

PAPER • OPEN ACCESS

# Physics-regularized neural network of the ideal-MHD solution operator in Wendelstein 7-X configurations

To cite this article: Andrea Merlo *et al* 2023 *Nucl. Fusion* **63** 066020

View the [article online](#) for updates and enhancements.

## You may also like

- [The Magnetic Mechanism for Hotspot Reversals in Hot Jupiter Atmospheres](#)  
A. W. Hindle, P. J. Bushby and T. M. Rogers
- [NONLINEAR FORCE-FREE EXTRAPOLATION OF THE CORONAL MAGNETIC FIELD BASED ON THE MAGNETOHYDRODYNAMIC RELAXATION METHOD](#)  
S. Inoue, T. Magara, V. S. Pandey et al.
- [The Solar Wind with Hydrogen Ion Exchange and Large-scale Dynamics \(SHIELD\) Code: A Self-consistent Kinetic-Magnetohydrodynamic Model of the Outer Heliosphere](#)  
A. T. Michael, M. Opher, G. Tóth et al.

# Physics-regularized neural network of the ideal-MHD solution operator in Wendelstein 7-X configurations

Andrea Merlo\* , Daniel Böckenhoff , Jonathan Schilling , Samuel Aaron Lazerson , Thomas Sunn Pedersen  and the W7-X Team<sup>a</sup>

Max-Planck-Institute for Plasma Physics, 17491 Greifswald, Germany

E-mail: [andrea.merlo@ipp.mpg.de](mailto:andrea.merlo@ipp.mpg.de)

Received 30 December 2022, revised 13 March 2023

Accepted for publication 29 March 2023

Published 20 April 2023



CrossMark

## Abstract

The computational cost of constructing 3D magnetohydrodynamic (MHD) equilibria is one of the limiting factors in stellarator research and design. Although data-driven approaches have been proposed to provide fast 3D MHD equilibria, the accuracy with which equilibrium properties are reconstructed is unknown. In this work, we describe an artificial neural network (NN) that quickly approximates the ideal-MHD solution operator in Wendelstein 7-X (W7-X) configurations. This model fulfils equilibrium symmetries by construction. The MHD force residual regularizes the solution of the NN to satisfy the ideal-MHD equations. The model predicts the equilibrium solution with high accuracy, and it faithfully reconstructs global equilibrium quantities and proxy functions used in stellarator optimization. We also optimize W7-X magnetic configurations, where desirable configurations can be found in terms of fast particle confinement. This work demonstrates with which accuracy NN models can approximate the 3D ideal-MHD solution operator and reconstruct equilibrium properties of interest, and it suggests how they might be used to optimize stellarator magnetic configurations.

Keywords: neural networks, surrogate models, ideal-MHD, Wendelstein 7-X

(Some figures may appear in colour only in the online journal)

## 1. Introduction

The computational cost of constructing 3D magnetohydrodynamic (MHD) equilibria is one major limiting factor in stellarator theory, experimentation and design. Depending on the desired resolution and accuracy of the solution, such computations require up to  $\mathcal{O}(10)$  CPUh [1, 2].

Fast and accurate equilibrium reconstructions are crucial to interpret experimental results in magnetically confined devices. At Wendelstein 7-X (W7-X), the Bayesian framework MINERVA [3, 4] reconstructs the plasma state (e.g. ion and electron temperature). For each forward evaluation of the simulation model, a free-boundary ideal-MHD equilibrium has to be constructed. For each reconstructed state (i.e. a single timestamp), between  $\mathcal{O}(10^3)$  and  $\mathcal{O}10^6$  forward evaluations are required [5, 6]. Due to the computational cost of constructing free-boundary ideal-MHD equilibria, a self-consistent Bayesian inversion of a single plasma state has been attempted only once [5] and is not part of the standard data evaluation procedure.

W7-X is currently the largest stellarator experiment in operation. Its scientific mission is to assess the stellarator line on the path towards a fusion reactor [7]. The configuration space

<sup>a</sup> See Sunn Pedersen *et al* 2022 (<https://doi.org/10.1088/1741-4326/ac2cf5>) for the W7-X Team.

\* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

of W7-X provides a large experimental test bed, however, only nine reference configurations are mainly investigated [8]. Experimental time is limited, and priorities must be established to fit financial and human resources. Flight simulators provide a parallel and cheaper path to explore and exploit current and future experiments [9, 10]. In a stellarator, 3D MHD equilibria bound the accuracy and cost of simulation.

Differentiable and less expensive MHD equilibrium solutions would allow a more extensive exploration of the stellarator optimization space, in which the number of degrees of freedom is usually large  $\sim \mathcal{O}(10^2)$  [11, 12]. The computational cost of the objective function, which is largely dominated by deriving ideal-MHD equilibria, limits the extent to which we can explore the optimization space. According to recent works, each optimization run requires up to  $\mathcal{O}(10^3)$  function evaluations [11, 13, 14].

Data-driven approaches (e.g. artificial neural networks (NNs)) can provide fast 3D ideal-MHD equilibria [15–20]. When benchmarked against equilibria from non-linear codes (e.g. the variational moments equilibrium code (VMEC)), the geometry of the flux surfaces is in excellent agreement. However, it is unclear how equilibrium properties (e.g. plasma stability) can be faithfully reproduced by the NN model. Furthermore, it is unclear how well the NN equilibria satisfy the MHD equations. To make use of such NN models in downstream tasks (e.g. Bayesian inference, optimization), it is necessary to investigate how well NNs capture equilibrium properties.

The ideal-MHD equilibrium problem is a system of partial differential equations (PDEs). In general, initial and boundary conditions, as well as input parameters define the PDE problem. The solution operator maps these variable terms to the corresponding PDE solution.

In this paper, we extend [20] by learning the solution operator of the ideal-MHD equilibrium problem in the subspace of W7-X magnetic configurations with a deep operator network (DeepONet) [21]. A DeepONet is a NN macro-architecture with two pathways: a branch network to encode the PDE problem inputs, and a trunk network to represent the domain of solution function. Here, the NN model maps an externally applied field (defined by the currents in the coil system) and a plasma state (defined by the pressure and toroidal current profiles) to an approximation of the ideal-MHD equilibrium solution. Within the boundary of the training data (i.e. W7-X magnetic configurations), the NN model approximates the solution of unseen equilibrium problems at a fraction of the computational cost currently required by a traditional code. Summarizing the paper's contributions, we:

- Propose a DeepONet like architecture to represent the solution operator of the ideal-MHD equilibrium problem (see section 3);
- Train the model to approximate the solution of ideal-MHD equilibrium problem obtained from a non-linear code (e.g. VMEC), while the ideal-MHD force residual regularizes the model's solution (see section 4);
- Investigate how well the model's solution satisfies the ideal-MHD equations, and propose a strategy to improve the

solution at inference time by approximating an equivalent fixed-boundary equilibrium solution (see section 5.3);

- Investigate to which degree equilibrium properties of interest are faithfully reproduced, in the context of ideal-MHD stability (e.g. magnetic well), neoclassical transport (e.g. the effective ripple), and fast particle confinement (e.g. extrema of the magnetic field strength along a field line) (see sections 5.4–5.6);
- Show the use of the model in the *a posteriori* optimization of W7-X magnetic configurations (see section 5.7);

## 2. Ideal-MHD equilibria

### 2.1. The ideal-MHD force balance

The equilibrium magnetic field under the ideal-MHD model is characterized by the force balance equation, Ampere's and Gauss's law:

$$\vec{F} = -\vec{J} \times \vec{B} + \vec{\nabla} p = 0, \quad (1)$$

$$\vec{\nabla} \times \vec{B} = \mu_0 \vec{J}, \quad (2)$$

$$\vec{\nabla} \cdot \vec{B} = 0, \quad (3)$$

where  $\vec{B}$  is the magnetic field,  $\vec{J}$  is the current density,  $p$  is the plasma pressure, and  $\mu_0$  is the permeability of free space.

In this work, we assume that a set of nested toroidal flux surfaces exists, and that the pressure is a flux function. Let  $\vec{x}^* = (R, \phi, Z)$  be a cylindrical coordinate system where  $R$  is the major radius,  $\phi$  is the cylindrical toroidal angle, and  $Z$  is the height above mid-plane.  $\vec{\alpha} = (s, \theta, \varphi)$  is a flux coordinate system where  $s = \frac{\psi}{\psi_{\text{edge}}}$  is a radial-like coordinate,  $\theta$  is a poloidal-like angle, and  $\varphi$  is a toroidal-like angle (in this work,  $\varphi = \phi$ ).  $2\pi\psi$  is the toroidal flux enclosed by a flux surface, and  $\Phi_{\text{edge}} = 2\pi\psi_{\text{edge}}$  is the toroidal flux enclosed by the plasma boundary (i.e.  $s = 1$ ). In the inverse equilibrium formulation [22],  $\vec{\alpha}$  describes the independent variables of the computational grid, and the flux surfaces locations are described by the coordinate transformation  $f: \vec{\alpha} \rightarrow \vec{x}^*$ .

Because magnetic fields are divergence free ( $\vec{\nabla} \cdot \vec{B} = 0$ ) and flux surfaces are assumed to be nested ( $\vec{B} \cdot \vec{\nabla} s = B^s = 0$ ), the magnetic field can be written in contravariant form following [23]:

$$\vec{B} = \vec{\nabla} \varphi \times \vec{\nabla} \chi + \vec{\nabla} \psi \times \vec{\nabla} \theta^* = B^\theta \vec{e}_\theta + B^\varphi \vec{e}_\varphi, \quad (4)$$

where  $2\pi\chi$  is the poloidal flux,  $\theta^* = \theta + \lambda(s, \theta, \varphi)$  is the poloidal angle for which the magnetic field lines are straight in  $(s, \theta^*, \varphi)$ , and  $\lambda$  is a periodic function of  $\theta$  and  $\varphi$  with zero average.  $\vec{e}_i = \frac{\partial \vec{x}^*}{\partial \alpha_i}$  are the covariant basis vectors, and  $\vec{e}^i = \vec{\nabla} \alpha_i$  are the contravariant basis vectors.

The contravariant components of the magnetic field ( $B^i = \vec{B} \cdot \vec{e}^i$ ) are then:

$$B^\theta = \frac{1}{\sqrt{g}} \Phi' \left( \epsilon - \frac{\partial \lambda}{\partial \varphi} \right), \quad (5)$$

$$B^\varphi = \frac{1}{\sqrt{g}} \Phi' \left( 1 + \frac{\partial \lambda}{\partial \theta} \right), \quad (6)$$

where  $\iota$  is the rotational transform, the prime denotes  $\partial/\partial s$ , and  $\sqrt{g} = (\vec{e}_s \cdot \vec{e}_\theta \times \vec{e}_\varphi) = (\vec{\nabla}s \cdot \vec{\nabla}\theta \times \vec{\nabla}\varphi)^{-1}$  is the Jacobian of the coordinate transformation  $f$ .

The covariant representation of  $\vec{B}$  ( $B_i = \vec{B} \cdot \vec{e}_i$ ) can be obtained from equations (5) and (6) and the metric tensor  $g_{ij} = \vec{e}_i \cdot \vec{e}_j = \frac{\partial \vec{x}^*}{\partial \alpha_i} \cdot \frac{\partial \vec{x}^*}{\partial \alpha_j}$  as follows:

$$B_s = B^\theta g_{\theta s} + B^\varphi g_{\varphi s}, \quad (7)$$

$$B_\theta = B^\theta g_{\theta\theta} + B^\varphi g_{\varphi\theta}, \quad (8)$$

$$B_\varphi = B^\theta g_{\theta\varphi} + B^\varphi g_{\varphi\varphi}. \quad (9)$$

The contravariant components of the current density ( $J^i = \vec{J} \cdot \vec{e}^i = \frac{1}{\mu_0} (\vec{\nabla} \times \vec{B}) \cdot \vec{e}^i$ ) are:

$$J^s = \frac{1}{\mu_0 \sqrt{g}} \left( \frac{\partial B_\varphi}{\partial \theta} - \frac{\partial B_\theta}{\partial \varphi} \right), \quad (10)$$

$$J^\theta = \frac{1}{\mu_0 \sqrt{g}} \left( \frac{\partial B_s}{\partial \varphi} - \frac{\partial B_\varphi}{\partial s} \right), \quad (11)$$

$$J^\varphi = \frac{1}{\mu_0 \sqrt{g}} \left( \frac{\partial B_\theta}{\partial s} - \frac{\partial B_s}{\partial \theta} \right). \quad (12)$$

Finally, inserting equations (5), (6) and (10)–(12) in equation (1), the covariant form of the MHD force residual is given by:

$$\vec{F} = F_s \vec{\nabla}s + F_\beta \vec{\beta}, \quad (13)$$

$$F_s = \frac{1}{\mu_0} \left( B^\theta \frac{\partial B_\theta}{\partial s} - B^\theta \frac{\partial B_s}{\partial \theta} - B^\varphi \frac{\partial B_s}{\partial \varphi} + B^\varphi \frac{\partial B_\varphi}{\partial s} \right) + p', \quad (14)$$

$$F_\beta = J^s = \frac{1}{\mu_0 \sqrt{g}} \left( \frac{\partial B_\varphi}{\partial \theta} - \frac{\partial B_\theta}{\partial \varphi} \right), \quad (15)$$

where  $\vec{\beta} = \sqrt{g} (B^\varphi \vec{\nabla}\theta - B^\theta \vec{\nabla}\varphi)$ .

If we define  $\vec{x} = (R, \lambda, Z)$ , the MHD force residual can be then computed from the mapping  $\vec{x}(s, \theta, \varphi)$ , and the  $p(s)$  and  $\iota(s)$  flux functions (see section ‘Derivation of MHD force residual’).

## 2.2. Fixed- and free-boundary equilibria

The ideal-MHD problem admits two kinds of boundary conditions: fixed- and free-boundary conditions. In the fixed-boundary case, the shape ( $R$  and  $Z$ ) of the outermost flux surface (i.e.  $s = 1$ ) is fixed. In the free-boundary case, the shape of the outermost flux surface is determined by the continuity of the total pressure  $B^2/2\mu_0 + p$  and by the vanishing of the normal component of the vacuum field at the plasma-vacuum interface (thus enforcing the  $s = 1$  surface to be a flux surface).

Multiple quantities parameterize the equilibrium problem:  $\Phi_{\text{edge}} \in \mathbb{R}$  defines the total toroidal flux enclosed by the plasma,  $p : [0, 1] \rightarrow \mathbb{R}$  defines the pressure profile, and  $\iota : [0, 1] \rightarrow \mathbb{R}$  defines the rotational transform profile. The toroidal current profile  $I_{\text{tor}} : [0, 1] \rightarrow \mathbb{R}$  can be also specified in place of the iota profile: since  $I_{\text{tor}}(s) = \int_0^s \int_0^{2\pi} J^\varphi \sqrt{g} ds' d\theta$ , and  $J^\varphi$  linearly depends on  $\iota$ , the rotational transform profile consistent with a given toroidal current profile can be computed by solving a linear algebraic equation.

In case of free-boundary equilibria, a vacuum field, which is usually generated by a set of external coils, can be provided. A reference field generated by each coil is commonly pre-computed on the cylindrical grid, and the total vacuum field is given as the superposition of each coil field  $\vec{B}_v = \sum_{i=1}^{N_c} i_i \vec{B}_i^0$ , where  $\vec{B}_i^0$  is the reference field generated by each coil given a reference current  $I_i^0$ ,  $N_c$  is the number of independent coils, and  $i_i = I_i/I_i^0$  is the current ratio between the actual and reference coil current. The set of current ratios  $\vec{i} \in \mathbb{R}^{N_c}$  additionally parameterizes the equilibrium problem.

In this work, since in equilibrium reconstruction routines the plasma boundary is not known *a priori*, we are only interested in free-boundary equilibria. In addition, no diagnostics to probe the iota profile are installed at W7-X, whereas continuous Rogowski coils measure the net toroidal current. Therefore, the toroidal current profile is specified instead of the iota profile.

## 2.3. The ideal-MHD solution operator

In this work, a set of scalar parameters ( $i_i$  and  $\Phi_{\text{edge}}$ ) and plasma profiles ( $p$  and  $I_{\text{tor}}$ ), represented by one-dimensional functions of the radial coordinate  $s$ , define an instance of the ideal-MHD equilibrium problem. The equilibrium solution is represented by the mapping  $(s, \theta, \varphi) \rightarrow \vec{x}(s, \theta, \varphi)$  (where  $\vec{x} = (R, \lambda, Z)$ ) and by the one-dimensional function  $s \rightarrow \iota(s)$ .

Let us define the solution operator as the operator that maps the set of scalar parameters and plasma profiles to the MHD solution:

$$\mathcal{O}(\vec{i}, \Phi_{\text{edge}}, p(s), I_{\text{tor}}(s))(s, \theta, \varphi) = [\vec{x}(s, \theta, \varphi), \iota(s)]^T, \quad (16)$$

where  $\vec{i} \in \mathbb{R}^{N_c}$ ,  $\Phi_{\text{edge}} \in \mathbb{R}$ ,  $\{p, I_{\text{tor}}, \iota\} : [0, 1] \rightarrow \mathbb{R}$  are one-dimensional functions defined on the radial plasma domain, and  $\vec{x} : [0, 1] \times [0, 2\pi] \times [0, 2\pi] \rightarrow \mathbb{R}^3$  is the mapping between flux coordinates and  $(R, \lambda, Z)$ .



### 3. NNs to approximate non-linear operators

NNs can approximate non-linear continuous functions [24–28], discontinuous functions [29, 30], and non-linear continuous operators [31–33]. In this work, we propose a NN model to learn the non-linear ideal-MHD operator. A set of free parameters  $\Theta \in \mathbb{R}^N$  (i.e. trainable weights) parameterizes the NN model:

$$\text{NN}(\vec{i}, \Phi_{\text{edge}}, p(s), I_{\text{tor}}(s); \Theta)(s, \theta, \varphi) = [\vec{x}(s, \theta, \varphi), \tilde{\epsilon}(s)]^T, \quad (17)$$

where  $\vec{x}(s, \theta, \varphi)$  and  $\tilde{\epsilon}(s)$  compose the equilibrium solution provided by the NN model.

Three sources of error characterize the accuracy of a NN: approximation, optimization and generalization errors [34]. The approximation error describes how well a NN represents the target function or operator, the optimization error describes how well the NN training procedure minimizes the training loss (i.e. empirical risk), and the generalization error describes how well the NN generalizes to unseen data.

The approximation theorem in [31] shows how any NN with a specific structure and enough capacity (i.e. number of free parameters) guarantees a small approximation error. However, no guarantees are given on the optimization and generalization errors. Also, from a practical standpoint, estimating the required number of free parameters to meet “enough capacity” is difficult. Fulfilling known symmetries of the solution by construction should ease the training process (i.e. it should reduce the optimization error), and the addition of a physics constraint loss should improve generalization (i.e. it should reduce the generalization error). The next section will introduce them separately.

#### 3.1. Imposing symmetries in the NN model

In a learning algorithm, the available computational budget is divided into three baskets: the capacity (i.e. size) of the model sets the approximation error, the length of the training affects the optimization error, and the size of the data set impacts the generalization error.

The space of possible functions that the model can represent is large. The training process is essentially a ‘search’ in this space. An inductive bias ‘allows a learning algorithm to prioritize one solution over another, independent of the observed data’ [35]. An inductive bias encodes the solution’s prior knowledge (e.g. a solution’s symmetry); it allows budget to be relocated from the capacity basket to the other two baskets, thus, reducing the optimization and generalization errors.

The equilibrium solution should satisfy multiple symmetries. Let us define  $(X^1, X^2, X^3) = (R, \lambda, Z)$ . Each quantity  $X^j$  (with  $j \in \{1, 2, 3\}$ ) should be a periodic function of the poloidal and toroidal angles:

$$X^j(s, \theta, \varphi) = X^j(s, \theta + 2\pi, \varphi), \quad (18)$$

$$X^j(s, \theta, \varphi) = X^j(s, \theta, \varphi + 2\pi/N_{\text{fp}}), \quad (19)$$

where  $N_{\text{fp}}$  is the number of toroidal field period.

Moreover, if stellarator symmetry is assumed (which is the case for most experimental devices, including W7-X), the solution should satisfy the following relationships [36]:

$$R(s, \theta, \varphi) = R(s, -\theta, -\varphi), \quad (20)$$

$$\lambda(s, \theta, \varphi) = -\lambda(s, -\theta, -\varphi), \quad (21)$$

$$Z(s, \theta, \varphi) = -Z(s, -\theta, -\varphi). \quad (22)$$

For example, the periodicity of  $X^j$  and stellarator symmetry can be imposed if the solution is expanded in Fourier series along the poloidal and toroidal directions:

$$R(s, \theta, \varphi) = \sum_{mn} R_{mn}(s) \cos(m\theta - N_{\text{fp}}n\varphi), \quad (23)$$

$$\lambda(s, \theta, \varphi) = \sum_{mn} \lambda_{mn}(s) \sin(m\theta - N_{\text{fp}}n\varphi), \quad (24)$$

$$Z(s, \theta, \varphi) = \sum_{mn} Z_{mn}(s) \sin(m\theta - N_{\text{fp}}n\varphi), \quad (25)$$

where  $X_{mn}^j$  are the associated Fourier coefficients,  $m$  is the poloidal mode number, and  $n$  is the toroidal mode number. Since  $\lambda$  is a periodic function with zero average,  $\lambda_{00} = 0$ .

At the magnetic axis ( $s = 0$ ), the flux surface geometry must be regular (i.e. infinitely differentiable).

If it is expanded in Fourier series, the Fourier coefficients  $X_{mn}^j$  must have the form [37]:

$$X_{mn}^j = \sqrt{s^m} \sum_{k \in \mathbb{N} \cup \{0\}} X_{kmn}^j s^k. \quad (26)$$

Section 3.2 introduces how the NN model inherently satisfies this property.

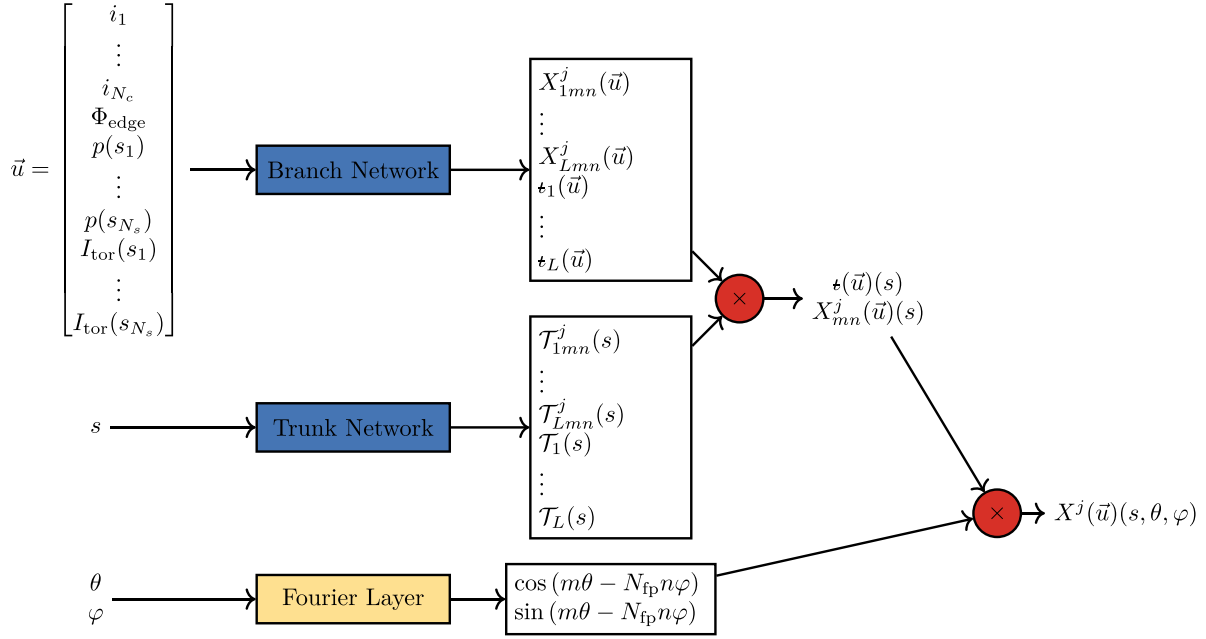
#### 3.2. DeepONet-like architecture

Figure 1 shows the NN macro architecture. The plasma profiles are observed at a set of equally spaced radial locations,  $s_i$  for  $i \in \{1, \dots, N_s\}$ , also called sensors. In addition to the independent variables ( $s$ ,  $\theta$ , and  $\varphi$ ), the model has two sets of inputs: the set of scalar equilibrium parameters ( $[i_1, \dots, i_{N_c}]^T$  and  $\Phi_{\text{edge}}$ ) and the plasma profiles ( $[p(s_0), \dots, p(s_{N_s})]^T$  and  $[I_{\text{tor}}(s_0), \dots, I_{\text{tor}}(s_{N_s})]^T$ ). Those inputs are concatenated into  $\vec{u} = [i_1, \dots, i_{N_c}, \Phi_{\text{edge}}, p(s_1), \dots, p(s_{N_s}), I_{\text{tor}}(s_1), \dots, I_{\text{tor}}(s_{N_s})]^T$ . The model represents the MHD equilibrium solution as:

$$X^j(\vec{u})(s, \theta, \varphi) = \sum_{mn} X_{mn}^j(\vec{u})(s) \mathcal{F}_{mn}(\theta, \varphi), \quad (27)$$

$$X_{mn}^j(\vec{u})(s) = \sqrt{s^m} \sum_{l=1}^L X_{lmn}^j(\vec{u}) \mathcal{T}_{lmn}^j(s), \quad (28)$$

$$\epsilon(\vec{u})(s) = \sum_{l=1}^L \epsilon_l(\vec{u}) \mathcal{T}_l(s), \quad (29)$$



**Figure 1.** The macro architecture of the NN model to approximate the ideal-MHD solution operator. In case of free-boundary equilibria, a set of parameters (the normalized coil currents  $[i_1, \dots, i_{N_c}]^T$  and the toroidal flux enclosed by the plasma  $\Phi_{\text{edge}}$ ) and two functions (the pressure  $p(s)$  and toroidal current  $I_{\text{tor}}(s)$  plasma profiles) define the ideal-MHD equilibrium problem. The variable  $u$  combines the set of parameters and the plasma profiles values at selected flux surfaces,  $[s_1, \dots, s_{N_s}]^T$ . The solution is represented by the set of flux surface coordinates  $(R, Z)$ , the renormalization parameter  $\lambda$ , and the iota profile  $\iota$ . The solution domain spans the radial  $s$ , poloidal  $\theta$ , and toroidal  $\varphi$  directions. The quantities  $X^j \in \{R, \lambda, Z\}$  are expanded in Fourier series along the poloidal and toroidal directions. A trunk network constructs a set of non-linear ‘basis’ functions of  $s$ , which radially expand the  $X^j$  Fourier coefficients and iota. A branch network parameterizes the set of trunk basis coefficients as a function of  $u$ .

where  $X^j \in \{R, \lambda, Z\}$ ,  $\mathcal{F}_{mn}(\theta, \varphi) = \cos(m\theta - N_{\text{fp}}n\varphi)$  in case of  $R$  and  $\mathcal{F}_{mn}(\theta, \varphi) = \sin(m\theta - N_{\text{fp}}n\varphi)$  in case of  $\lambda$  or  $Z$ . The Fourier terms  $\mathcal{F}_{mn}(\theta, \varphi)$  expand each  $X^j$ , where  $X_{mn}^j$  represent the Fourier coefficients.

Up to  $m = 8$  poloidal and  $|n| = 11$  toroidal modes represent each  $X^j$ . For W7-X, previous works showed that up to six modes are sufficient to represent the geometry of the flux surfaces [18–20]. However, using VMFC equilibria with up to  $m = 11$  and  $|n| = 12$  as a ‘high-fidelity’ reference, a Fourier resolution scan revealed that up to  $m = 8$  and  $|n| = 11$  modes are required to represent the equilibrium field with a field error below 1% (see section ‘Equilibria Fourier scaling’).

For each Fourier coefficients and for iota, a ‘trunk’ network parameterizes a set of non-linear functions of  $s$ ,  $[\mathcal{T}_{1mn}^j(s), \dots, \mathcal{T}_{Lmn}^j(s)]^T \in \mathbb{R}^L$  and  $[\mathcal{T}_1(s), \dots, \mathcal{T}_L(s)]^T \in \mathbb{R}^L$ , respectively. A ‘branch’ network outputs a set of scalar coefficients as a function of  $u$ ,  $[X_{1mn}^j, \dots, X_{Lmn}^j]^T \in \mathbb{R}^L$  and  $[\iota_1, \dots, \iota_L]^T \in \mathbb{R}^L$ . Finally, the iota and the Fourier coefficients are obtained as a linear combination of trunk network functions, which are weighted by the branch network coefficients.

Iota highly depends on the toroidal current [38]. To let the model exploit this dependency, the toroidal current profile substitutes the last trunk network function in the expansion of the iota profile:  $\mathcal{T}_L(s) = I_{\text{tor}}(s)$ .

The trunk network expands the radial coordinate  $s$  into a set of ‘basis’ functions, independent for each Fourier coefficient  $X_{mn}^j$  and iota  $\iota$ . We inappropriately refer to them as a set of basis functions on the interval  $[0, 1]$  since these functions

are not guaranteed to form a basis (a set of orthogonal functions). To obey the required form needed to represent  $X^j$  (see section 3.1), each function<sup>1</sup> is multiplied by a  $\sqrt{s}^m$  factor, and the two first trunk network functions are the constant and identity functions, respectively:  $\mathcal{T}_{1mn}^j(s) = \mathcal{T}_1(s) = 1$  and  $\mathcal{T}_{2mn}^j(s) = \mathcal{T}_2(s) = s$ . Imposing the identity function as one of the trunk network functions effectively bypasses the trunk network layers, and it preserves the model ability to build simple functions (e.g. the identity function). This construction is known as a ‘skip connection’ in machine learning [39]. A multilayer perceptron (MLP) parameterizes the trunk network. Please refer to section ‘Model architecture and hyperparameters (HPs)’ for a detailed description of the trunk network architecture.

The branch network maps the scalar parameters and the plasma profiles of the equilibrium problem to the set of coefficients associated to the trunk functions. The one-dimensional plasma profiles are observed at a set of sensors  $[s_1, \dots, s_{N_s}]^T$ . A set of one-dimensional convolutional layers extract high-level features from the plasma profiles into a compressed latent representation, which is then concatenated to the set of scalar parameters. Finally, a MLP outputs the coefficients associated to the trunk functions. Please refer to section ‘Model architecture and HPs’ for a detailed description of the branch network architecture.

<sup>1</sup> Only the functions that expand the Fourier coefficients  $X_{mn}^j$ .

The model can represent a global ideal-MHD solution (i.e. a solution that is defined everywhere in the plasma volume) and fulfils the required symmetries and regularity condition at the magnetic axis by construction. Because of the Fourier expansion along the poloidal and toroidal directions, the periodic boundary conditions and stellarator symmetry are satisfied. The required radial form of each Fourier coefficients is satisfied due to the  $\sqrt{s^m}$  factor and the trunk network skip connection.

## 4. Training

The model is trained in a supervised-learning fashion: VMEC provides the ground-truth ideal-MHD equilibria. To restrict the space in which to approximate the ideal-MHD solution operator, only W7-X configurations are considered.

### 4.1. Data set

The boundaries of the training data limit the application of any data-driven method. In this work, the training data are sampled from a multivariate distribution that approximates the experimental distribution (table 1). The coil current ratios define the vacuum magnetic configuration, the plasma volume determines the total toroidal flux at the edge. One-dimensional Gaussian processes (GPs) on the  $[0, 1]$  interval represent the normalized pressure and toroidal current profiles. The volume averaged plasma beta informs the scale of the pressure profile, and the total integrated toroidal current defines the scale of the current profile. See [20, section 2.2] for a more in depth description of the methodology.


- **Vacuum configuration.** The current ratios between the W7-X coils define the vacuum magnetic configuration. Two symmetric sets of coils compose each half toroidal field period: five non-planar coils, two planar (tilted) coils, and one control coil constitute each set. Defining  $I_1$  as the current flowing in the first non-planar coil,  $i_i = \frac{I_i}{I_1}$  for  $i \in \{1, \dots, N_c\}$  are the coil current ratios. In the data set, the current of the first non-planar coil is held fixed at  $I_1 = 13\,068$  A (current value to obtain  $B_{\text{axis}}(\varphi = 0) \simeq 2.52$  T in case of the standard configuration), and the current in the control coil is null. The coil current ratios are uniformly sampled to cover a broad configuration space around the nine reference W7-X configurations [8].
- **Plasma volume.** For a given magnetic field strength,  $\Phi_{\text{edge}}$  sets the plasma volume:  $\Phi_{\text{edge}} = \int_0^1 \int_0^{2\pi} \vec{B} \cdot \vec{\nabla} \varphi \sqrt{g} ds d\theta \propto Ba\pi^2 \simeq BV_p/R$ , where  $a$  is the effective minor radius,  $R$  the major radius, and  $V_p$  the plasma volume. The plasma volume is sampled from a normal distribution with mean  $30\text{ m}^3$  and standard variation  $3\text{ m}^3$  [40]. Once a plasma volume has been sampled, a linear fit estimates  $\Phi_{\text{edge}}$ .
- **Plasma beta.** For a given magnetic field strength and pressure profile shape, the pressure on axis sets the volume averaged plasma beta:  $\langle \beta \rangle = 2\mu_0 \langle p/B^2 \rangle_{\text{vol}} = 2\mu_0 p_0 \langle \bar{p}/B^2 \rangle_{\text{vol}}$ , where  $p_0$  is the pressure on axis and  $\bar{p} = p/p_0$  is the normalized pressure profile. In previous operational campaigns,

**Table 1.** Distributions from which the training data have been sampled.  $U(a, b)$  represents a uniform distribution on the  $[a, b]$  interval, and  $N(\mu, \sigma)$  represents a normal distribution with  $\mu$  mean and  $\sigma$  standard variation. One-dimensional GPs with a non-stationary covariance function model the plasma profiles. The GP covariance function scale factor  $\sigma_f$ , core length scale  $l_{\text{core}}$ , edge length scale  $l_{\text{edge}}$ , core-edge transition location  $s_0$ , and core-edge transition width  $s_w$  parameterize each GP [20, section 2.2]. The references from which the distribution parameters have been informed are listed.

Quantity magnetic configuration	Distribution	Unit	References
Non-planar coil currents	$U(0.6, 1.2)$	—	[8, 18–20]
Planar coil currents $i_{[A,B]}$	$U(-0.8, 0.8)$	—	[8, 18–20]
Plasma parameters			
Plasma volume $V_p$	$N(30, 3.0)$	$\text{m}^3$	[40]
Plasma beta $\beta$	$U(0, 5)$	%	[41]
Total toroidal current $I_{\text{tor}}$	$U(-20, 20)$	kA	[40]
Pressure profile GP			
Covariance function scale factor $\sigma_f$	$U(0.1, 5.0)$	—	[20]
Core length scale $l_{\text{core}}$	$U(0.1, 1.0)$	—	[20]
Edge length scale $l_{\text{edge}}$	$U(0.1, 2.0)$	—	[20]
Core-edge transition location $s_0$	$U(0.1, 1.0)$	—	[20]
Core-edge transition width $s_w$	$U(0.1, 0.3)$	—	[20]
Toroidal current profile GP			
Covariance function scale factor $\sigma_f$	$U(1.0, 5.0)$	—	[20]
Core length scale $l_{\text{core}}$	$U(0.1, 1.0)$	—	[20]
Edge length scale $l_{\text{edge}}$	$U(0.5, 1.0)$	—	[20]
Core-edge transition location $s_0$	$U(0.001, 0.8)$	—	[20]
Core-edge transition width $s_w$	$U(0.1, 0.5)$	—	[20]

W7-X has reached a maximum  $\langle \beta \rangle$  of roughly 1%. However, W7-X has been optimized for higher plasma beta (up to  $\langle \beta \rangle \simeq 5\%$  [7]). Therefore, the volume averaged plasma beta is sampled from a uniform distribution between 0% and 5%. Once a plasma beta has been sampled, a linear fit estimates  $p_0$ .

- **Net toroidal current.** The plasma bootstrap current and the externally induced current (e.g. via electron cyclotron current drive (ECCD)) mainly define the net toroidal current. The net toroidal current should be minimized for the robust operation of an island divertor. In W7-X operations, two current-free scenarios are envisioned [40]: a scenario in which the bootstrap current is minimized via the magnetic configuration, and a scenario in which the bootstrap current is balanced by a strong ECCD current. During the first operational campaign, the measured bootstrap current ranged between  $-7\text{ kA}$  and  $17\text{ kA}$  [42]. Therefore, the total toroidal current is sampled from a uniform distribution between  $-20\text{ kA}$  and  $20\text{ kA}$ .



**Figure 2.** Histograms of the volume averaged beta (*a*), the pressure peaking factor (*b*), and the peak toroidal current density (*c*) for all converged equilibria in the data set.

- Pressure profile shape.** For electrons and ions, the gradient of the density and temperature ( $p \propto nT$ ) profiles can be substantially different in the core and edge regions. To account for such difference, a GP with a non-stationary covariance function models the pressure profile. The GP is parameterized by a set of HPs [20, section 2.2]: the covariance function scale factor  $\sigma_f$ , the core length scale  $l_{\text{core}}$ , the edge length scale  $l_{\text{edge}}$ , the transition location  $s_0$ , and the transition width  $s_w$ . In W7-X, the pressure profile is generally flat, however, high-performance operations are correlated with a peaked density profile [43]. To account for both, flat and peaked profiles are included in the data set. Profiles with a non-physical positive pressure gradient (i.e.  $\nabla p > 0$ ) are discarded.
- Toroidal current profile shape.** In W7-X, the toroidal current has two main contributions: the ECCD induced current and the bootstrap current. The bootstrap current density is generally flat across the plasma volume, however, the ECCD current density is peaked at the deposition location. To capture both shapes, a GP with a non-stationary covariance function is used to model the toroidal current profile.

14281 equilibria have been sampled in the data set, but equilibria for which VMEC was not able to minimize the MHD variational forces up to a tolerance of  $1 \times 10^{-7} \text{NHm}^{-1}$

were discarded<sup>2</sup>. The data set thus contains 9857 equilibria. Figures 2(a)–(c) show the histograms of a subset equilibrium properties for all converged equilibria: the volume averaged beta  $\langle \beta \rangle$ , the pressure peaking factor  $p_0/\langle p \rangle$ , and the peak toroidal current density  $\max_s |j_{\text{tor}}(s)|$ . The  $\langle \beta \rangle$  distribution is not uniform due to the not converged equilibria at high beta (figure 2(a)). Nevertheless, the data set holds both flat (i.e. small peaking factor) and peaked (i.e. large peaking factor) pressure profiles (figure 2(b)).

The data set has been split into a training set (80%), a validation set (10%), and a test set (10%). The training set has been used to train the model, the validation set has been used to assess the best model HPs, and the test set has been used to assess the model performance on held-out data.

To improve the training procedure, the model inputs have been normalized with the mean and standard deviation computed on 20% of the training data.

#### 4.2. Loss function

Three terms form the model loss function: a data term, a gradient term, and a physics-based regularization term. The data

<sup>2</sup>  $\text{ftol} = 1 \times 10^{-14}$  in VMEC.

term biases the model to provide equilibria that are close to the ground truth (computed by VMEC), the gradient term regularizes the model equilibria to exhibit the same radial derivatives as the ground truth, and the physics-based term regularizes the model equilibria to satisfy the ideal-MHD equations.

**4.2.1. Data loss.** The model equilibrium solution is globally defined in the whole plasma volume. However, to numerically compute a residual measure against the ground truth, an equally spaced grid is adopted:  $N_s = 99$  flux surfaces,  $N_\theta = 32$  poloidal locations, and  $N_\varphi = 36$  toroidal locations compose the loss grid. The iota profile, being a flux surface quantity, is evaluated only at the radial grid locations. On each equilibrium, the data loss term is:

$$L_d = \sum_{j=1}^3 \alpha_{X^j} \sum_{i=1}^{N_s} \sum_{t=1}^{N_\theta} \sum_{z=1}^{N_\varphi} [\tilde{X}^j(s_i, \theta_t, \varphi_z) - X^j(s_i, \theta_t, \varphi_z)]^2 + \alpha_t \sum_{i=1}^{N_s} [\tilde{t}(s_i) - t(s_i)]^2, \quad (30)$$

where  $\tilde{X}^j$  and  $\tilde{t}$  are the model outputs, and  $\alpha_{X^j}$  and  $\alpha_t$  are the set of coefficients that weight the data loss terms.

**4.2.2. Gradient loss.** An equilibrium solution is fully defined by the mapping  $X^j(s, \theta, \varphi)$  for  $j \in \{1, 2, 3\}$  and by  $t(s)$ . However, many equilibrium properties depend also on equilibrium first and second radial, poloidal and toroidal derivatives. For example, the flux surface Jacobian  $\sqrt{g}$ , which is the local linear approximation of the coordinate transformation  $(s, \theta, \varphi) \rightarrow (R, \phi, Z)$ , is:

$$\sqrt{g} = R \left( \frac{\partial R}{\partial \theta} \frac{\partial Z}{\partial s} - \frac{\partial R}{\partial s} \frac{\partial Z}{\partial \theta} \right). \quad (31)$$

To faithfully reconstruct the equilibrium properties, the gradients of the equilibrium solution need to match the gradients of the ground truth. Due to the Fourier expansion of the solution in the poloidal and toroidal directions, the poloidal and toroidal derivatives are already constrained by the Fourier coefficients: for example,  $\frac{\partial R}{\partial \theta} = \frac{\partial}{\partial \theta} [\sum_{mn} R_{mn} \cos(m\theta - N_{fp}n\varphi)] = -\sum_{mn} m R_{mn} \sin(m\theta - N_{fp}n\varphi)$ . A similar relationship holds for any poloidal and toroidal derivatives of  $(R, \lambda, Z)$ .

However, the radial derivatives of the equilibrium solution are not guaranteed to match those of the ground truth. The radial derivatives depend on the learnt trunk network functions, which are derived during training to minimize a  $L^2$  residual with the ground truth on a finite-spacing radial grid. Even if the  $L^2$  residual vanishes, numerical artefacts (e.g. numerical oscillations) are known to occur [44]. Therefore, the radial derivatives of the model solution can substantially differ from the ground truth.

To ameliorate these phenomena, a gradient term regularizes the radial derivatives of the model equilibrium:

$$L_g = \sum_{j=1}^3 \alpha_{X^j} \sum_{i=1}^{N_s} \sum_{t=1}^{N_\theta} \sum_{z=1}^{N_\varphi} \left[ \frac{\partial \tilde{X}^j}{\partial s}(s_i, \theta_t, \varphi_z) - \frac{\partial X^j}{\partial s}(s_i, \theta_t, \varphi_z) \right]^2 + \alpha_{t'} \sum_{i=1}^{N_s} \left[ \frac{\partial \tilde{t}}{\partial s}(s_i) - \frac{\partial t}{\partial s}(s_i) \right]^2, \quad (32)$$

where  $\alpha_{X^j}$  and  $\alpha_{t'}$  are a set of coefficients that weight the gradient loss terms. The radial derivatives are approximated with a central finite difference scheme for both the model and ground truth solutions.

**4.2.3. Physics regularization.** The ideal-MHD equilibrium problem exhibits a high condition number [45]. In the framework of MHD equilibria, the condition number is defined as ratio between the largest and the smallest eigenvalue of the MHD force residual linearized around an equilibrium solution. Due to the presence of second-order radial derivatives in the residual, the condition number scales with the square of the number of radial grid locations [45], namely,  $\mathcal{P} \sim N_s^2$  ( $\mathcal{O}(10^4)$  in this work).

A high condition number means that the MHD force residual is highly sensitive to the equilibrium solution: small deviations of the solution away from the ground truth may lead to a large MHD residual. Even if the model equilibria are close (in the  $L^2$  sense) to the ground truth equilibria, they are not guaranteed to well minimize the residual.

The MHD force residual in the loss function regularizes the model to provide equilibria that better satisfy the ideal-MHD equations. For an equilibrium solution,  $\|\vec{F}\|_2 = 0$  everywhere in the domain. The volume-averaged residual norm can be computed as:

$$\langle \|\vec{F}\|_2 \rangle_{\text{vol}} = \int_0^1 \int_0^{2\pi} \int_0^{2\pi} \|\vec{F}\|_2 \sqrt{g} ds d\theta d\varphi, \quad (33)$$

where  $\langle A \rangle_{\text{vol}} = \int_0^1 \int_0^{2\pi} \int_0^{2\pi} A \sqrt{g} ds d\theta d\varphi$  denotes a volume-averaged operator.

In this work, a proxy for the MHD force residual is used instead. VMEC (the code that provides the ground truth equilibria) does not directly minimize the force residual, instead, it minimizes the total energy of the system (i.e. it uses a variational approach). It has been shown that equilibria computed by VMEC only poorly minimize the ideal-MHD force residual [2]. Therefore, adding the full  $\langle \|\vec{F}\|_2 \rangle_{\text{vol}}$  as regularization term may cause instabilities in the training procedure. The volume averaged radial force balance assuming a vanishing helical component ( $F_\beta = 0$ ) is used instead. This quantity is also used in VMEC to determine the goodness of a converged solution. Let us firstly define

$$f_* = \frac{\mu_0}{(2\pi)^2} \langle F_s |_{F_\beta=0} \rangle_{F_s}, \quad (34)$$

where  $\langle A \rangle_{F_s} = \int_0^{2\pi} \int_0^{2\pi} A \sqrt{g} d\theta d\varphi$  is a flux surface averaged operator. Then, the model is regularized with the following term:

$$L_{\text{MHD}} = \alpha_{\text{MHD}} \int_{s_{\min}}^1 f_*^2 ds, \quad (35)$$



where  $s_{\min} = 0.02$  to avoid the sensitivity of the MHD force residual at the magnetic axis. Please refer to section ‘Physics regularization’ for the derivation of  $f_*$ .

#### 4.3. Training stages

Three subsequent training stages compose the training procedure. In every training stage, early stopping is employed during training [46], and the best model accordingly to the validation loss is used as a warm start for the next stage.

- **Data stage.** At the beginning of the training, only the data loss  $L_d$  is used. A set of flux surfaces randomly sampled across all equilibria in the training set compose each batch (i.e. individual flux surfaces from different equilibria are randomly shuffled into the same batch).
- **Gradient stage.** In this stage, the gradient loss  $L_g$  is included in the overall loss. To compute the radial derivatives of the solution, the flux surfaces of the same equilibrium are aggregated together in the same batch. However, multiple equilibria are still randomly shuffled into each batch.
- **Physics regularization stage.** Finally, the physics regularization term  $L_{\text{MHD}}$  is included in the loss function. As in the gradient stage, and to compute  $f_*$ , the flux surfaces of the same equilibrium are aggregated together in the same batch. Still, multiple equilibria form each batch.

During the physics regularization stage, a curriculum learning approach is used [47]. The high condition number limits the convergence rate and imposes a superior bound on the learning rate (i.e. the step size): even if the model commits a finite, but small error on the equilibrium solution, the regularization term might lead to a diverging loss.

To mitigate this issue, the model is gradually regularized, starting with well-predicted equilibria. At every batch, the physics regularization term is included into the loss function only for the subset of equilibria where the mean-squared error (mse) on  $(R, Z)$  (the key quantities affecting the MHD residual) is below a given threshold. When the regularization term has not decreased in the previous 50 epochs, the threshold is progressively increased by a factor of 10. More formally, the physics regularization term for the generic batch  $\mathcal{B}$  at the  $k$ th step is:

$$L_{\text{MHD},\mathcal{B}}^k = \frac{1}{|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} w_i^k L_{\text{MHD},i}, \quad (36)$$

$$w_i^k = H(L^k - L_{(R,Z),i}), \quad (37)$$

where  $L_{\text{MHD},i}$  is the physics regularization term for the  $i$ th equilibrium in the batch,  $H(\cdot)$  is the Heaviside function,  $L^k$  is the loss threshold value at the  $k$ th step,  $L_{(R,Z),i}$  is the mse on  $(R, Z)$  for the  $i$ th equilibrium. The loss threshold is progressively increased, i.e.  $L^k = 10^k L^0$ . An initial threshold of  $L^0 = 3 \times 10^{-7} \text{ m}^2$  is used. At the fourth step, the physics regularization term is computed on all the equilibria in the batch, namely,  $L^4 = +\infty$ .

#### 4.4. HPs

Due to the computational cost of the whole training procedure ( $\simeq 8$  days on a single NVIDIA Tesla-V100 graphical processing unit (GPU)), a simplified scenario is used to identify the best model HPs. Given a fixed budget of  $\simeq 6\text{M}$  free parameters, HP grid search is performed for the depth and width of both trunk and branch networks. In addition, only half of the training data is used. The optimum HPs combination is chosen based on the loss on the validation set in the data stage.

The trunk and branch networks are then proportionally upscaled to reach  $\simeq 40\text{M}$  free parameters. Tables 7 and 8 summarize the HPs of the final model.

The AdamW [48, 49] optimizer is used to train the model. In each stage, due to the different loss function, the learning rate, the learning rate scheduler, and the  $L_2$  regularization term are derived with a limited grid search. The loss on the validation set is used to inform the search. In addition, early stopping [46] is employed during training. Tables 9–11 report the HPs used in each training stage.

## 5. Results

In the following chapter, the model’s accuracy is evaluated by how well it predicts the equilibrium problem solution and reconstructs equilibrium properties of interest (which are not explicitly included in the loss function) across many physics domains (table 2): flux surface geometry (section 5.1); iota profile, shear profile and magnetic field structure (section 5.2); MHD force residual (section 5.3); ideal-MHD stability (section 5.4); neoclassical transport (section 5.5); fast particle confinement (section 5.6).

The root-mean-square error (rmse) and mean absolute percentage error (mape) summarize the error on the test set. Given a quantity of interest  $y \in \mathbb{R}^K$  (e.g. the magnetic well), the rmse evaluates the average residual between the ground truth and predicted values for that quantity. If  $\mathcal{Y} = \{y \in \mathbb{R}^K\}$  is the set of true values, and  $\hat{\mathcal{Y}} = \{\hat{y} \in \mathbb{R}^K\}$  is the set of related model predictions, the rmse on  $y$  is:

$$\text{rmse}_y = \sqrt{\frac{1}{K|\mathcal{Y}|} \sum_{i=1}^{|\mathcal{Y}|} \sum_{k=1}^K (\hat{y}_i^k - y_i^k)^2}. \quad (38)$$

Given a quantity of interest  $y \in \mathbb{R}^K$ , the mape measures the average relative deviation of ground truth with respect to predicted values for that quantity. Using the same notation as for the rmse, the mape on  $y$  is:

$$\text{mape}_y = \frac{1}{K|\mathcal{Y}|} \sum_{i=1}^{|\mathcal{Y}|} \sum_{k=1}^K \left| \frac{\hat{y}_i^k - y_i^k}{y_i^k} \right|. \quad (39)$$

To further investigate the model accuracy, and to validate its use in downstream applications, section 5.7 shows the *a posteriori* optimization of W7-X coil currents to search for ideal-MHD unstable and improved fast particle confinement configurations. In this task, the model replaces a traditional



**Table 2.** Model error metrics on the test set across multiple equilibrium properties (which are not explicitly included in the loss function). For the description of each quantity, please see sections 5.1–5.6. Each row represents either an error metric (e.g. the flux surface error  $\epsilon_{Fs}$ ) or an equilibrium property (e.g. the proxy for the radial force balance  $f_*$ ). In both cases, the table shows the average over the whole test set. In case of an equilibrium property, the bracket indicates the ground truth value as computed by VMEC.

Flux surfaces geometry	
Flux surface error $\epsilon_{Fs}$ (mm)	$6.16 \times 10^{-1}$
$R_{axis}$ rmse (mm)	$5.84 \times 10^{-1}$
$\lambda$ rmse (mrad)	$3.90 \times 10^0$
Magnetics	
Iota mape	$9.31 \times 10^{-4}$
Shear mape	$8.44 \times 10^{-1}$
Mean shear mape	$8.16 \times 10^{-2}$
Magnetic field rmse (mT)	$2.36 \times 10^1$
Volume averaged plasma beta mape	$1.37 \times 10^{-3}$
Ideal-MHD force balance	
MHD force residual proxy $\epsilon_{f_*}$ (VMEC) ( $\text{NHm}^{-1}$ )	$1.97 \times 10^{-2}$ ( $2.43 \times 10^{-5}$ )
Normalized MHD force residual proxy $\eta_{f_*}$ (VMEC)	$5.01 \times 10^{-1}$ ( $5.29 \times 10^{-4}$ )
Ideal-MHD stability	
Vacuum magnetic well depth $\bar{W}_{vacuum}$ mape	$3.43 \times 10^{-2}$
Finite- $\langle\beta\rangle$ magnetic well depth $\bar{W}$ mape	$2.61 \times 10^{-1}$
Vacuum magnetic well $W_{vacuum}$	$5.12 \times 10^{-1}$
Finite- $\langle\beta\rangle$ magnetic well $W$ mape	$1.93 \times 10^0$
Mercier well $D_W$ mape	$1.12 \times 10^0$
Neoclassical transport	
Effective ripple $\epsilon_{eff}$ ( $s \leq 0.33$ ) mape	$5.46 \times 10^{-1}$
Effective ripple proxy $\hat{\epsilon}_{eff}$ mape	$1.33 \times 10^{-1}$
Fast particle confinement	
Standard deviation of the maxima of $B\sigma_{B_{max}}$ rmse (mT)	$9.43 \times 10^0$
Standard deviation of the minima of $B\sigma_{B_{min}}$ rmse (mT)	$6.57 \times 10^0$

non-linear ideal-MHD code (e.g. VMEC) in constructing free-boundary equilibria.

The results of the model on the test set are now presented. Because of the adopted finite differences scheme, numerical noise affects quantities that rely on the radial derivative of  $(R, \lambda, Z)$  at the magnetic axis and on the nearest flux surface; those values are invalid. A note is present in the caption of all affected figures.

### 5.1. Flux surface geometry

The geometry of the flux surfaces  $(R, \lambda, Z)$  and  $t$  represent the solution of the ideal-MHD equilibrium problem: they fully describe the equilibrium magnetic field and its properties. Therefore, the accuracy with which the model predicts  $(R, \lambda, Z)$  and  $t$  bounds the accuracy with which it reconstructs equilibrium properties.

The flux surface geometry is predicted with an exceptional accuracy (figure 3). The flux surface error between the ground truth and the model prediction is evaluated as:

$$\epsilon_{Fs} = \langle \|\Delta R_i, \Delta Z_i\|_2 \rangle_i, \quad (40)$$

where  $\Delta R_i$  and  $\Delta Z_i$  are the differences in flux surfaces relative to the ground truth, and the brackets denote the mean across all equilibria and grid locations. In other words,  $\epsilon_{Fs}$  represents the average absolute distance  $\sqrt{(\Delta R)^2 + (\Delta Z)^2}$  across all equilibria and grid locations. The model achieves  $\epsilon_{Fs} = 616 \mu\text{m}$ . Similarly, the model successfully predicts the angle renormalization parameter  $\lambda$ , achieving  $\text{rmse}_\lambda = 2.41$  mrad.


Low-beta equilibria are better regressed than high-beta equilibria (figure 4). As expected from the distribution of the volume averaged beta in the training data set (see figure 2(a)), the flux surface error increases towards the boundary of the training set (i.e. where the plasma beta increases). No clear patterns are visible in terms of the net toroidal current enclosed by the plasma. The flux surface error does not considerably change when the physics regularization is applied (see table 3).

As the plasma beta increases, the plasma column moves outwards. This effect is called the Shafranov shift, and it is usually measured as  $\Delta/a$ , where  $\Delta = R_{axis}(\beta) - R_{axis}(0)$  and  $a$  is the effective minor radius. The Shafranov shift defines the equilibrium beta limit ( $\Delta/a \simeq 0.5$ ) [50], and it causes an outward movement of the last closed flux surface, which affects the edge magnetic field topology. In W7-X, the minimization of the Shafranov shift was an optimization criterion [51]. Nevertheless, a Shafranov shift of  $\Delta \simeq 5$  cm has been observed in high- $\beta$  operations [52].

The model faithfully predicts this finite-beta effect (figure 5): the axis location at  $\varphi = 0$  is resolved with a relative error of  $\text{mape}_{R_{axis}} < 0.01\%$ . A two-dimensional histogram visualizes the distribution of ground truth and predicted values. The color-bar indicates the counts in each two-dimensional histogram bins. To guide the eye, a red dashed line indicates the case of perfect regression: the more the density of points lies on the line, the more the model provides accurate equilibrium properties. The coefficient of determination  $R^2$  captures the reconstruction quality.

### 5.2. Magnetics

$t$  sets the location of the resonant surfaces in the plasma volume. It also defines the island chain and magnetic edge topology, which affects how the W7-X island divertor operates. Although equilibria with the assumption of nested flux surfaces cannot resolve islands, VMEC equilibria are an



**Figure 3.** The ground truth (solid) and predicted (dashed) flux surfaces geometry,  $(R, Z)$ , for the median predicted equilibrium. Four toroidal locations at  $\varphi = 0^\circ$  (pink),  $\varphi = 18^\circ$  (orange),  $\varphi = 36^\circ$  (green), and  $\varphi = 54^\circ$  (red) are shown. For each toroidal location, ten flux surfaces equally spaced in  $\rho = \sqrt{s}$  are shown. The model approximates the geometry of the flux surfaces with high accuracy: the average distance from the ground truth solution is below 1 mm. To highlight the model prediction, the region in which the flux surfaces have the largest curvature is zoomed in. Indeed, the two curves overlap each other.

important requirement in the investigation of plasma phenomena involving resonant surfaces [53–55].

The model effectively predicts the  $\iota$  profile (figure 6(a)). A relative error of  $\text{mape}_\iota < 0.1\%$  is achieved. In case of the median regressed equilibrium, the predicted and ground truth  $\iota$  profiles overlap (figure 6(b)).

The magnetic shear is the radial derivative of  $\iota$ :

$$\iota' = \frac{d\iota}{d\psi} = \frac{\partial\iota}{\partial s} \frac{\partial s}{\partial\psi} = \frac{1}{\psi_{\text{edge}}} \frac{\partial\iota}{\partial s}, \quad (41)$$


where  $2\pi\psi(s) = \Phi_{\text{edge}}s$  is the toroidal flux. It represents the change of direction of the magnetic field lines from one flux surface to another. The magnetic shear influences multiple

properties of the equilibrium: it affects MHD stability [56, 57], it is inversely proportional to the island width (i.e. large islands are present at low shear) [56], and a low shear leads to long connection lengths allowing efficient divertor operation [58].


The model qualitatively reconstructs the shear profile (figure 7(a)). Figure 7(b) shows that the ground truth and predicted shear profile only partially overlap. Across all equilibria in the test set, the relative error on the magnetic shear is high  $\text{mape}_{\iota'} = 84.4\%$ .

However, the model captures the average magnetic shear  $\iota'_*$  (figure 8):  $\text{mape}_{\iota'_*} = 8.16\%$ , where the average magnetic shear is the shear of the least squares linear fit of the iota profile.

The structure of the magnetic field plays a crucial role in many physics aspects: neoclassical and turbulence transport,



**Figure 4.** Distribution of the flux surface error  $\varepsilon_{Fs}$  as a function of the equilibrium volume averaged beta  $\langle\beta\rangle$  and the net toroidal current enclosed by the plasma  $I_{tor}$  in the test set.



**Figure 5.** Reconstruction of the location of the magnetic axis at  $R_{axis}(\varphi = 0)$  across all equilibria in the test set. The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction.

confinement of energetic particles, plasma instabilities, and deposition of electron cyclotron resonance heating power.

The model correctly reconstructs the magnetic field structure (figure 9). The model achieves a  $rmse_B = 23.6$  mT on all equilibria in the test set. When compared with an average W7-X field of  $B_0 = 2.5$  T, this value represents a relative error below 1%.

However, the magnetic field strength error differs along the radial profile (figure 9). On average, the field strength error is below 10 mT for a considerable fraction of the plasma volume, up to  $s \simeq 0.3$  (i.e.  $\rho \simeq 0.55$ ). The accuracy of the magnetic field

strength in the core region is sufficient to navigate the magnetic configuration space of W7-X to obtain fast particle optimized configurations (see section 5.7.2). Moreover, an accuracy of 10 mT is expected to be enough to perform equilibrium reconstruction routines [60].

The model introduces artificial field ripples in the equilibrium solution (figure 10): at the plasma edge, and also to a less extent close to the magnetic axis, the predicted  $B$  field is not as smooth as the ground truth solution. The flux surface geometry Jacobian  $\sqrt{g}$  and the metric tensor elements  $g_{ij}$  mainly define the magnetic field  $B$  (section 2.1). These quantities are highly sensitive to the radial derivative of the flux surface coordinates  $(R, \lambda, Z)$ . High frequency components in the learned trunk network *basis* functions might be the cause of such field ripples.

### 5.3. Ideal-MHD loss

A high condition number, which has several negative implications, characterizes the ideal-MHD equilibrium problem (see section 4.2.3).

Firstly, a high condition number limits the rate of convergence. During the physics regularization training stage, the weights of the NN model are adjusted to minimize the MHD force residual for the equilibria in the training set. The condition number sets the convergence rate of this gradient descent method. To accelerate the descent, non-linear MHD codes adopt a preconditioning matrix, which decreases the condition number of the problem [45]. In this work, this technique is not available: in general, preconditioning algorithms require an approximation of the Hessian matrix, which is impractical due to the large number of free parameters of the NN model.

Secondly, a high condition number implies that the MHD force residual is highly sensitive to the equilibrium solution. For example, if we define  $\vec{x}_0$  to be that solution that satisfies the MHD force residual  $F_{MHD}(\vec{x}_0) = 0$  and  $\kappa$  as the condition number, then  $F_{MHD}(\vec{x}_0 + \vec{\epsilon}) \simeq \kappa \|\vec{\epsilon}\|_2 \gg 1$ , where  $\vec{\epsilon}$  is a small displacement to the equilibrium solution. Even if the model provides a good approximation of the equilibrium solution, the MHD force residual is not guaranteed to be null.

To quantify how well the equilibria satisfy the MHD equations, two measures are adopted. The proxy and full MHD force residuals:

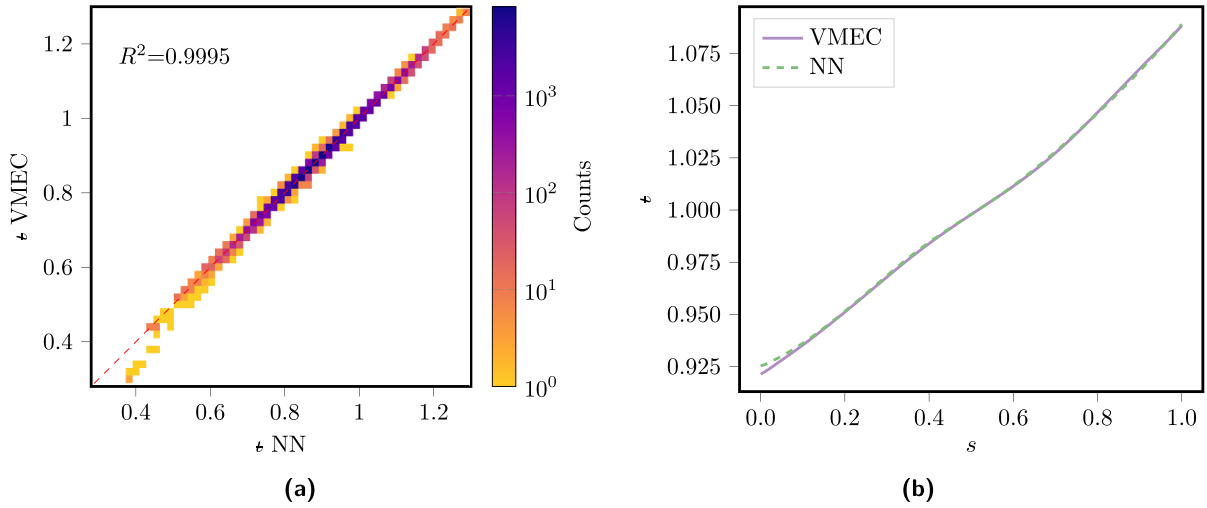
$$\varepsilon_{f_*} = \frac{\mu_0}{(2\pi)^2} \langle F_s |_{F_\beta=0} \rangle_{Vol} = \int_{s_{min}}^1 f_* ds, \quad (42)$$

$$\varepsilon_F = \frac{\mu_0}{(2\pi)^2} \langle \|\vec{F}\|_2 \rangle_{Vol} = \frac{\mu_0}{(2\pi)^2} \int_{s_{min}}^1 \int_0^{2\pi} \int_0^{2\pi} \|\vec{F}\|_2 \sqrt{g} ds d\theta d\varphi. \quad (43)$$

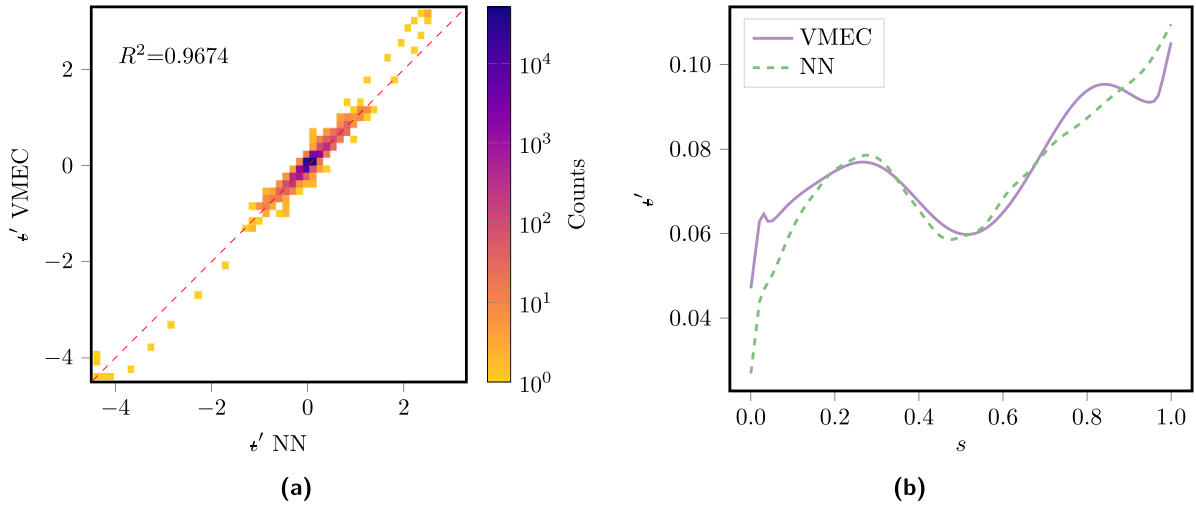
And their normalized versions:

$$\eta_{f_*} = \frac{\int_{s_{min}}^1 f_* ds}{\int_{s_{min}}^1 \int_0^{2\pi} \int_0^{2\pi} \|\vec{\nabla} p\|_2 \sqrt{g} ds d\theta d\varphi}, \quad (44)$$

$$\eta_F = \frac{\int_{s_{min}}^1 \int \|\vec{F}\|_2 \sqrt{g} ds d\theta d\varphi}{\int_{s_{min}}^1 \int_0^{2\pi} \int_0^{2\pi} \|\vec{\nabla} p\|_2 \sqrt{g} ds d\theta d\varphi}. \quad (45)$$



**Figure 6.** (a) Reconstruction of the iota profile  $\epsilon$  across all equilibria in the test set. The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction. (b) The true (solid pink) and predicted (dashed green) iota profile for the median regressed equilibrium in the test set.



**Figure 7.** (a) Reconstruction of the local magnetic shear  $\epsilon'$  across all equilibria in the test set. The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction. (b) The true (solid pink) and predicted (dashed green) magnetic shear profile for the median regressed equilibrium in the set.


The model trained with the physics regularization term better satisfies the ideal-MHD force residual compared with a model that has not been regularized (table 3). Both models show similar flux surface and iota errors, however, the physics regularized model shows a reduced ideal-MHD force proxy residual. However, this number is still three orders of magnitude higher than the value of ground truth equilibria from VMEC.

Even though the model provides the solution of the ideal-MHD equilibrium problem (the geometry of the flux surfaces and the iota profile) with high accuracy, the ideal-MHD force residual is poorly satisfied. In case of the median predicted equilibrium (at  $\langle\beta\rangle = 1.39\%$ ), the normalized MHD force proxy residual is  $\eta_{i_*} = 23.3\%$  (figure 11). For comparison, the ground truth equilibrium has a normalized residual of  $3.04 \times 10^{-2}\%$ .

**5.3.1. Fine-tuning at inference time.** Let us call a *single-pass* equilibrium an equilibrium solution provided by the model with a single forward pass.

*Single-pass* equilibria poorly satisfy the ideal-MHD equations (section 5.3). However, without access to the ground truth solution, equilibria can be improved at inference time. The original free-boundary equilibrium problem is cast into an equivalent fixed-boundary one: the plasma boundary of the *single-pass* free-boundary equilibrium defines the boundary condition on  $(R, Z)$  for the equivalent fixed-boundary equilibrium. The error on  $(R, Z)$  of the *single-pass* equilibrium is below 1 mm (section 5.1), therefore, the plasma boundary can be safely assumed to be close to the correct.

The solution of the equivalent fixed-boundary equilibrium problem is approximated by minimizing the MHD force residual at inference time. The loss function has three terms: the



**Figure 8.** Reconstruction of the average magnetic shear across all equilibria in the test set. The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction. The average magnetic shear is defined as the shear of the least squares linear fit of the iota profile [59].

full MHD force residual, a distance between the model and the fixed boundary  $(R_b, Z_b)$  derived from the *single-pass* equilibrium, and the satisfaction of the required toroidal current profile  $I_{\text{tor}}(s)$ . The loss function is then:

$$L = \alpha_{\text{MHD}} L_{\text{MHD}} + \alpha_b L_b + \alpha_I L_I, \quad (46)$$


$$L_{\text{MHD}} = \frac{\mu_0^2}{(2\pi)^4} \int_{s_{\min}}^1 \langle \|\vec{F}\|_2 \rangle_{\text{Fs}}^2 ds, \quad (47)$$

$$L_b = \langle (\Delta R_b^b)^2 \rangle_i + \langle (\Delta Z_b^b)^2 \rangle_i, \quad (48)$$

$$L_I = \mu_0^2 \int_{s_{\min}}^1 (\tilde{I}_{\text{tor}} - I_{\text{tor}})^2 ds, \quad (49)$$

where  $\alpha_{\text{MHD}} = 1 \times 10^{-3}$ ,  $\alpha_b = 0.2495$ ,  $\alpha_I = 0.5$ , and  $s_{\min} = 0.02$ . The *single-pass* equilibrium is used as initial guess for the minimization problem. The AdamW algorithm minimizes the loss function. The initial learning rate is set to  $10^{-4}$ , and it is decreased to  $10^{-6}$  at  $10^4$  steps with an exponential decay. The training is halted after  $2 \times 10^4$  steps. The branch network, apart from the bias of the last layer, is frozen. For a complete description of the HPs, see section ‘Model architecture and HPs’. We call this optimization at the inference step *fine-tuning*.

In terms of the full MHD force residual, the model can provide better-than-ground-truth equilibria (table 4, figures 12(a) and (b)). As an example, we consider a W7-X standard configuration at  $\langle \beta \rangle = 2\%$ , which is not included in the training data. Both the *single-pass* and fine-tuned equilibrium show low flux surface errors, however, the fine-tuned equilibrium satisfies the ideal-MHD equations much better than the *single-pass* equilibrium: the normalized error is only  $\eta_{\epsilon_*} \simeq 1\%$ . Surprisingly, the fine-tuned equilibrium minimizes



**Figure 9.** The error on the magnetic field strength as a function of the radial location, averaged across all grid points on the flux surface. The solid line represents the mean across all equilibria in the test set, and the shaded regions represent the 95% and 5% quantiles. To guide the eye, two reference values are shown: 25 mT (dashed black) represents a 1% error (W7-X has a field of 2.5 T on average on axis), and 10 mT (dashed red) represents the expected accuracy needed on the magnetic field strength for equilibrium reconstruction routines [60]. Numerical noise in the finite difference scheme invalidates the values on the magnetic axis (see section 5).

the full MHD force residual better than the ground truth equilibrium computed by VMEC.

In addition, the fine-tuning procedure also improves the reconstruction of the magnetic field strength: the average error is below 10 mT for the entire plasma volume (figure 13).

The fine-tuning procedure is currently computationally expensive: on a single GPU, the fine-tuning procedure reaches the MHD force residual of the VMEC ground truth equilibrium after 4807 iterations (i.e. gradient update step) and 3610 s, whereas VMEC converges after 13 783 iterations and 127 s (on 20 central processing unit cores). Both the fine-tuning procedure and VMEC use a gradient descent method, however, even if the model requires fewer iterations to achieve the same MHD force residual, the model runtime is longer: the network and loss function are described in python, while VMEC is written in FORTRAN (thus resulting in compiled code).


The fine-tuning procedure simply shows how the single-pass model predictions can be improved in a self-supervised fashion (i.e. without the use of a ground truth equilibrium). The reduction of the fine-tuning runtime (e.g. just-in-time compilation of both the network and loss function) and the improvement of the fine-tuning convergence (e.g. higher learning rate) are left for future investigations.

#### 5.4. Ideal-MHD stability

The magnetic well  $W$  [61, 62] represents a fast proxy for ideal-MHD stability [59, 63]. The configuration is taken to be stable if [64]:

$$p'W < 0. \quad (50)$$





**Figure 10.** True (left) and predicted (right)  $B$  magnetic field strength structure at two radial locations:  $s = 0.06$  (top, close to the magnetic axis) and  $s = 0.75$  (bottom, close to the plasma boundary). The median predicted equilibrium in the test set is shown.

**Table 3.** Comparison between the not-regularized and regularized model in minimizing the MHD force residual proxy across all equilibria in the test set. NN denotes a model that has not been regularized with the MHD force residual (i.e. trained only in the data and gradient stages),  $\text{NN}_{\text{regularized}}$  denotes the model that has been regularized (i.e. trained also in the physics regularization stage), and VMEC denotes the value for the ground truth equilibria.

	NN	$\text{NN}_{\text{regularized}}$	VMEC
$\text{rmse}_{F_s}$ (mm)	$6.08 \times 10^{-1}$	$6.16 \times 10^{-1}$	—
$\text{mape}_t$	$9.24 \times 10^{-4}$	$9.31 \times 10^{-4}$	—
$\varepsilon_{f_*}$ ( $\text{NHm}^{-1}$ )	$7.27 \times 10^{-2}$	$1.97 \times 10^{-2}$	$2.43 \times 10^{-5}$


Therefore, for a standard pressure profile with  $p' < 0$ , a positive  $W$  is favorable for stability.

In the literature, multiple definitions can be found. When considering vacuum configurations, the magnetic well can be expressed as the second radial derivative of the plasma volume:

$$W_{\text{vacuum}}(s) = \frac{d^2 V_p}{d\psi^2} = \frac{d^2 V_p}{ds^2} \frac{d^2 s}{d\psi^2} = \frac{1}{\psi_{\text{edge}}^2} \frac{d^2 V_p}{ds^2}, \quad (51)$$

where  $V_p(s)$  is the plasma volume within the flux surface labelled  $s$ :

$$V_p(s) = \int_0^s \int_0^{2\pi} \int_0^{2\pi} |\sqrt{g}| ds' d\theta d\varphi. \quad (52)$$



**Figure 11.** True (solid pink) and predicted (dashed green) normalized MHD force proxy residual for the median regressed equilibrium. Numerical noise in the finite difference scheme invalidates the values at  $s \in [0, 0.01]$  (see section 5).


When a finite plasma pressure is introduced, the expression is modified into [61]:

$$W(s) = \frac{V_p}{\langle B^2 \rangle} \frac{d}{dV_p} (2\mu_0 p + \langle B^2 \rangle). \quad (53)$$



**Table 4.** Comparison between the fine-tuned, not-regularized and regularized (*single-pass*) model in minimizing the MHD force residual in case of a W7-X standard configuration at  $\langle\beta\rangle = 2\%$ , which is not part of the training data.  $\varepsilon_{f_*}$  and  $\varepsilon_F$  denote the proxy and full MHD force residuals,  $\eta_{f_*}$  and  $\eta_F$  denote the proxy and full normalized MHD force residuals. Please refer to section 5.3 for the definition of these quantities. NN denotes a model that has not been regularized with the MHD force residual (i.e. trained only in the data and gradient stages), NN<sub>regularized</sub> denotes the model that has been regularized with MHD force residual (i.e. the *single-pass* model), and NN<sub>fine-tuned</sub> denotes the fine-tuned model. VMEC denotes the value for the ground truth equilibrium. For each quantity, the best value is highlighted in bold.

	NN	NN <sub>regularized</sub>	NN <sub>fine-tuned</sub>	VMEC
rmse <sub>F<sub>s</sub></sub> (mm)	$3.11 \times 10^{-4}$	$3.21 \times 10^{-4}$	$3.63 \times 10^{-3}$	—
$\varepsilon_{f_*}$ (NHm <sup>-1</sup> )	$4.90 \times 10^{-2}$	$1.18 \times 10^{-2}$	$8.11 \times 10^{-4}$	$9.86 \times 10^{-6}$
$\varepsilon_F$ (NHm <sup>-1</sup> )	$8.83 \times 10^{-1}$	$8.28 \times 10^{-1}$	$1.24 \times 10^{-2}$	$5.43 \times 10^{-2}$
$\eta_{f_*}$	$7.17 \times 10^{-1}$	$1.73 \times 10^{-1}$	$1.19 \times 10^{-2}$	$1.44 \times 10^{-4}$
$\eta_F$	$3.51 \times 10^0$	$3.29 \times 10^0$	$4.92 \times 10^{-2}$	$2.16 \times 10^{-1}$



**Figure 12.** The comparison between the *single-pass* and fine-tuned model in minimizing the proxy and full MHD force residuals. (a) The normalized MHD force residual proxy for the *single-pass* (dashed green), fine-tuned (solid green), and ground truth equilibrium (solid pink). A W7-X standard configuration at  $\langle\beta\rangle = 2\%$  equilibrium, which was not part of the training data, is shown here. Numerical noise in the finite difference scheme invalidates the values at  $s \in [0, 0.01]$  (see section 5). (b) The normalized full MHD force residual for the *single-pass* (dashed green), fine-tuned (solid green), and ground truth equilibrium (solid pink). A W7-X standard configuration at  $\langle\beta\rangle = 2\%$  equilibrium, which was not part of the training data, is shown here. Numerical noise in the finite difference scheme invalidates the values at  $s \in [0, 0.01]$  (see section 5).

Moreover, the magnetic well term as defined in the Mercier stability criterion is [64, 65]:

$$D_W(s) = \frac{dp}{ds} \left( \frac{d^2 V_p}{ds^2} - \frac{dp}{ds} \int_0^{2\pi} \int_0^{2\pi} \frac{|\sqrt{g}| d\theta d\varphi}{B^2} \right) \times \int_0^{2\pi} \int_0^{2\pi} |\sqrt{g}| d\theta d\varphi \frac{B^2}{|\nabla s|^2}, \quad (54)$$

where the  $\frac{s^2}{\pi^2}$  prefactor common to all Mercier terms has been omitted, as employed in VMEC.

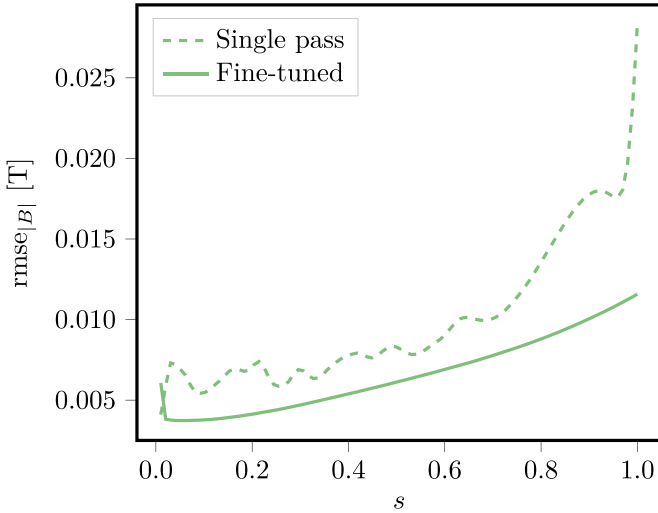
The model only qualitatively reconstructs the local magnetic well (figures 14(a)–(c)); it correctly provides the trend of the profile, however, it introduces artificial wiggles: because the magnetic well depends on the solution's second radial derivatives, the magnetic well is particularly sensitive to the solution's radial dependency. The relative error for all three expressions of the magnetic well is high (see table 2).

The MHD force residual regularization aids the model to faithfully reconstruct equilibrium properties. As an example,

figure 15 shows the reconstructed local magnetic well by the not-regularized, the regularized (i.e. *single-pass*), and the fine-tuned models. We again consider the W7-X standard configuration equilibrium at  $\langle\beta\rangle = 2\%$  previously examined (see section 5.3.1). The amplitude of the radial wiggle decreases from the not-regularized to the regularized models, and it almost vanishes in the fine-tuned model. Since the MHD force residual depends on the equilibrium solution's second radial derivatives, the MHD force residual regularization also implicitly regularizes the equilibrium solution's radial derivatives.

To assess the MHD stability of an equilibrium, a globally defined magnetic well is usually considered. For example, stellarator optimization frameworks often use the magnetic well depth as a fast proxy for plasma stability [8, 12]:

$$\bar{W}_{\text{vacuum}} = \frac{V_p'(s)|_{s=0} - V_p'(s)|_{s=1}}{V_p'(s)|_{s=0}}, \quad (55)$$



**Figure 13.** The magnetic field strength error as a function of the radial profile, averaged across a flux surface, in case of a W7-X standard configuration at  $\langle\beta\rangle = 2\%$  equilibrium, which was not part of the training data. The *single-pass* (dashed green), and fine-tuned (solid green) equilibria are shown. Numerical noise in the finite difference scheme invalidates the values on the magnetic axis (see section 5).

where  $\bar{W}_{\text{vacuum}}$  does not consider finite- $\langle\beta\rangle$  effects. A magnetic well depth that considers also a finite pressure can be defined as:

$$\bar{W} = -\frac{p^*(s)|_{s=0} - p^*(s)|_{s=1}}{p^*(s)|_{s=0}}, \quad (56)$$

where  $p^* = 2\mu_0 p + \langle B^2 \rangle$  is the sum of the fluid and the magnetic pressure.

Despite only a qualitative agreement of the local magnetic well, the model faithfully reconstructs the magnetic well depth (figures 16(a) and (b)). In case of the vacuum magnetic well depth, the relative error is remarkably low:  $\text{mape}_{\bar{W}_{\text{vacuum}}} = 3.43\%$ . The accuracy on the magnetic well depth is sufficient to effectively navigate the magnetic configuration space of W7-X in finding precise, negative-well configurations (see section 5.7.1).

### 5.5. Neoclassical transport

In a stellarator, one critical transport regime is the so called  $1/\nu$  regime, where the neoclassical transport increases with decreasing collision frequency. A measure of such transport is the effective ripple coefficient,  $\epsilon_{\text{eff}}$  [66]. In this work, the drift-kinetic code NEO [67, 68] evaluates the  $\epsilon_{\text{eff}}$ .

No acceptable agreement is present between the predicted and ground truth  $\epsilon_{\text{eff}}$  values (figure 17(a)). The effective ripple is greatly affected by the local structure of the magnetic field, which the model struggles to smoothly reconstruct (figure 10). It seems coherent that the artificial field ripples in the model equilibrium solution affect the accuracy on the effective ripple.

A poor, qualitative agreement between the ground truth and predicted  $\epsilon_{\text{eff}}$  is present only in the core region (figure 17(b)). In such region, namely for  $s \leq 0.33$  ( $\rho \leq 0.57$ ), the magnetic

field strength is reconstructed with an accuracy below 10 mT. However, the  $\epsilon_{\text{eff}}$  relative error is still high,  $\text{mape}_{\epsilon_{\text{eff}}} = 54.6\%$ .

The fine-tuned model can faithfully resolve  $\epsilon_{\text{eff}}$ . As an example, figure 18 shows the reconstructed  $\epsilon_{\text{eff}}$  by the not-regularized, the regularized, and the fine-tuned models in case of the W7-X standard configuration equilibrium at  $\langle\beta\rangle = 2\%$  previously examined (see section 5.3.1). Both the not-regularized and regularized model fail to smoothly reconstruct  $\epsilon_{\text{eff}}$ .

Simplified proxies of the neoclassical transport do exist. For example, in case of W7-X magnetic configurations with a low mirror term ( $b_{01} \simeq 0$ ), the effective ripple can be approximated as [40]:

$$\epsilon_{\text{eff}} \simeq \hat{\epsilon}_{\text{eff}} = -b_{11}\kappa^{4/3}, \quad (57)$$

where  $b_{mn} = B_{mn}/B_0$ ,  $B_{mn}$  are the Fourier coefficients of the magnetic field strength  $B$  in Boozer coordinates,  $B_0$  is a reference magnetic field value (e.g.  $B_{00}(s=1)$ ),  $\kappa = -b_{10}\frac{R}{r_{\text{eff}}}$  is the toroidal curvature term,  $R$  is the major radius, and  $r_{\text{eff}} = \rho a$  is the effective radius.

The model well reconstructs the epsilon effective proxy  $\hat{\epsilon}_{\text{eff}}$  (figure 19).  $\hat{\epsilon}_{\text{eff}}$  relies on few geometrical quantities and on the leading Fourier components of the magnetic field strength  $B$ , which the model faithfully reconstructs (see sections 5.1 and 5.2). In case of the epsilon effective proxy  $\hat{\epsilon}_{\text{eff}}$ , the model shows a relative error of  $\text{mape}_{\hat{\epsilon}_{\text{eff}}} = 13.3\%$ , lower than compared with  $\epsilon_{\text{eff}}$ .

### 5.6. Fast particle confinement

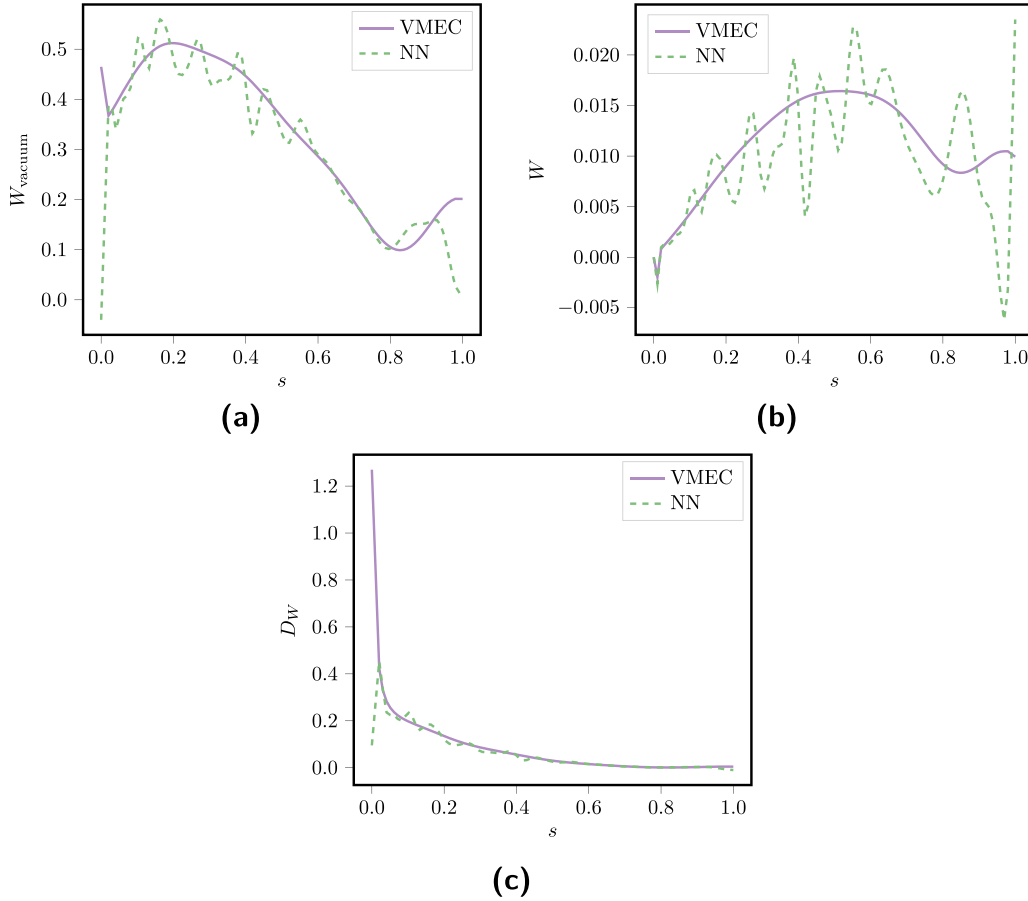
Another critical measure of transport is the confinement of fast  $\alpha$ -particles (energetic ions), which are the products of fusion reactions. Any fusion facility should confine fast particles so that they can heat the bulk plasma during the plasma burn, and to avoid damage to the first wall as a result of their losses.

Codes like the symplectic integration methods for particle loss estimation (SIMPLE) compute the loss fraction of test particles, for both trapped and passing particles, when injected with an initial velocity in the plasma volume. The collisionless guiding centers of particle motion are followed in time, till a stopping condition is met or the particle leaves the plasma volume.


However, such computations are expensive. Stellarator optimization frameworks instead use explicit symmetries (e.g. quasi-symmetry), equilibrium properties (e.g. omnigenity) or simpler proxies that correlate with the confinement of fast particles (e.g.  $\Gamma_c$ ) [63, 69–73].

For example, W7-X is a quasi-isodynamic (QI) stellarator: a quasi-omnigenous (QO) stellarator with magnetic field strength contours that close poloidally. Omnigenity means that the bounce-averaged radial drift of locally trapped particles vanishes [74]. For a confining field  $B$ , the conditions of omnigenity are [75]:

1. The contours of  $B_{\text{max}}$  are straight lines in Boozer coordinates;
2. The magnetic field strength maxima and minima are the same across all field lines;



**Figure 14.** True (solid pink) and predicted (dashed green) values of the vacuum magnetic well (a), the finite pressure magnetic well (b), and the magnetic well term in the Mercier criterion (c), in case of the median regressed equilibrium. Numerical noise in the finite difference scheme invalidates the values at  $s \in [0, 0.01]$  (see section 5).



**Figure 15.** Comparison on how the not-regularized, the regularized (i.e. single-pass), and the fine-tuned model reconstructs the local magnetic well  $W$ . Numerical noise in the finite difference scheme invalidates the values at  $s \in [0, 0.01]$  (see section 5).

3. Equal Boozer angular separation  $\delta$  along a field line between contours of constant magnetic field strength  $B$  (i.e. bounce points) across all field lines;

Does the model reconstruct these equilibrium properties?

To assess property (1), the standard deviation of the toroidal location of the contours of  $B_{\max}$ ,  $\sigma_{\varphi_{B_{\max}}}$ , is considered. In a QO configuration, the  $B_{\max}$  contours are straight lines in Boozer coordinates. Namely,  $\sigma_{\varphi_{B_{\max}}} = 0$ .


The model only qualitatively reconstructs the toroidal variance of the contours of  $B_{\max}$  (figure 20). The contours of  $B_{\max}$  is a local equilibrium property that is affected by how well the model reconstructs the local magnetic field strength.

To assess property (2), the standard deviation of the extrema of  $B$  on a flux surface across all field lines are assessed,  $\sigma_{B_{\max}}$  and  $\sigma_{B_{\min}}$ . In a QO configuration,  $\sigma_{B_{\max}}(s) = \sigma_{B_{\min}}(s) = 0$  across the whole plasma volume.

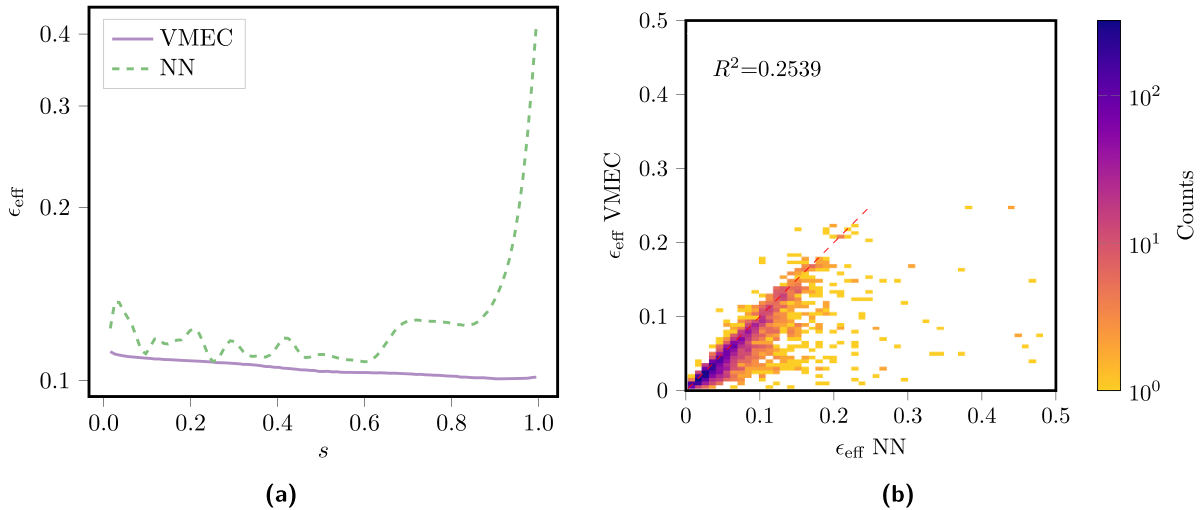
The model equilibria correctly reproduce the standard deviation of the extrema of  $B$  (figures 21(a) and (b)):  $\text{rmse}_{\sigma_{B_{\max}}} = 9.43 \text{ mT}$ , and  $\text{rmse}_{\sigma_{B_{\min}}} = 6.57 \text{ mT}$ .

The Boozer angular separation  $\delta(s, B, \alpha)$  is the angular distance along the field line  $\alpha$  between the  $B$  contours (see figure 28). In a QO configuration, the angular separation  $\delta(s, B, \alpha)$  should be independent of the field line label  $\alpha$ , namely,  $\frac{\partial \delta(s, B, \alpha)}{\partial \alpha} \Big|_{s, B} = 0$ . Numerically, such condition can be verified by observing  $\sigma_{\delta}(s, B) = 0$ .

The Boozer angular separation  $\delta$  relates to the bounce-averaged radial drift that trapped particles experience. Therefore, it makes sense to assess it only if a clear  $B$  field



**Figure 16.** Reconstruction of the vacuum (a) and finite pressure (b) magnetic well depth across all equilibria in the test set. The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction. The magnetic well depth is a measure of an average magnetic well value [8].




**Figure 17.** (a) True (solid pink) and predicted (dashed green) values of the effective ripple  $\epsilon_{\text{eff}}$  for the median regressed equilibrium. (b) Reconstruction of the effective ripple  $\epsilon_{\text{eff}}$  across all equilibria in the test set. The  $\epsilon_{\text{eff}}$  is computed on five equally spaced flux surfaces in the plasma region for  $s \leq 0.33$  ( $\rho \leq 0.57$ ). The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction.

well, in which particles can be trapped, exists. Not all W7-X magnetic configurations exhibit such a  $B$  field well: the median regressed equilibrium in the test set is not guaranteed to possess it. Therefore, only to assess property (3), the model is evaluated on a small (58 equilibria) out-of-sample set of high-mirror W7-X configurations. The high-mirror W7-X configuration has been constructed to feature a magnetic-mirror like  $B$  field structure, and it is the closest configuration to QI out of the nine W7-X reference configurations. This test set has been constructed as the main data set (see section 4), however, the coil current ratios have been fixed to generate only high-mirror W7-X configurations. Still, finite- $\beta$  and toroidal current effects are present in the data set. For completeness, section ‘Boozer angular separation  $\delta$  on test set’ presents the  $\delta$  reconstruction in case of the median regressed equilibrium in the larger test set.


Even if the predicted magnetic field structure does not perfectly overlap with the ground truth field (figure 10), the model preserves the Boozer angular separation (figure 22). However, the reconstruction accuracy is better in the core region (figure 23).

### 5.7. Optimization in the configuration space of W7-X

The exploration and optimization of the W7-X magnetic configuration space serve to further investigate how faithfully the NN model reproduces equilibrium properties, and to validate its use in downstream tasks. For example, [70, 76, 77] proposed the *a posteriori* optimization of coil currents to further improve the equilibrium properties, and to recover inevitable errors that occur during the coils manufacturing and placement. Moreover, especially in case of an island divertor that



**Figure 18.** Comparison on how the not-regularized, the regularized (i.e. single-pass), and the fine-tuned model reconstructs  $\epsilon_{\text{eff}}$ . Numerical noise in the finite difference scheme invalidates the values at  $s \in [0, 0.01]$  (see section 5).




**Figure 19.** Reconstruction of the effective ripple proxy  $\hat{\epsilon}_{\text{eff}}$  across all equilibria in the test set. The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction.

relies on a robust edge magnetic topology, it is desirable to include finite- $\beta$  effects in the optimization.

A fast, free-boundary ideal-MHD equilibrium model benefits such optimization. In this work, W7-X configurations are targeted due to the limitations of the training data. However, the proposed technique can be applied to any present or future device once a training data set is generated. The optimizations carried out in this work represent toy examples of real-world applications. Indeed, the proposal of new, optimized W7-X configurations is beyond the scope of this paper.

The equilibrium solution provided by the NN is differentiable. The analytic gradient of equilibrium properties (e.g. the magnetic well) with respect to the magnetic configuration



**Figure 20.** Reconstruction of the standard deviation of the toroidal location of the contours of  $B_{\text{max}}$ ,  $\sigma_{\varphi_{B_{\text{max}}}}$ . The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction.

and plasma profiles is available at no additional cost via automatic differentiation (AD). In a gradient-based optimization, a single equilibrium evaluation is sufficient in each optimization step (regardless of the number of free parameters). In contrast, when finite-differences approximate the objective function gradient, the number of equilibrium evaluations each step scales linearly with the number of free parameters.

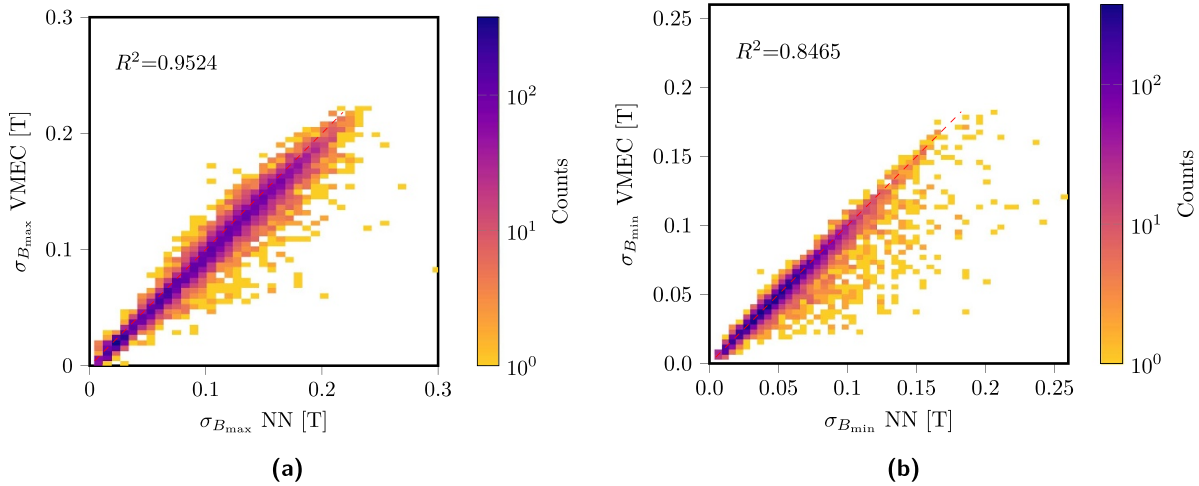
Both magnetic configuration and finite- $\beta$  effects are included in the optimization. The search space is a subset of the NN model input parameters: the W7-X non-planar  $i_{[2, \dots, 5]}$  and planar  $i_{[A, B]}$  coil current ratios, the toroidal magnetic flux at the edge  $\Phi_{\text{edge}}$ , and the pressure on axis  $p_0$ . The pressure profile shape is fixed, and it is assumed to be of the form  $p(s) = p_0(1-s)^2$ . Moreover, we assume the toroidal plasma current to vanish (i.e.  $I_{\text{tor}}(s) = 0$ ).

To avoid extrapolation, the search space is bounded by the distribution of the input parameters as seen during training. Denoting  $\vec{p} = [p_1, \dots, p_d] \in \mathbb{R}^d$  the optimization vector, each input parameter  $p_i$  is bounded by  $[\mu_i - c\sigma_i, \mu_i + c\sigma_i]$ , where  $\mu_i$  and  $\sigma_i$  are the mean and the standard variation of the distribution of  $p_i$ , respectively.  $c = \sqrt{3}$  for the parameters that have been uniformly sampled in the training data, and  $c = 2$  for all the others. The configuration with  $p_i = \mu_i \forall i$  serves as initial guess (i.e. the *average* configuration seen during training).

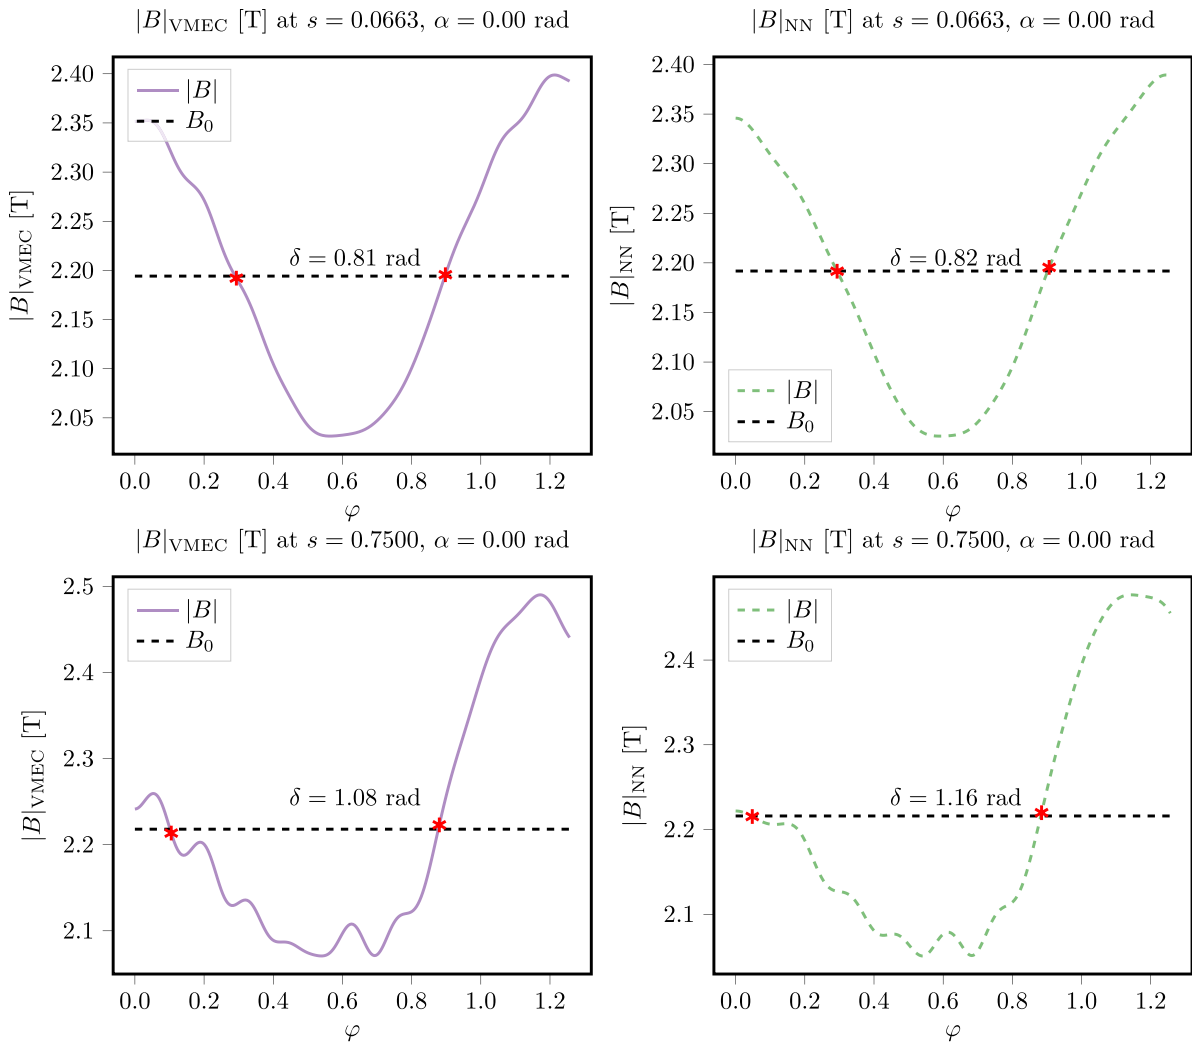
The objective function is the weighted sum of the squared residual between the proposed and target equilibrium properties:

$$f(\vec{p}) = \sum_i w_i (f_i(\vec{p}) - f_i^*)^2, \quad (58)$$

where  $f_i^*$  is a target equilibrium property,  $f_i(\vec{p})$  is the proposed equilibrium property, and  $w_i$  is the weight of each residual term. Such form of the objective function is usually employed in stellarator optimization frameworks [63, 78, 79].




**Figure 21.** Reconstruction of the standard deviation of the (a) maxima  $\sigma_{B_{\max}}$  and (b) minima  $\sigma_{B_{\min}}$  of the magnetic field strength across all field lines on a flux surface. The standard deviation evaluated on ten equally spaced flux surfaces is considered in these figures. The color-bar indicates the counts in each histogram bin. To guide the eye, a red dashed line indicates the case of perfect reconstruction.



**Figure 22.** True (left) and predicted (right) values for the magnetic field strength  $B$  in Boozer coordinates along a field line in case of the median predicted equilibrium in the high-mirror out-of-sample test set. The  $\alpha = 0$  field line at the  $s = 0.06$  (top) and at the  $s = 0.75$  flux surface (bottom) are depicted. The dashed black horizontal line indicates the average magnetic field strength on the flux surface  $B_0$ . The red crosses highlight the intersections of the field line with the  $B_0$  contours, and the Boozer angular separation  $\delta(s, B_0, \alpha)$  between them is shown just beneath.





**Figure 23.** True (solid pink) and predicted (dashed green) values of the standard deviation of Boozer separation angle

$\sigma_\delta(s, B_0) = \sqrt{\langle \delta(s, B_0, \alpha)^2 \rangle_\alpha}$  for the median regressed equilibrium in the high-mirror out-of-sample test set. Without loss of generality, the contour of the average magnetic field strength  $B_0$  on each flux surface is considered.


A combination of the global tree-structured Parzen estimator (TPE) algorithm from *hyperopt* [80] and the local limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) [81] algorithm from *scipy* [82] drives the optimization. Whenever possible, AD computes the analytical Jacobian of the objective function. On the other hand, if the analytical Jacobian is not available, (e.g. an objective function term requires an external software package that breaks the model computational graph), finite differences approximate the Jacobian.

**5.71. Negative well.** Almost all W7-X configurations in the training set feature a positive magnetic well (figure 24). W7-X was optimized with respect to good MHD stability [7], therefore, in general, a positive magnetic well is to be expected.

Can the model generalize to negative well configurations? To search for a negative well configuration, a small, but negative magnetic well depth is targeted:  $\overline{W}_{\text{vacuum}}^* = -0.01$  and  $w_{\overline{W}_{\text{vacuum}}} = 10^4$ . Finite- $\beta$  effects deepen the magnetic well, therefore, the search is performed in the limit of vacuum configurations. After 500 iterations of TPE, L-BFGS refines the best configuration.

Although the NN model has rarely seen a negative well configuration during training, it can reliably predict equilibria with a negative well (table 5). The optimization converged to a configuration for which the model predicts a magnetic well depth of  $-0.01$ . The configuration has been evaluated also with VMEC, which reports a magnetic well depth of  $-0.009$ . Indeed, the vacuum magnetic well ( $W(\psi) = \frac{d^2 V_p(\psi)}{d\psi^2}$ ) is negative (figure 25).

**5.72. Improved fast particle confinement.** A classical stellarator does not confine fast particles as good as a tokamak,



**Figure 24.** The distribution of the vacuum magnetic well depth  $\overline{W}_{\text{vacuum}}$  for all equilibria in the training set.


**Table 5.** Relative magnitude of the non-planar and planar W7-X coil currents for the negative well configuration.  $\overline{W}_{\text{vacuum}}^* = -0.01$  is the small, negative magnetic well depth targeted in the optimization. In case of the obtained configuration, the NN model predicts  $\overline{W}_{\text{vacuum}} = -0.010$ , and VMEC computes  $\overline{W}_{\text{vacuum}} = -0.009$ .

Quantity	Value
$i_1$	1.000
$i_2$	1.027
$i_3$	0.689
$i_4$	0.614
$i_5$	0.614
$i_A$	0.205
$i_B$	0.263
$\overline{W}_{\text{vacuum}}^*$	$-0.010$
$\overline{W}_{\text{vacuum}} \text{ NN}$	$-0.010$
$\overline{W}_{\text{vacuum}} \text{ VMEC}$	$-0.009$

therefore, the magnetic geometry has to be optimized to do so. W7-X standard and high-mirror configurations have been optimized to improve fast particle confinement at high- $\beta$ . Moreover, [70, 76] have also found improved configurations.

Can we use the model to search for a W7-X fast particle optimized configuration? To address this question, we employ a similar objective function to the function used for the *min-max* configuration in [70], and we investigate whether the optimized configuration shows the same performance in terms of fast particle confinement.

In a perfect QI configuration, the variance of the extrema of the magnetic field strength across all field lines vanishes. As in the case of the *min-max* configuration, the variance of  $B$  at the bean-shaped cross-section (i.e.  $\varphi = 0$  or the  $B_{\text{max}}$ -contour) and at the triangular cross-section (i.e.  $\varphi = \pi/N_{\text{fp}}$  or the  $B_{\text{min}}$ -contour) is minimized. The flux surface at  $s = 0.015$  is targeted (roughly 6 cm from the magnetic axis).



**Figure 25.** The magnetic well of the obtained W7-X negative well configuration (green) as evaluated by VMEC. For comparison, the magnetic well of the W7-X standard reference configuration (blue) is also shown (the W7-X standard configuration has a positive magnetic well). Numerical noise in the finite difference scheme invalidates the values at  $s \in [0, 0.01]$  (see section 5).

The optimization is performed at a fixed  $\Phi_{\text{edge}}$  value. No toroidal current is assumed in the plasma.

The objective function is:


$$f = w_B(\text{Var}_{B_{\text{max}}} + \text{Var}_{B_{\text{min}}}) + w_\beta(\langle\beta\rangle - \langle\beta\rangle^*)^2 + w_{mr}(mr - mr^*)^2, \quad (59)$$

where  $\text{Var}_{B_{\text{max}}}$  and  $\text{Var}_{B_{\text{min}}}$  are the variances of the magnetic field strength at the bean-shaped and triangular cross-section, respectively, and

$$mr = \left\langle \frac{B(\theta, \varphi = 0) - B(\theta, \varphi = \pi/N_{\text{fp}})}{B(\theta, \varphi = 0) + B(\theta, \varphi = \pi/N_{\text{fp}})} \right\rangle_\theta, \quad (60)$$

is a mirror factor to ensure that the extrema of the magnetic field strength are located at the correct angular position (maxima at  $\varphi = 0$  and minima at  $\varphi = \pi/N_{\text{fp}}$ ).  $\langle\beta\rangle^* = 2\%$  and  $mr^* = 10\%$  (like in the W7-X *high-mirror* configuration). Finally,  $w_B = 10^2$ ,  $w_{\langle\beta\rangle} = 10^3$ , and  $w_{mr} = 10^2$ . The optimization is carried out for 500 iterations of the TPE algorithm.

SIMPLE [83, 84] assess the fraction of lost particles from the plasma volume.  $10^4$  particles at 60 keV are launched from  $s = 0.06$ , and their trajectories are followed for 0.1 s, or till they are lost from the computational domain. The optimized configuration is compared against the W7-X *standard*, *high-mirror* and *min-max* configurations. All configurations have been scaled to have the same magnetic field strength on axis  $B(\varphi = 0) \simeq 2.52$  T, the same plasma beta  $\langle\beta\rangle \simeq 2\%$ , and the same plasma volume  $V_p \simeq 30\text{m}^3$ .



**Figure 26.** The fraction of lost particles for the W7-X standard, high-mirror, *min-max*, and optimized configuration at  $\langle\beta\rangle = 2\%$ .

**Table 6.** Relative magnitude of the non-planar and planar W7-X coil currents for the  $\langle\beta\rangle = 2\%$  fast particle confinement optimized configuration. The total fraction of lost 60 keV fast ions initialized at  $s = 0.06$  is reported.

Quantity	Value
$i_1$	1.000
$i_2$	0.737
$i_3$	0.889
$i_4$	0.776
$i_5$	0.899
$i_A$	0.011
$i_B$	-0.149
Lost fraction	8.86%

The optimized configuration confines fast particles better than the *standard* and *high-mirror* configurations, and as well as the *min-max* configuration (figure 26): the total fraction of lost particles after 0.1 s is only 8.86% (table 6).

## 6. Discussion and conclusion

We propose a NN model that quickly approximates the ideal-MHD solution operator in a stellarator geometry. The model correctly predicts the equilibrium solution, and it faithfully reconstructs global equilibrium properties and proxy functions that are commonly used in stellarator optimization frameworks. Compared with previous data-driven approaches, the model fulfils equilibrium symmetries by construction. In addition, the MHD force residual regularizes the model to better satisfy the ideal-MHD equations.

The error between the model and the ground truth equilibrium solution is less than 1 mm in case of the flux surface

locations. However, mainly due to the high condition number of the equilibrium problem, the model equilibria only poorly minimize the MHD force residual. Yet, the model predictions can be improved at inference time to yield better than ground truth equilibria (i.e. equilibria that minimize the MHD force residual better than the ground truths).

Equilibrium properties that are highly sensitive to the magnetic field structure cannot be precisely reconstructed. The magnetic field depends on first-order derivatives of the equilibrium solution, therefore, even small errors in the equilibrium geometry might lead to large discrepancies in the computed magnetic field. Nevertheless, the magnetic field strength is smoothly reconstructed in the core region, but the model introduces artificial field ripples at the edge. As a consequence, the model fails to meet the required magnetic field accuracy to faithfully compute some quantities of interest (e.g. the effective ripple, which is a measure of neoclassical transport). Still, the accuracy of the model is sufficient to explore and optimize W7-X magnetic configurations, in terms of ideal-MHD stability and fast particle confinement. This result suggests that NN-based surrogate models can be used to find optimized configurations for current and future stellarator devices.

The model has still many limitations. The NN is trained in a supervised learning fashion, therefore, the training data limit the applicability of the model (the model can only approximate W7-X equilibria). Moreover, noise-free plasma profiles are assumed. Finally, the model accuracy on the magnetic field limits its use in tasks that are highly sensitive to it.

This work can be improved along multiple dimensions. The robustness of the model with respect to experimental noise can be investigated and enhanced with established robustness and regularization techniques: injecting noise to the input plasma profiles and training the convolutional layers with dropout. The artificial field ripples introduced by the model can be ameliorated by adding a regularization term on the second-order radial derivative (e.g.  $\frac{\partial^2 R}{\partial s^2}$ ), or, by replacing the learned trunk network functions with an orthonormal set of basis functions (e.g. Legendre polynomials). In addition, the model can be trained on an extended space of stellarator equilibria, relaxing the constraint of W7-X configurations. Finally, equilibrium reconstructions in a stellarator geometry must be demonstrated.

Contrary to a tokamak, the computation of MHD equilibria is a limiting factor in stellarator research and design. Fast, differentiable and accurate three-dimensional MHD equilibrium solutions would reduce the gap between tokamaks and stellarators. In particular, they would enable: real-time equilibrium reconstructions [85, 86], first-principles flight simulators [9], data-driven plasma control [87], and extensive understanding and exploration of the stellarator optimization space.

## Acknowledgments

We wish to acknowledge the helpful discussions on the limitations of this work with Ralf Schneider. Furthermore, we are indebted to the communities behind the multiple open-source software packages on which this work depends: hydra [88],

hyperopt [80], matplotlib [89], numpy [90], pytorch [91], pytorch lightning [92], scipy [82].

The data sets were generated on the Max Planck computing and data facility (MPCDF) cluster ‘COBRA’, Germany. Financial support by the European Social Fund (ID: ESF/14-BM-A55-0007/19) and the Ministry of Education, Science and Culture of Mecklenburg-Vorpommern, Germany via the project ‘NEISS’ is gratefully acknowledged. This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No. 101052200—EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

## Author contributions

The contributions to this paper are described using the CRediT taxonomy [93]:

**Andrea Merlo:** Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Software, Visualization, Writing—original draft, Writing—review & editing.

**Daniel Böckenhoff:** Methodology, Software, Supervision, Validation, Writing—review & editing.

**Jonathan Schilling:** Formal Analysis, Methodology.

**Samuel Aaron Lazerson:** Supervision, Writing—review & editing.

**Thomas Sunn Pedersen:** Conceptualization, Funding acquisition, Supervision.


## Appendix

### Equilibria Fourier scaling

A trade-off must be struck between model capacity and computational complexity. Given the poloidal and toroidal expansions of the equilibrium solution in Fourier series, the number of Fourier modes utilized to represent the solution limits its accuracy: a large number of Fourier modes enhances the model accuracy. On the other hand, it increases the computational complexity of the model, as well as the training and inference times. How many Fourier modes are adequate?

To examine such trade-off, the approximation error on magnetic field strength is investigated. One of the objectives of this study is to investigate the precision with which the model reconstructs equilibrium properties. The magnetic field strength is one of the critical equilibrium properties. Therefore, the Fourier resolution of the equilibrium solution should not be the largest source of error in representing the field strength.

To represent the magnetic field strength with an error of less than 1%, up to  $m = 8$  poloidal and  $|n| = 11$  toroidal models are required. Figure 27 shows the magnetic field approximation error when the Fourier resolution is varied, using VMEC equilibria with up to  $m = 11$  and  $|n| = 12$  Fourier modes as ‘high-fidelity’ references.



**Figure 27.** The approximation error on the magnetic field strength when representing the equilibrium solution  $\{R, \lambda, Z\}$  with a reduced number of Fourier modes. The equilibrium with up to  $m = 11$  poloidal and  $|n| = 12$  toroidal mode is taken as a reference. The solid line represents the average error across 128 randomly sampled equilibria in the data set, and the shaded region represents the 95% and 5% quantiles. W7-X has a mean field  $\simeq 2.5$  T, therefore, a field error of 25 mT represents a 1% error.

**Table 7.** HPs of the trunk network. See section 3.2 for details on the model macro-architecture.

HP	Value
Depth	4
Width	2048
Trunk basis	8
Activation function	SiLU

### Model architecture and HPs

The trunk network maps the radial dimension  $s$  into a set of non-linear basis functions, in which the  $X_{mn}^j$  Fourier coefficients and  $\iota$  profile are radially expanded (see section 3.2). Each Fourier coefficient and  $\iota$  has its own set of basis functions. The input dimension is  $N_{\text{ti}} = 1$ , and the output dimension is  $N_{\text{to}} = (N_{X_{mn}^j} + N_{\iota})N_{\text{tb}} = 4696$ , where  $N_{X_{mn}^j} = 3[(m_{\text{pol}} + 1)(2n_{\text{tor}} + 1) - n_{\text{tor}}] - 2 = 586$  is the number of Fourier modes used to represent  $(R, \lambda, Z)$ ,  $m_{\text{pol}} = 8$  is the highest Fourier poloidal mode,  $n_{\text{tor}} = 11$  is the highest Fourier toroidal mode,  $N_{\iota} = 1$  is the single output needed to represent the iota profile, and  $N_{\text{tb}} = 8$  is the number of trunk basis functions per output.

The trunk network is a MLP network with sigmoid linear unit (SiLU) as non-linear activation functions (table 7). Each layer is initialized accordingly to [94].

The branch network maps the MHD parameters and input plasma profiles into the weighting coefficients of the trunk basis functions (see section 3.2). A profile head extracts high-level features from the input profiles into a latent representation. A set of blocks made of one 1D convolutional and pooling layers composes the profile head. The branch body then processes the concatenation of the profile latent representation

**Table 8.** HPs of the branch network. See section 3.2 for details on the model macro-architecture.

HP	Value
Profile head	
Depth	2
Convolution kernel size	5
Convolution stride	2
Convolution padding	0
Filters	16
Pooling kernel size	3
Pooling stride	2
Pooling padding	0
Body	
Depth	5
Width	2048
Activation function	SiLU

**Table 9.** HPs for the data training stage.

HP	value
Batch size	128
Initial learning rate	$8.0 \times 10^{-4}$
$\gamma$	$9.9 \times 10^{-1}$
$L^2$ regularization	$6.0 \times 10^{-2}$
$\alpha_R$	$2.5 \times 10^{-1}$
$\alpha_Z$	$2.5 \times 10^{-1}$
$\alpha_\lambda$	$2.5 \times 10^{-1}$
$\alpha_\iota$	$2.5 \times 10^{-1}$
Patience epochs	150
Gradient clip	$1.0 \times 10^{-2}$

with the input scalar parameters. The output dimension is  $N_{\text{bo}} = (N_{X_{mn}^j} + N_{\iota})(N_{\text{tb}} + 2) = 5870$ , where the '+2' is for the constant and identity functions.

The branch body is a MLP network with SiLU as non-linear activation functions (table 7). Each layer is initialized accordingly to [94]. The last layer is initialized at the median value of the model outputs evaluated on 20% of the training data set.

The overall training process is divided in three different stages (see section 4.3): data, gradient, and physics regularization stages. During the data stage, the learning rate is decreased accordingly to an exponential decay schedule with  $\gamma$  as decay constant. During the gradient and physics regularization stages, the learning is kept constant. Only after the first step of the curriculum learning procedure in the physics regularization stage, the learning rate is decreased by a factor of 10. In addition, in every training stage, the norm of the model gradients is clipped. Tables 9–11 list the HPs in each training stage. Table 12 lists the HPs for the fine-tuning procedure.

### Derivation of MHD force residual

This section shows how the MHD force residual can be derived once the mapping  $\vec{x} = (R, \lambda, Z)$ , as well as the  $p(s)$  and  $\iota(s)$

**Table 10.** HPs for the gradient training stage.

HP	Value
Batch size	32
Initial learning rate	$1.0 \times 10^{-5}$
$\gamma$	1
$L^2$ regularization	$6.0 \times 10^{-2}$
$\alpha_R$	$2.4 \times 10^{-1}$
$\alpha_Z$	$2.4 \times 10^{-1}$
$\alpha_\lambda$	$2.4 \times 10^{-1}$
$\alpha_t$	$2.4 \times 10^{-1}$
$\alpha_{R_s}$	$1.0 \times 10^{-2}$
$\alpha_{Z_s}$	$1.0 \times 10^{-2}$
$\alpha_{\lambda_s}$	$1.0 \times 10^{-2}$
$\alpha_{t'}$	$1.0 \times 10^{-2}$
Patience epochs	150
Gradient clip	$1.0 \times 10^{-2}$

**Table 11.** HPs for the physics regularization training stage.

HP	Value
Batch size	32
Initial learning rate	$1.0 \times 10^{-6}$
$\gamma$	1
$L^2$ regularization	$6.0 \times 10^{-2}$
$\alpha_R$	$2.4 \times 10^{-1}$
$\alpha_Z$	$2.4 \times 10^{-1}$
$\alpha_\lambda$	$2.4 \times 10^{-1}$
$\alpha_t$	$2.399 \times 10^{-1}$
$\alpha_{R_s}$	$1.0 \times 10^{-2}$
$\alpha_{Z_s}$	$1.0 \times 10^{-2}$
$\alpha_{\lambda_s}$	$1.0 \times 10^{-2}$
$\alpha_{t'}$	$1.0 \times 10^{-2}$
$\alpha_{\text{MHD}}$	$1.0 \times 10^{-4}$
Patience epochs	50
Gradient clip	$1.0 \times 10^{-2}$

**Table 12.** HPs for the fine-tuning procedure.

HP	Value
Initial learning rate	$1.0 \times 10^{-4}$
$\gamma$	$9.996 \times 10^{-1}$
$L^2$ regularization	$1.0 \times 10^{-1}$
$\alpha_b$	$2.495 \times 10^{-1}$
$\alpha_I$	$5.0 \times 10^{-1}$
$\alpha_{\text{MHD}}$	$1.0 \times 10^{-3}$
Gradient clip	$1.0 \times 10^{-2}$

profiles, are known. As an example,  $F_\beta$  is derived;  $F_s$  can be derived in a similar manner.

Let us recall that:

$$F_\beta = J^s = \frac{1}{\mu_0 \sqrt{g}} \left( \frac{\partial B_\varphi}{\partial \theta} - \frac{\partial B_\theta}{\partial \varphi} \right), \quad (61)$$

where  $\sqrt{g}$  is the Jacobian of the 3D coordinate transformation  $f: (s, \theta, \varphi) \rightarrow (R, \phi, Z)$ , and it is expressed as

$\sqrt{g} = (\vec{e}_s \cdot \vec{e}_\theta \times \vec{e}_\varphi)$ . Recalling the definition of the covariant basis vectors (see section 2.1), the choice of  $\varphi = \phi$  effectively yields a 2D Jacobian:

$$\sqrt{g} = R \left( \frac{\partial R}{\partial \theta} \frac{\partial Z}{\partial s} - \frac{\partial R}{\partial s} \frac{\partial Z}{\partial \theta} \right). \quad (62)$$

The covariant magnetic field components (see equations (8) and (9)) are connected to the contravariant magnetic field components (see equations (5) and (6)) through the covariant metric tensor elements:

$$g_{ij} = \vec{e}_i \cdot \vec{e}_j = \frac{\partial R}{\partial \alpha_i} \frac{\partial R}{\partial \alpha_j} + R^2 \frac{\partial \phi}{\partial \alpha_i} \frac{\partial \phi}{\partial \alpha_j} + \frac{\partial Z}{\partial \alpha_i} \frac{\partial Z}{\partial \alpha_j}. \quad (63)$$

Finally, the expression for  $F_\beta$  reads:

$$F_\beta = \frac{\Phi'}{\mu_0 \sqrt{g}} \left\{ \frac{\partial}{\partial \theta} \left[ \frac{1}{\sqrt{g}} \left[ \left( t - \frac{\partial \lambda}{\partial \varphi} \right) \left( \frac{\partial R}{\partial \theta} \frac{\partial R}{\partial \varphi} + \frac{\partial Z}{\partial \theta} \frac{\partial Z}{\partial \varphi} \right) + \left( 1 + \frac{\partial \lambda}{\partial \theta} \right) \left( \frac{\partial R}{\partial \varphi} \frac{\partial R}{\partial \varphi} + R^2 + \frac{\partial Z}{\partial \varphi} \frac{\partial Z}{\partial \varphi} \right) \right] \right. \right. \quad (64)$$

$$\left. - \frac{\partial}{\partial \varphi} \left[ \frac{1}{\sqrt{g}} \left[ \left( t - \frac{\partial \lambda}{\partial \varphi} \right) \left( \frac{\partial R}{\partial \theta} \frac{\partial R}{\partial \theta} + \frac{\partial Z}{\partial \theta} \frac{\partial Z}{\partial \theta} \right) + \left( 1 + \frac{\partial \lambda}{\partial \theta} \right) \left( \frac{\partial R}{\partial \varphi} \frac{\partial R}{\partial \theta} + \frac{\partial Z}{\partial \varphi} \frac{\partial Z}{\partial \theta} \right) \right] \right] \right\}. \quad (65)$$

$F_\beta$  depends only on  $\vec{x}$ , its first and second derivatives, and  $t$ .

#### Physics regularization

The proxy of the MHD force residual  $f_*$  is derived as follows.

$$\vec{F} = -\vec{J} \times \vec{B} + \vec{\nabla} p, \quad (66)$$

$$\vec{\nabla} \times \vec{B} = \mu_0 \vec{J}, \quad (67)$$

$$\vec{\nabla} \cdot \vec{B} = 0. \quad (68)$$

The magnetic field can be written in contravariant form as:

$$\vec{B} = B^\theta \vec{e}_\theta + B^\varphi \vec{e}_\varphi, \quad (69)$$

$$B^\theta = \frac{1}{\sqrt{g}} \left( \chi' - \Phi' \frac{\partial \lambda}{\partial \varphi} \right), \quad (70)$$

$$B^\varphi = \frac{\Phi'}{\sqrt{g}} \left( 1 + \frac{\partial \lambda}{\partial \theta} \right). \quad (71)$$

By substituting the expression of  $\vec{B}$  into  $\vec{F}$ , the covariant form of the MHD force is:

$$\vec{F} = F_s \vec{\nabla} s + F_\beta \vec{\beta}, \quad (72)$$

$$F_s = \sqrt{g} (J^\varphi B^\theta - J^\theta B^\varphi) + p', \quad (73)$$



$$F_\beta = J^s, \quad (74)$$

where  $\vec{\beta} = \sqrt{g}(B^\varphi \vec{\nabla}\theta - B^\theta \vec{\nabla}\varphi)$ .

The contravariant components of the current density,  $J^i = \vec{J} \cdot \vec{\nabla}\alpha_i = \frac{1}{\mu_0}(\vec{\nabla} \times \vec{B}) \cdot \vec{\nabla}\alpha_i$ , are:

$$J^s = \frac{1}{\mu_0\sqrt{g}} \left( \frac{\partial B_\varphi}{\partial\theta} - \frac{\partial B_\theta}{\partial\varphi} \right) \quad (75)$$

$$J^\theta = \frac{1}{\mu_0\sqrt{g}} \left( \frac{\partial B_s}{\partial\varphi} - \frac{\partial B_\varphi}{\partial s} \right) \quad (76)$$

$$J^\varphi = \frac{1}{\mu_0\sqrt{g}} \left( \frac{\partial B_\theta}{\partial s} - \frac{\partial B_s}{\partial\theta} \right). \quad (77)$$

By substitute the contravariant current density components in the radial force component, it follows that:

$$F_s = \frac{1}{\mu_0} \left( B^\theta \frac{\partial B_\theta}{\partial s} - B^\theta \frac{\partial B_s}{\partial\theta} - B^\varphi \frac{\partial B_s}{\partial\varphi} + B^\varphi \frac{\partial B_\varphi}{\partial s} \right) + p'. \quad (78)$$

Let us define the proxy of the MHD residual as the flux surface averaged force residual norm when  $F_\beta = 0$ :

$$f_* = \frac{1}{(2\pi)^2} \langle \mu_0\sqrt{g} \|F\|_2 |_{F_\beta=0} \rangle = \frac{1}{(2\pi)^2} \langle \mu_0\sqrt{g} F_s \rangle \quad (79)$$

where  $\langle A \rangle = \int_0^{2\pi} \int_0^{2\pi} A d\theta d\varphi$  denotes a flux surface average operator (note that the  $\sqrt{g}$  factor is not included in the average operator here).

Substituting the contravariant magnetic field components (see equations (5) and (6)) in equation (79), the following terms result:

$$\begin{aligned} \left\langle \sqrt{g} B^\theta \frac{\partial B_\theta}{\partial s} \right\rangle &= \left\langle \left( \chi' - \Phi' \frac{\partial\lambda}{\partial\varphi} \right) \frac{\partial B_\theta}{\partial s} \right\rangle \\ &= \chi' \left\langle \frac{\partial B_\theta}{\partial s} \right\rangle + \Phi' \left\langle \lambda \frac{\partial B_\theta}{\partial s \partial \varphi} \right\rangle, \end{aligned} \quad (80)$$

$$\begin{aligned} \left\langle \sqrt{g} B^\theta \frac{\partial B_s}{\partial\theta} \right\rangle &= \left\langle \left( \chi' - \Phi' \frac{\partial\lambda}{\partial\varphi} \right) \frac{\partial B_s}{\partial\theta} \right\rangle \\ &= \chi' \left\langle \frac{\partial B_s}{\partial\theta} \right\rangle + \Phi' \left\langle \lambda \frac{\partial B_s}{\partial\theta \partial \varphi} \right\rangle = \Phi' \left\langle \lambda \frac{\partial B_s}{\partial\theta \partial \varphi} \right\rangle, \end{aligned} \quad (81)$$

$$\begin{aligned} \left\langle \sqrt{g} B^\varphi \frac{\partial B_\varphi}{\partial s} \right\rangle &= \left\langle \Phi' \left( 1 + \frac{\partial\lambda}{\partial\theta} \right) \frac{\partial B_\varphi}{\partial s} \right\rangle \\ &= \Phi' \frac{\partial \langle B_\varphi \rangle}{\partial s} - \Phi' \left\langle \lambda \frac{\partial B_\varphi}{\partial s \partial \theta} \right\rangle, \end{aligned} \quad (82)$$

$$\begin{aligned} \left\langle \sqrt{g} B^\varphi \frac{\partial B_s}{\partial\varphi} \right\rangle &= \left\langle \Phi' \left( 1 + \frac{\partial\lambda}{\partial\theta} \right) \frac{\partial B_s}{\partial\varphi} \right\rangle \\ &= \Phi' \left\langle \frac{\partial B_s}{\partial\varphi} \right\rangle - \Phi' \left\langle \lambda \frac{\partial B_s}{\partial\varphi \partial \theta} \right\rangle \\ &= -\Phi' \left\langle \lambda \frac{\partial B_s}{\partial\varphi \partial \theta} \right\rangle. \end{aligned} \quad (83)$$

The  $\frac{\partial\lambda}{\partial\varphi}$  and  $\frac{\partial\lambda}{\partial\theta}$  terms have been integrated by parts exploiting the fact that  $\lambda$  and the covariant component of the B field are periodic functions in the poloidal and toroidal directions. The second and fourth terms cancel each other.

In addition, when  $F_\beta = 0$  and  $\sqrt{g} \neq 0$ , it follows that:

$$F_\beta = J^s = \frac{1}{\mu_0\sqrt{g}} \left( \frac{\partial B_\varphi}{\partial\theta} - \frac{\partial B_\theta}{\partial\varphi} \right) = 0 \rightarrow \frac{\partial B_\varphi}{\partial\theta} = \frac{\partial B_\theta}{\partial\varphi}. \quad (84)$$

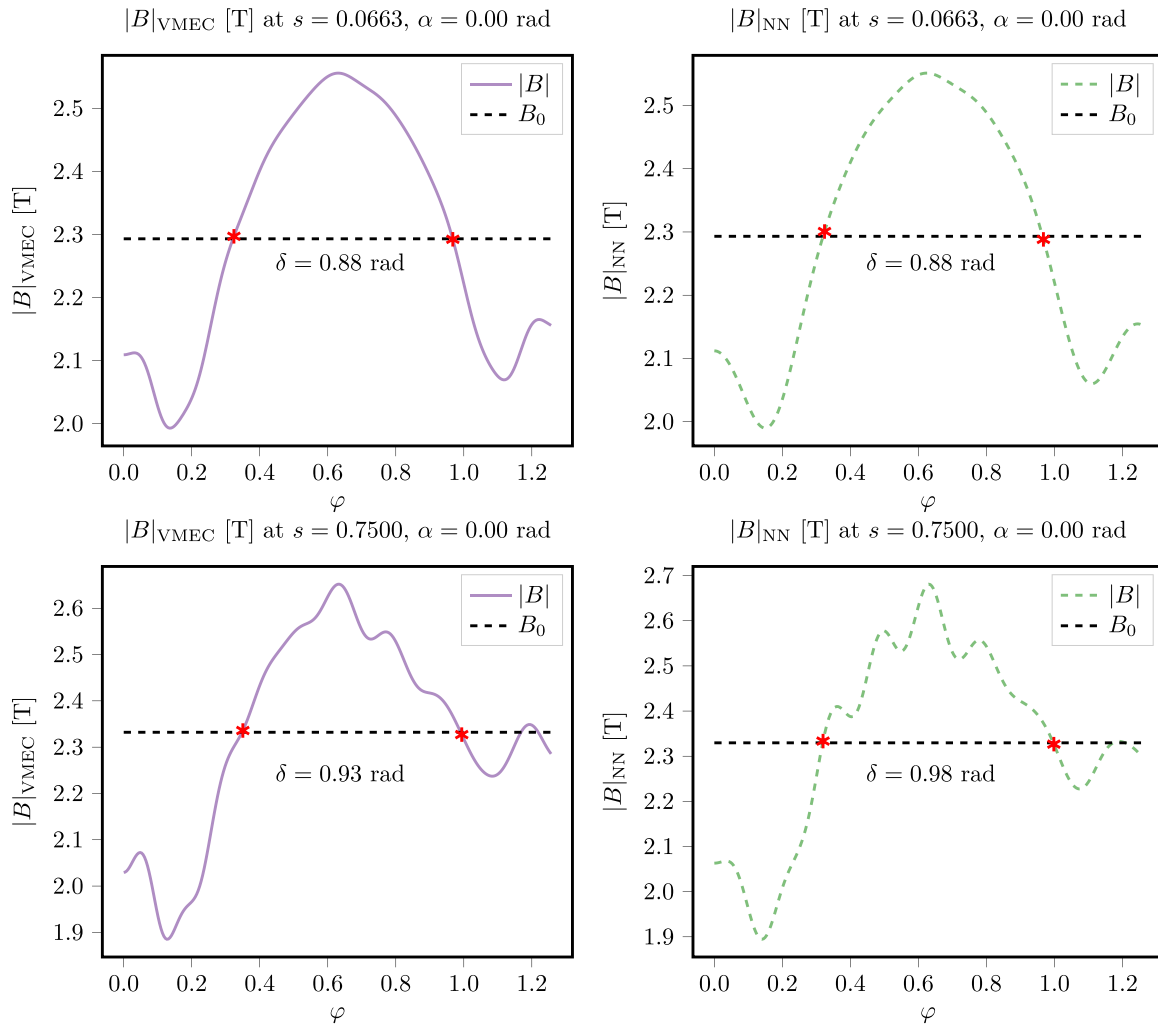
Therefore, also the second addenda of the first and third terms cancel each other. Then, the MHD force residual proxy is:

$$\begin{aligned} (2\pi)^2 f_* = \langle \mu_0\sqrt{g} F_s \rangle &= \chi' \frac{\partial \langle B_\theta \rangle}{\partial s} + \Phi' \frac{\partial \langle B_\varphi \rangle}{\partial s} \\ &+ \mu_0 p' \langle \sqrt{g} \rangle = 0. \end{aligned} \quad (85)$$

#### Boozer angular separation $\delta$ on test set

The evaluation of the reconstruction of the Boozer angular separation  $\delta$  on the test set is shown here. The median regressed equilibrium in the test set does not have a minimum of the magnetic well strength for  $\varphi = \frac{\pi}{N_{ip}}$ , therefore, the Boozer angular separation  $\delta$  depicted in figure 28 does not represent the angular separation between bounce points. Nevertheless, figure 28 shows how the model smoothly reconstructs the local magnetic field strength close to the magnetic axis, and it introduces artificial field ripples in the edge region.





**Figure 28.** True (left) and predicted (right) values for the magnetic field strength  $B$  in Boozer coordinates along a field line in case of the median predicted equilibrium in the test set. The  $\alpha = 0$  field line at  $s = 0.06$  (top) and  $s = 0.75$  flux surface (bottom) are depicted. The dashed black horizontal line indicates the average magnetic field strength on the flux surface  $B_0$ . The red crosses highlight the  $B_0$  contours, and the Boozer angular separation  $\delta(s, B_0, \alpha)$  between them is shown just beneath.

## ORCID iDs

Andrea Merlo <https://orcid.org/0000-0001-8359-2731>

Daniel Böckenhoff <https://orcid.org/0000-0003-1033-4648>

Jonathan Schilling <https://orcid.org/0000-0002-6363-6554>

Samuel Aaron Lazerson <https://orcid.org/0000-0001-8002-0121>

Thomas Sunn Pedersen <https://orcid.org/0000-0002-9720-1276>

## References

- [1] Seal S.K., Cianciosa M.R., Hirshman S.P., Wingen A., Wilcox R.S. and Unterberg E.A. 2017 Parallel reconstruction of three dimensional magnetohydrodynamic equilibria in plasma confinement devices *2017 46th Int. Conf. on Parallel Processing (ICPP) (Bristol, United Kingdom, 14–17 August 2017)* (IEEE) pp 282–91
- [2] Panici D., Conlin R., Dudt D.W. and Kolemen E. 2022 The DESC stellarator code suite part I: quick and accurate equilibria computations (arXiv:2203.17173 [physics])
- [3] Svensson J., Ford O., Werner A., Von Nessi G., Hole M., McDonald D.C., Appel L., Beurskens M., Boboc A. and Brix M. 2010 Connecting physics models and diagnostic data using Bayesian Graphical Models *37th EPS Conf. on Plasma Physics (Dublin, Ireland, 21–25 June 2010)* (European Physical Society) (available at: [https://pure.mpg.de/rest/items/item\\_2140911/component/file\\_3325641/content](https://pure.mpg.de/rest/items/item_2140911/component/file_3325641/content))
- [4] Svensson J. and Werner A. 2007 Large scale Bayesian data analysis for nuclear fusion experiments *2007 IEEE Int. Symp. on Intelligent Signal Processing (Alcala de Henares, Spain, 3–5 October 2007)* (IEEE) pp 1–6
- [5] Schilling J. 2018 Experimental MHD equilibrium analysis of magnetic configurations in the Wendelstein 7-X stellarator *Master Thesis IPP 2018-20* Max-Planck-Institute for Plasma Physics
- [6] Hoefel U. *et al* 2019 Bayesian modeling of microwave radiometer calibration on the example of the Wendelstein 7-X electron cyclotron emission diagnostic *Rev. Sci. Instrum.* **90** 043502

- [7] Grieger G. and Milch I. 1993 Das Fusionsexperiment WENDELSTEIN 7-X *Phys. Blätter* **49** 1001–5
- [8] Andreeva T. 2002 Vacuum magnetic configurations of Wendelstein 7-X (available at: [https://scholar.google.com/scholar?cluster=13899734479194957494&hl=it&as\\_sdt=0,5](https://scholar.google.com/scholar?cluster=13899734479194957494&hl=it&as_sdt=0,5))
- [9] Fable E. *et al* 2021 The modeling of a Tokamak plasma discharge, from first principles to a flight simulator *Plasma Phys. Control. Fusion* **64** 044002
- [10] Morris A.W., Akers R.J., Cox M., Militello F., Surrey E., Waldon C.W., Wilson H.R. and Zohm H. 2022 Towards a fusion power plant: integration of physics and technology *Plasma Phys. Control. Fusion* **64** 064002
- [11] Xanthopoulos P., Mynick H.E., Helander P., Turkin Y., Plunk G.G., Jenko F., Görler T., Told D., Bird T. and Proll J.H.E. 2014 Controlling turbulence in present and future stellarators *Phys. Rev. Lett.* **113** 155001
- [12] Landreman M. and Paul E. 2022 Magnetic fields with precise quasisymmetry for plasma confinement *Phys. Rev. Lett.* **128** 035001
- [13] Paul E.J., Landreman M., Bader A. and Dorland W. 2018 An adjoint method for gradient-based optimization of stellarator coil shapes *Nucl. Fusion* **58** 076015
- [14] Paul E.J., Landreman M. and Antonsen T. 2021 Gradient-based optimization of 3D MHD equilibria *J. Plasma Phys.* **87** 905870214
- [15] van Milligen B.P., Tribaldos V. and Jiménez J.A. 1995 Neural network differential equation and plasma equilibrium solver *Phys. Rev. Lett.* **75** 3594–7
- [16] Callaghan H.P., McCarthy P.J. and Geiger J. 1999 Fast equilibrium interpretation on the W7-AS stellarator using function parameterization *Nucl. Fusion* **39** 509–23
- [17] Callaghan H.P., McCarthy P.J. and Geiger J. 2000 Pressure profile recovery on W7-AS with idealized external magnetic measurements using function parameterization *Plasma Phys. Control. Fusion* **42** 1013–22
- [18] Sengupta A., McCarthy P., Geiger J. and Werner A. 2004 Fast recovery of vacuum magnetic configuration of the W7-X stellarator using function parameterization and artificial neural networks *Nucl. Fusion* **44** 1176–88
- [19] Sengupta A., Geiger J. and Carthy P.J.M. 2007 Statistical analysis of the equilibrium configurations of the W7-X stellarator *Plasma Phys. Control. Fusion* **49** 649–73
- [20] Merlo A., Böckenhoff D.F., Schilling J., Höfel U., Kwak S., Svensson J., Pavone A., Lazerson S.A. and Sunn Pedersen T. 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
- [21] Lu L., Jin P. and Karniadakis G.E. 2020 DeepONet: learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators (arXiv:1910.03193 [cs, stat])
- [22] Bauer F., Betancourt O. and Garabedian P. 1978 *A Computational Method in Plasma Physics* (Berlin: Springer)
- [23] Kruskal M.D. and Kulsrud R.M. 1958 Equilibrium of a magnetically confined plasma in a toroid *Phys. Fluids* **1** 265
- [24] Cybenko G. 1989 Approximation by superpositions of a sigmoidal function *Math. Control Signals Syst.* **2** 303–14
- [25] Hornik K. 1991 Approximation capabilities of multilayer feedforward networks *Neural Netw.* **4** 251–7
- [26] Pinkus A. 1999 Approximation theory of the MLP model in neural networks *Acta Numer.* **8** 143–95
- [27] Eldan R. and Shamir O. 2016 The power of depth for feedforward neural networks *29th Annual Conference on Learning Theory (23–26 June 2016) Conf. on Learning Theory* vol 49 (PMLR) pp 907–40 (available at: <https://proceedings.mlr.press/v49/eldan16>)
- [28] Lu Z., Pu H., Wang F., Hu Z. and Wang L. 2017 The expressive power of neural networks: a view from the width (4–9 December 2017) *Advances in Neural Information Processing Systems* vol 30 pp 6231–9 (available at: <https://proceedings.neurips.cc/paper/2017/hash/32cbf687880eb1674a07bf717761dd3a-Abstract.html>)
- [29] Llanas B., Lantarón S. and Sáinz F.J. 2008 Constructive approximation of discontinuous functions by neural networks *Neural Process. Lett.* **27** 209–26
- [30] Jagtap A.D., Kawaguchi K. and Karniadakis G.E. 2020 Adaptive activation functions accelerate convergence in deep and physics-informed neural networks *J. Comput. Phys.* **404** 109136
- [31] Chen T. and Chen H. 1995 Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems *IEEE Trans. Neural Netw.* **6** 911–7
- [32] Chen T. and Chen H. 1995 Approximation capability to functions of several variables, nonlinear functionals and operators by radial basis function neural networks *IEEE Trans. Neural Netw.* **6** 904–10
- [33] Lu L., Jin P., Pang G., Zhang Z. and Karniadakis G.E. 2021 Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators *Nat. Mach. Intell.* **3** 218–29
- [34] Poggio T., Banburski A. and Liao Q. 2019 Theoretical issues in deep networks: approximation, optimization and generalization (arXiv:1908.09375 [cs, stat])
- [35] Battaglia P.W. *et al* 2018 Relational inductive biases, deep learning, and graph networks (arXiv:1806.01261)
- [36] Hirshman S.P. and Whitson J.C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria *Phys. Fluids* **26** 3553
- [37] Lewis H.R. and Bellan P.M. 1990 Physical constraints on the coefficients of Fourier expansions in cylindrical coordinates *J. Math. Phys.* **31** 2592–6
- [38] Strand P.I. and Houlberg W.A. 2001 Magnetic flux evolution in highly shaped plasmas *Phys. Plasmas* **8** 2782–92
- [39] He K., Zhang X., Ren S. and Sun J. 2016 Deep residual learning for image recognition *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (26 June–1 July)* pp 770–8 (available at: [https://openaccess.thecvf.com/content\\_cvpr\\_2016/html/He\\_Deep\\_Residual\\_Learning\\_CVPR\\_2016\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html))
- [40] Geiger J., Beidler C.D., Feng Y., Maaßberg H., Marushchenko N.B. and Turkin Y. 2014 Physics in the magnetic configuration space of W7-X *Plasma Phys. Control. Fusion* **57** 014004
- [41] Erckmann V. *et al* 1997 The W7-X project: scientific basis and technical realization *17th IEEE/NPSS Symp. Fusion Engineering (Cat. No.97CH36131) (San Diego, CA, USA, 06–10 October 1997)* vol 1 pp 40–48
- [42] Neuner U. *et al* (The W7-X Team) 2021 Measurements of the parameter dependencies of the bootstrap current in the W7-X stellarator *Nucl. Fusion* **61** 036024
- [43] Bozhnikov S.A. *et al* 2020 High-performance plasmas after pellet injections in Wendelstein 7-X *Nucl. Fusion* **60** 066011
- [44] Gresho P.M. and Lee R.L. 1981 Don't suppress the wiggles—they're telling you something! *Comput. Fluids* **9** 223–53
- [45] Hirshman S.P. and Betancourt O. 1991 Preconditioned descent algorithm for rapid calculations of magnetohydrodynamic equilibria *J. Comput. Phys.* **96** 99–109
- [46] Morgan N. and Bourlard H. 1989 Generalization and parameter estimation in feedforward nets: some experiments *Advances in Neural Information Processing Systems* vol 2 (Morgan-Kaufmann)

- [47] Bengio Y., Louradour J., Collobert R. and Weston J. 2009 Curriculum learning *Proc. 26th Annual Int. Conf. on Machine Learning (Montreal Quebec Canada, 14–18 June 2009)* pp 41–48
- [48] Kingma D.P. and Ba J. 2017 Adam: a method for stochastic optimization (arXiv:1412.6980 [cs])
- [49] Loshchilov I. and Hutter F. 2017 Decoupled weight decay regularization (arXiv:1711.05101)
- [50] Hayashi T., Sato T. and Takei A. 1990 Three-dimensional studies of helical equilibria and magnetic surface breaking due to the finite beta effect *Phys. Fluids B* **2** 329–37
- [51] Beidler C. *et al* 1990 Physics and engineering design for Wendelstein VII-X *Fusion Technol.* **17** 148–68
- [52] Bozhakov S., Kazakov Y., Baldzuhn J., Laqua H.P., Alonso J.A., Beurskens M.N.A., Brandt C., Brunner K.J., Damm H. and Fuchert G. 2018 High density and high performance operation with pellet injection in W7-X 27th IAEA Fusion Energy Conf. (Gandhinagar, 22–27 October 2018) (available at: [https://inis.iaea.org/search/search.aspx?orig\\_q=RN:50052453](https://inis.iaea.org/search/search.aspx?orig_q=RN:50052453))
- [53] Lazerson S.A., Loizu J., Hirshman S. and Hudson S.R. 2016 Verification of the ideal magnetohydrodynamic response at rational surfaces in the VMEC code *Phys. Plasmas* **23** 012507
- [54] Zanini M. *et al* 2020 ECCD-induced sawtooth crashes at W7-X *Nucl. Fusion* **60** 106021
- [55] Zanini M. *et al* (The W7-X Team) 2021 Confinement degradation and plasma loss induced by strong sawtooth crashes at W7-X *Nucl. Fusion* **61** 116053
- [56] Boozer A.H. 2005 Physics of magnetically confined plasmas *Rev. Mod. Phys.* **76** 1071–141
- [57] Aleynikova K., Zocco A. and Geiger J. 2022 Influence of magnetic configuration properties on kinetic ballooning modes in W7-X *J. Plasma Phys.* **88** 905880411
- [58] Feng Y., Sardei F., Grigull P., McCormick K., Kisslinger J. and Reiter D. 2006 Physics of island divertors as highlighted by the example of W7-AS *Nucl. Fusion* **46** 807–19
- [59] Landreman M., Medasani B., Wechsung F., Giuliani A., Jorge R. and Zhu C. 2021 SIMSOPT: a flexible framework for stellarator optimization *J. Open Source Softw.* **6** 3525
- [60] Hoefel U. 2021 private communication
- [61] Greene J.M. 1997 A brief review of magnetic wells *Comments on Plasma Physics and Controlled Fusion* vol 17 (Princeton, NJ: Citeseer) pp 389–402
- [62] Freidberg J.P. 2014 *Ideal MHD* (New York: Cambridge University Press)
- [63] Drevlak M., Beidler C.D., Geiger J., Helander P. and Turkin Y. 2018 Optimisation of stellarator equilibria with ROSE *Nucl. Fusion* **59** 016010
- [64] Bauer F., Betancourt O. and Garabedian P. 1984 *Magnetohydrodynamic Equilibrium and Stability of Stellarators* (New York: Springer)
- [65] Carreras B.A., Dominguez N., Garcia L., Lynch V.E., Lyon J.F., Cary J.R., Hanson J.D. and Navarro A.P. 1988 Low-aspect-ratio torsatron configurations *Nucl. Fusion* **28** 1195–207
- [66] Nemov V.V., Kasilov S.V., Kernbichler W. and Heyn M.F. 1999 Evaluation of  $1/\nu$  neoclassical transport in stellarators *Phys. Plasmas* **6** 4622–32
- [67] Belli E.A. and Candy J. 2008 Kinetic calculation of neoclassical transport including self-consistent electron and impurity dynamics *Plasma Phys. Control. Fusion* **50** 095010
- [68] Belli E.A. and Candy J. 2011 Full linearized Fokker–Planck collisions in neoclassical transport simulations *Plasma Phys. Control. Fusion* **54** 015015
- [69] Nemov V.V., Kasilov S.V., Kernbichler W. and Leitold G.O. 2008 Poloidal motion of trapped particle orbits in real-space coordinates *Phys. Plasmas* **15** 052501
- [70] Drevlak M., Geiger J., Helander P. and Turkin Y. 2014 Fast particle confinement with optimized coil currents in the W7-X stellarator *Nucl. Fusion* **54** 073002
- [71] Bader A., Drevlak M., Anderson D.T., Faber B.J., Hegna C.C., Likin K.M., Schmitt J.C. and Talmadge J.N. 2019 Stellarator equilibria with reactor relevant energetic particle losses *J. Plasma Phys.* **85** 905850508
- [72] Bader A. *et al* 2020 Advancing the physics basis for quasi-helically symmetric stellarators *J. Plasma Phys.* **86** 905860506
- [73] Bader A., Anderson D.T., Drevlak M., Faber B.J., Hegna C.C., Henneberg S., Landreman M., Schmitt J.C., Suzuki Y. and Ware A. 2021 Modeling of energetic particle transport in optimized stellarators *Nucl. Fusion* **61** 116060
- [74] Hall L.S. and McNamara B. 1975 Three-dimensional equilibrium of the anisotropic, finite-pressure guiding-center plasma: theory of the magnetic plasma *Phys. Fluids* **18** 552–65
- [75] Cary J.R. and Shasharina S.G. 1997 Omnigenity and quasihelicity in helical plasma confinement systems *Phys. Plasmas* **4** 3323–33
- [76] Mikhailov M.I., Nührenberg J., Strumberger E. and Zille R. 2011 Collisionless fast-particle confinement in high- $\beta$  W7-X equilibria *Plasma Phys. Rep.* **37** 391–4
- [77] Wechsung F., Landreman M., Giuliani A., Cerfon A. and Stadler G. 2022 Precise stellarator quasi-symmetry can be achieved with electromagnetic coils *Proc. Natl Acad. Sci.* **119** e2202084119
- [78] Hirshman S.P., Spong D.A., Whitson J.C., Lynch V.E., Batchelor D.B., Carreras B.A. and Rome J.A. 1998 Transport optimization and MHD stability of a small aspect ratio toroidal hybrid stellarator *Phys. Rev. Lett.* **80** 528–31
- [79] Spong D.A., Hirshman S.P., Whitson J.C., Batchelor D.B., Carreras B.A., Lynch V.E. and Rome J.A. 1998 J\* optimization of small aspect ratio stellarator/tokamak hybrid devices *Phys. Plasmas* **5** 1752–8
- [80] Bergstra J., Yamins D. and Cox D. 2013 Making a science of model search: hyperparameter optimization in hundreds of dimensions for vision architectures 30th Int. Conf. on Machine Learning (16–21 June 2013) vol 28 (PMLR) pp 115–23 (available at: <https://proceedings.mlr.press/v28/bergstra13.html>)
- [81] Nocedal J. 1980 Updating quasi-Newton matrices with limited storage *Math. Comput.* **35** 773–82
- [82] Virtanen P. *et al* 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python *Nat. Methods* **17** 261–72
- [83] Albert C.G., Kasilov S.V. and Kernbichler W. 2020 Accelerated methods for direct computation of fusion alpha particle losses within, stellarator optimization *J. Plasma Phys.* **86** 815860201
- [84] Albert C.G., Kasilov S.V. and Kernbichler W. 2020 Symplectic integration with non-canonical quadrature for guiding-center orbits in magnetic confinement devices *J. Comput. Phys.* **403** 109065
- [85] Ferron J.R., Walker M.L., Lao L.L., John H.E.S., Humphreys D.A. and Leuer J.A. 1998 Real time equilibrium reconstruction for tokamak discharge control *Nucl. Fusion* **38** 1055–66
- [86] Yue X.N., Xiao B.J., Luo Z.P. and Guo Y. 2013 Fast equilibrium reconstruction for tokamak discharge control based on GPU *Plasma Phys. Control. Fusion* **55** 085016
- [87] Degraeve J. *et al* 2022 Magnetic control of tokamak plasmas through deep reinforcement learning *Nature* **602** 414–9
- [88] Yadan O. 2019 Hydra—a framework for elegantly configuring complex applications (available at: <https://github.com/facebookresearch/hydra>)

- [89] Hunter J.D. 2007 Matplotlib: a 2D graphics environment *Comput. Sci. Eng.* **9** 90–95
- [90] Harris C.R. *et al* 2020 Array programming with NumPy *Nature* **585** 357–62
- [91] Paszke A. *et al* 2019 PyTorch: an imperative style, high-performance deep learning library (8–14 December 2019) *Advances in Neural Information Processing Systems* vol 32 (Curran Associates, Inc.) (available at: <https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>)
- [92] Falcon W. 2019 Pytorch lightning vol 3 (<https://doi.org/10.5281/zenodo.3530844>)
- [93] Brand A., Allen L., Altman M., Hlava M. and Scott J. 2015 Beyond authorship: attribution, contribution, collaboration and credit *Learn. Publ.* **28** 151–5
- [94] Glorot X. and Bengio Y. 2010 Understanding the difficulty of training deep feedforward neural networks *Proc. 13th Int. Conf. on Artificial Intelligence and Statistics (JMLR Workshop Conf. Proc. vol 9) (3–15 May 2010)* pp 249–56 (available at: <http://proceedings.mlr.press/v9/glorot10a>)