XANTHOPOULOS, P. 2019 A comparison of turbulent transport in a quasi-helical and a quasi-axisymmetric stellarator. *J. Plasma Phys.* **85**, 905850503.

MENEGHINI, O., LUNA, C.J., SMITH, S.P. & LAO, L.L. 2014 Modeling of transport phenomena in tokamak plasmas with neural networks. *Phys. Plasmas* **21**, 060702.

MENEGHINI, O., *et al.* 2017 Self-consistent core-pedestal transport simulations with neural network accelerated models. *Nucl. Fusion* **57**, 086034.

MERLO, A., BÖCKENHOFF, D., SCHILLING, J., HÖFEL, U., KWAK, S., SVENSSON, J., PAVONE, A., LAZERSON, S.A. & PEDERSEN, T.S. 2021 Proof of concept of a fast surrogate model of the VMEC code via neural networks in Wendelstein 7-X scenarios. *Nucl. Fusion* **61**, 096039.

MYNICK, H.E., POMPHREY, N. & XANTHOPOULOS, P. 2010 Optimizing stellarators for turbulent transport. *Phys. Rev. Lett.* **105**, 095004.

MYNICK, H., XANTHOPOULOS, P., FABER, B., LUCIA, M., RORVIG, M. & TALMADGE, J.N. 2014 Turbulent optimization of toroidal configurations. *Plasma Phys. Control. Fusion* **56**, 094001.

NAKATA, M. & MATSUOKA, S. 2022 Impact of geodesic curvature on zonal flow generation in magnetically confined plasmas. *Plasma Fusion Res.* **17**, 1203077–1203077.

NAKAYAMA, T., NAKATA, M., HONDA, M., NARITA, E., NUNAMI, M. & MATSUOKA, S. 2023 A simplified model to estimate nonlinear turbulent transport by linear dynamics in plasma turbulence. *Sci. Rep.-UK* **13**, 2319.

NARITA, E., HONDA, M., NAKATA, M., YOSHIDA, M., HAYASHI, N. & TAKENAGA, H. 2019 Neural-network-based semi-empirical turbulent particle transport modelling founded on gyrokinetic analyses of JT-60U plasmas. *Nucl. Fusion* **59**, 106018.

PEDREGOSA, F., *et al.* 2011 Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830.

VAN DE PLASSCHE, K.L., *et al.* 2020 Fast modeling of turbulent transport in fusion plasmas using neural networks. *Phys. Plasmas* **27**, 022310.

PROLL, J.H.E., MYNICK, H.E., XANTHOPOULOS, P., LAZERSON, S.A. & FABER, B.J. 2015 TEM turbulence optimisation in stellarators. *Plasma Phys. Control. Fusion* **58**, 014006.

ROBERG-CLARK, G.T., PLUNK, G.G., XANTHOPOULOS, P., NÜHRENBERG, C., HENNEBERG, S.A. & SMITH, H.M. 2023 Critical gradient turbulence optimization toward a compact stellarator reactor concept. *Phys. Rev. Res.* **5**, L032030.

ROSENBLUTH, M.N. & HINTON, F.L. 1998 Poloidal flow driven by ion-temperature-gradient turbulence in tokamaks. *Phys. Rev. Lett.* **80**, 724–727.

SHAPLEY, L.S. 1953 A value for n-person games. *Contrib. Theory Games* **2**, 307–318.

STROTEICH, S., XANTHOPOULOS, P., PLUNK, G. & SCHNEIDER, R. 2022 Seeking turbulence optimized configurations for the Wendelstein 7-X stellarator: ion temperature gradient and electron temperature gradient turbulence. *J. Plasma Phys.* **88**, 175880501.

WAKASA, A., MURAKAMI, S., ITAGAKI, M. & OIKAWA, S.-I. 2007 Construction of neoclassical transport database for large helical device plasma applying neural network method. *Japan. J. Appl. Phys.* **46**, 1157.

WHITNEY, A.W. 1971 A direct method of nonparametric measurement selection. *IEEE Trans. Comput.* **100**, 1100–1103.

XANTHOPOULOS, P., MISCHCHENKO, A., HELANDER, P., SUGAMA, H. & WATANABE, T.-H. 2011 Zonal flow dynamics and control of turbulent transport in stellarators. *Phys. Rev. Lett.* **107**, 245002.

XANTHOPOULOS, P., MYNICK, H.E., HELANDER, P., TURKIN, YU, PLUNK, G.G., JENKO, F., GÖRLER, T., TOLD, D., BIRD, T. & PROLL, J.H.E. 2014 Controlling turbulence in present and future stellarators. *Phys. Rev. Lett.* **113**, 155001.

YEO, I.-K. & JOHNSON, R.A. 2000 A new family of power transformations to improve normality or symmetry. *Biometrika* **87**, 954–959.

ZHANG, D. *et al.* 2023 Observation of impurity accumulation and its compatibility with high plasma performance in W7-X. *Plasma Phys. Control. Fusion* **65**, 105006.