

Reconstructions of electron-temperature profiles from EUROfusion pedestal database using turbulence models and machine learning

L.-P. Turica^{1,2,3} , A.R. Field², A.A. Schekochihin^{1,4} , L. Frassinetti⁵,
JET Contributors^a and the EUROfusion Tokamak Exploitation Team^b

¹Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford OX13PU, UK

²United Kingdom Atomic Energy Authority, Culham Campus, Abingdon OX14 3DB, UK

³University College, Oxford OX1 4BH, UK

⁴Merton College, Oxford OX1 4JD, UK

⁵KTH Royal Institute of Technology, Stockholm SE 100 44, Sweden

Corresponding author: L.-P. Turica, leonard-petru.turica@physics.ox.ac.uk

(Received 22 April 2025; revision received 26 July 2025; accepted 5 August 2025)

This study makes use of plasma-profile data from the EUROfusion pedestal database (Frassinetti *et al.* 2020 *Nucl. Fusion* vol. **61**, p. 016001), focusing on the electron-temperature and electron-density profiles in the edge region of H-mode ELMy JET ITER-Like-Wall (ILW) pulses. We make systematic predictions of the electron-temperature pedestal, taking engineering parameters of the plasma pulses and the density profiles as inputs. We first present a machine-learning (ML) algorithm which, given more inputs than theory-based modelling, is able to reconstruct unseen temperature profiles within 20 % of the experimental values. We find a hierarchy of the most consequential engineering parameters for such predictions. This result confirms the conceptual possibility of accurate data-driven prediction. Next, taking a simple theoretical approach that assumes a definite local relationship between the electron-density (R/L_{n_e}) and electron-temperature (R/L_{T_e}) gradients, we find that a range of power-law scalings $R/L_{T_e} = A(R/L_{n_e})^\alpha$ with $\alpha \approx 0.4$ correctly capture the behaviour of the electron-temperature in the steep-gradient region. Fitting A and α independently for each pedestal reveals a clear one-to-one correlation, suggesting an underlying constraint in pedestal physics. The measured $\eta_e = L_{n_e}/L_{T_e}$ values across the pedestal exhibit a wide distribution, significantly exceeding the slab-ETG linear stability threshold, implying either a non-linear threshold shift or a measurably supercritical saturated turbulent state. Finally, we fit parameters for scalings that relate the turbulent heat flux to the gradients R/L_{T_e} and R/L_{n_e} , similarly to models extracted from gyrokinetic simulations. The inclusion of more experimental parameters is necessary for such models to match the accuracy of our ML results.

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

^aSee the author list of Mailloux *et al.* (2022).

^bSee the author list of Joffrin *et al.* (2024).

1. Introduction

The operation of tokamaks with a divertor can result in an edge-transport barrier (ETB), which occurs once the heating power exceeds a certain threshold. The appearance of this ETB marks the onset of the high-confinement mode (the H-mode, see, e.g., Wagner *et al.* 1982 – as opposed to a lack of an ETB in the L-mode, see, e.g., Solano *et al.* 2022). This barrier results in enhanced confinement and correspondingly higher plasma temperatures and densities in the core region, which are desirable for the optimisation of fusion power. As a result, it is the H-mode that is also planned to be achieved in future tokamaks such as ITER.

The ETB is located just inside the edge of the confined plasma, which is identified by the last closed flux surface (LCFS) bounded by the separatrix. The ETB is characterised by high temperature and density gradients, so the plasma profiles appear to be ‘raised up’ upon a pedestal. The pedestal exhibits a stability limit, which constrains the pressure gradient (via unstable ballooning modes) and the bootstrap current (via unstable peeling modes) (Wilson *et al.* 1999). When this peeling-ballooning threshold is crossed, the equilibrium becomes subject to magnetohydrodynamic (MHD) instabilities that are thought to trigger edge-localised modes (ELMs). Peeling and ballooning instabilities extend radially on scales larger than the typical pedestal width (Snyder *et al.* 2004), and, therefore, the stability of these modes imposes a ‘global’ condition on the pedestal’s width and height. Such stability conditions are used by the semi-empirical EPED model (Snyder *et al.* 2009) in order to predict the pedestal width and height in fusion devices.

Alongside the MHD limitations on the plasma equilibrium, the pedestal profile is also believed to be determined by the transport caused by turbulence saturated in these steep-gradient regions. We assume that the action of this turbulence is local, i.e. that the gradients of the temperature and density are sources of free energy that linearly destabilise turbulent modes on micro-scales (much smaller than the extent of the pedestal), and that the nonlinear saturation of these unstable modes determines the effective transport at each location in the pedestal.

The exact relationship between the transport properties of such turbulence and the underlying equilibrium is still largely unknown. This represents a significant shortcoming in the modelling and prediction of tokamak plasmas: while the linear stability of a profile is generally well understood, this must be coupled consistently with energy and particle transport (turbulent, or otherwise) to predict experimental equilibria correctly; in particular, the pedestal is a boundary layer that significantly affects the overall confinement of the plasma. Here, we approach this question using data-driven methods, focusing on the pedestal structure. Previous studies have found that a distinctive feature of pedestals in several machines is a general radial shift between the density-pedestal location and the temperature-pedestal location: the steep-temperature-gradient region occurs further inside the LCFS than the steep-density-gradient region. This has been found on JET, ASDEX, and DIII-D by Dunne *et al.* (2016), Stefanikova *et al.* (2018), Wang *et al.* (2018), Frassinetti *et al.* (2020, 2021). Of these works, Frassinetti *et al.* (2020, 2021) used the EUROfusion database, which contains the pedestal characteristics (position, height, width) of over 2000 JET-C and JET-ILW type-I ELMy H-mode pulses, alongside their engineering and magnetic-equilibrium parameters. The pedestal profile data are obtained by compiling high-resolution Thomson-scattering (HRTS) measurements of density and

temperature pedestals, which are in a near-steady state. Because of the slow evolution time scale of the plasma profiles compared with that of turbulent fluctuations, the transport properties of the pedestal plasma ought to be set ‘quasistatically’ by a saturated state of the local turbulence.

Here, we will make use of the pedestal profile data for pulses in the EUROfusion database in order to identify correlations between the electron-temperature (T_e) and electron-density (n_e) profiles of 1251 JET-ILW pedestals and the corresponding pulse parameters. These pedestal profiles constitute information with which it should be possible to characterise the nature of edge transport in JET-ILW and the effect of this transport on confinement. In order to do this, we will reconstruct electron-temperature profiles using machine learning (ML) or a number of ‘physics-based’ turbulence models, and we will use the quality of these reconstructions to determine the accuracy of each method.

Section 2 provides an overview of the pedestal profiles and the functional fits that represent these profiles. Formal definitions of the pedestal height, width, and position are given in §2.2; we identify how these values correlate with the T_e and n_e profiles and the corresponding gradients in §2.3. In §2.4 we define reference positions at characteristic locations across the pedestal (the locations of the density-pedestal top, temperature-pedestal top, steep-gradient point, and separatrix). In §2.5 we describe a method to align these locations across the database. Section 3 explains our usage of the contents of the EUROfusion database. We justify the choice of the 1251-pulse subset of the database used in our analysis in §3.1, and then outline the database parameters and their values in §3.2. Appendix A defines these parameters formally.

Our first method of reconstructing pedestal T_e profiles, presented in §4, uses an ML approach. This is a natural step to take in order to use nonlinear correlations in the multidimensional parameter space of the EUROfusion pedestal database. We will conclude that it is indeed possible to infer electron-temperature profiles statistically, taking the electron-density profiles and various pulse parameters as inputs. The ML method fully exploits the information in the pedestal database. Therefore, it gives us an effective upper bound for the accuracy of such a pedestal prediction in our database using ML, physics-based models, or otherwise.

We build two different families of models, which generate ‘global’ and ‘local’ predictions. Global predictions, presented in §4.1, reconstruct the full temperature-pedestal profile using the knowledge of the full density-pedestal profile and different groups of database parameters. Local models, presented in §4.2, assemble the temperature profile from predictions of the local electron temperature, using the corresponding local electron density and various database parameters. The local approach allows us to mimic a local-transport model and compare the outcome to the global predictions of §4.1. In §4.3, we scan systematically through all combinations of database parameters and identify those that result in the most accurate predictions. This exhaustive approach of training a different network for each parameter combination reveals the most important parameters for a successful local or global prediction. Appendix B describes the numerical methods employed for our ML models.

As an alternative to ML, we turn to physics-based turbulence models to reproduce the electron-temperature pedestal. Section 5 depicts the loci of normalised gradients of the electron density (R/L_{n_e}) and electron temperature (R/L_{T_e}) over the pedestal region. With this information, we discuss the relevance of linear instabilities and

sources of turbulence in this physical regime. The plasma gradients are found to lie in a regime unstable to electron-temperature-gradient (ETG) modes. In particular, an experimentally relevant proportion of the turbulent transport is believed to be caused by slab-ETG modes (Field *et al.* 2023), which should have a strong influence on the nature of local turbulence, as we will discuss in § 5.1.

In § 6, we investigate the experimental distribution of $\eta_e \equiv L_{n_e}/L_{T_e}$, a parameter that governs the linear instability of slab-ETG modes. In § 6.1, we find that in the steep-gradient region of the pedestal, the mean value of this parameter is ≈ 2 . This mean value is consistent with past findings, however, we find a considerable spread in the values of η_e . In § 6.2, by assuming that the turbulence in the pedestal region can be described as ‘pinned’ to a nominal value of η_e , we reconstruct electron-temperature profiles by numerically inverting and integrating the relationship between the temperature and density gradients, with the density gradient as an input (as described in Appendix C). Doing this by minimising the discrepancy between the reconstructed and the experimental profile in the steep-gradient region yields an ‘effective’ $\eta_e = 1.9$ across the database. If, instead, the nominal η_e is allowed to vary between pulses, then this optimal pulse-dependent η_e agrees well with the measured local η_e in the steep-gradient region of the pedestal.

In § 7, we complicate the model by using the experimental distribution of R/L_{n_e} and R/L_{T_e} and an assumed relation between them of the form $R/L_{T_e} = A(R/L_{n_e})^\alpha$ to reconstruct the T_e -pedestal profiles, with A and α being fitting constants. Section 7.1 highlights that the exponent α maintains values below 1.5 across the pedestal when calculated locally. Such an approach can indeed produce correct plasma profiles, since any value of $\alpha \neq 1$ allows the gradient ratio η_e to adjust over the span of the pedestal. Section 7.2 identifies an effective value of $\alpha \approx 0.4$ required for an accurate full-database reproduction, however, most such reconstructions of the T_e profiles do not produce a pedestal, instead growing exponentially towards the core and failing to reproduce the correct gradient inside the temperature-pedestal top. If A and α are fit independently for each pulse, a clear relationship between A and α emerges, meaning that identifying either of these parameters for each pulse is sufficient for excellent reconstructions of the pedestal profile up to the position of the electron-temperature top.

Having found in § 4 that the power transported through the separatrix is a crucial parameter, we use the power balance in the pedestal and the knowledge of turbulent transport for more accurate reconstructions of T_e profiles in § 8. In § 8.1, we discuss recent results on modelling electron-channel transport in the pedestal by Guttenfelder *et al.* (2021), Chapman-Oplopoiou *et al.* (2022), Hatch *et al.* (2022, 2024), Farcaş *et al.* (2024). The models of turbulent transport by Guttenfelder *et al.* (2021) and Chapman-Oplopoiou *et al.* (2022) represent a physical picture of ETG transport: the electron-channel turbulent heat flux is thought to increase with the difference between the actual gradient ratio η_e and some nonlinear critical value of it, $\eta_{e,\text{crit}}$ (Field *et al.* 2023). More general models of heat transport, for which the physical picture is more uncertain, should still be considered because of the existence in the pedestal of ETG modes with low parallel wavenumbers (toroidal ETG) and other instabilities that survive suppression by flow shear, such as micro-tearing modes (MTMs) (Told *et al.* 2008; Parisi *et al.* 2020; Leppin *et al.* 2023). These other modes can alter the saturation of pedestal turbulence, and could therefore heavily affect the resulting heat fluxes. No comprehensive comparison between

the predictions of various heat-flux scaling laws and experimental data has been carried out before on the EUROfusion database. We carry this out in § 8.3 by fitting the free parameters of the models proposed by Chapman-Oploupoiou *et al.* (2022) and Hatch *et al.* (2022). We obtain models that yield far better pedestal reconstructions than those of §§ 6 and 7, however not better than the ML results of § 4, as evidenced in Appendix D. We also discuss the differences in the analytical forms of the heat-flux models and the consequences these differences have on the reconstructed profiles. We find that more recent heat-flux models, which also account for the local properties of the local magnetic field (Farcaş *et al.* 2024; Hatch *et al.* 2024), comply with the requirements for good T_e -profile reconstructions, showing promise for a study that has access to magnetic-field equilibrium reconstructions.

We provide a summary and a discussion of our results in § 9.

2. Pedestal profiles

For the purpose of this analysis, the electron temperature T_e and density n_e were measured for type-I ELMy JET-ILW pulses using the HRTS diagnostic (Pasqualotto *et al.* 2004). The HRTS laser path crosses the magnetic flux surfaces on the device midplane, from the outboard side through the centre of the plasma core, and back out. The HRTS diagnostic uses a laser beam fired at 20 Hz throughout the duration of the pulse. The scattered light from the beam is measured at different locations by detectors (Pasqualotto *et al.* 2004); this information was then used to determine the electron temperature and density at these locations (Frassinetti *et al.* 2012). The profiles that occur within the last 20 % of each inter-ELM period was then selected manually. Only inter-ELM periods longer than $2\tau_E$ were used, where the profiles of temperature and density reach near-steady state. Here, $\tau_E = W/P_{\text{Loss}}$ is the energy-confinement time, with W the plasma energy and P_{Loss} the total power loss. This data selection obtained profiles close to the peeling-ballooning stability threshold, which is reached many times over the duration of a pulse. Once this threshold is crossed, an ELM is triggered. More information regarding type-I ELMs and their relevance for the pedestal profiles can be found in § 3.1.

The profiles obtained from the HRTS diagnostic will be referred to as `raw` profiles (scatter points in figure 1a,b). These `raw` profiles were used by Frassinetti *et al.* (2020) to determine the deconvolved `mtanh` fits used in our analysis (which we will introduce in § 2.2). These `raw` profiles will be used in our analysis to estimate the experimental confidence intervals of the data for the ML methods described in § 4, avoiding the systematic biases of the `mtanh` fits.

2.1. Radial coordinates

The `raw` profiles and other calculated plasma profiles are cast as functions of the radial coordinate: either R (the radial distance from the axis of JET-ILW) or ψ_N (the poloidal-magnetic-flux label). The normalisation convention of ψ_N and the radial shift applied to each pulse will be described below.

Due to uncertainties in the exact underlying equilibrium and in the location of the separatrix, the HRTS profiles are shifted radially to pin them at a specific prescribed value of the electron temperature at the separatrix,

$$T_e(R^{\text{Sep}}) \equiv T_e^{\text{Sep}}, \quad (2.1)$$

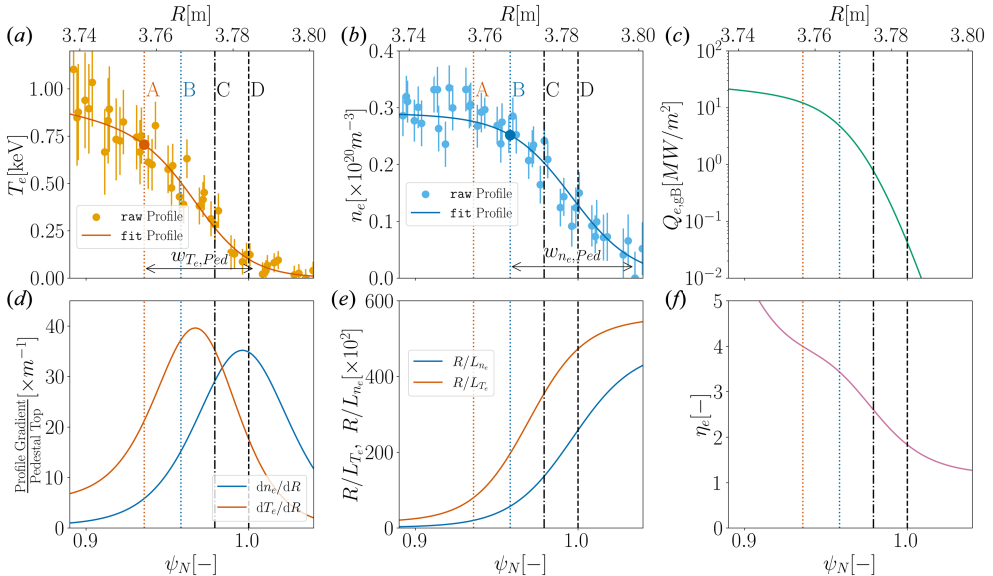


FIGURE 1. The fit (lines) and raw (dots) profiles for pulse #90339: (a) the electron density n_e and (b) the electron temperature T_e , with the pedestal-top points highlighted; (c) the gyro-Bohm heat flux $Q_{e,gB}$; (d) the normalised radial gradients of the profiles, $(dT_e/dR)/T_e^{\text{Ped}}$ and $(dn_e/dR)/n_e^{\text{Ped}}$, which are the gradients of a normalised \tanh fit with height 1 at the pedestal top; (e) the R/L_{T_e} and R/L_{n_e} profiles; (f) the $\eta_e \equiv L_{n_e}/L_{T_e}$ profile. The horizontal axes represent the normalised poloidal-magnetic-flux coordinate ψ_N (lower) and the radial position R (upper). These profiles are plotted alongside the four characteristic pedestal locations introduced in § 2.4 (vertical lines): A is the T_e top position $\psi_{T_e,\text{Top}}$, B is the n_e top position $\psi_{n_e,\text{Top}}$, C is the steep-gradient point ψ_{Steep} , D is the separatrix position ψ_{Sep} . The steep-gradient region is bounded by points B and D, and the exact middle of this region is marked by ψ_{Steep} (point C). We define the core as the locations inside of point A, and the SOL as those outside of D. The horizontal arrows in panels (a) and (b) represent the radial extent of the pedestal region for T_e and n_e .

a relationship used to determine the location of the separatrix R^{Sep} . The same separatrix location applies to both the electron-temperature and density profiles. The value of T_e^{Sep} is estimated to be 100 eV in JET-ILW using the two-point model for heat flux, as described in Chapter 5.2 of Stangeby (2000). The two-point model is a simplified model that describes the transport of heat along the scrape-off layer (SOL) between a point on the divertor and a point in the outboard plasma midplane. In this context, the model treats the SOL as being in a dominantly conductive (as opposed to convective) regime, which is consistent with the usage of stationary inter-ELM plasma profiles. This model yields a scaling of the SOL heat flux $Q_e \propto T_e^{7/2}(R^{\text{Sep}})$, which implies that T_e^{Sep} is very insensitive to changes in Q_e . The EDGE2D simulations consistently show that T_e^{Sep} varies in the range 80–110 eV (Simpson *et al.* 2019). As shown by Frassinetti *et al.* (2020), this uncertainty of the value of T_e^{Sep} has a negligible effect on the radial location of the separatrix as a result of the steepness of the profiles around $R = R^{\text{Sep}}$. This can also be generally inferred from the shapes of the electron-temperature profiles shown in figure 1(a).

The poloidal-magnetic-flux label ψ_N is a normalised poloidal-magnetic-flux coordinate $\psi(R)$. It is defined to be the poloidal-magnetic flux subtended between the

magnetic axis and the radial position R . Consistent with this definition, on the magnetic axis, $\psi_{\text{axis}} = 0$; and $\psi(R)$ is obtained by reconstructing the equilibrium magnetic field using the MHD code HELENA (Huysmans *et al.* 1999). For this, HELENA is provided with the plasma profiles and the shape of the plasma boundary reconstructed using EFIT (Lao *et al.* 1985). Then ψ_N is obtained by normalising $\psi(R)$ so that $\psi_N = 1.0$ at the radial location of the separatrix, *viz.*,

$$\psi_N = \frac{\psi(R)}{\psi(R^{\text{Sep}})}. \quad (2.2)$$

2.2. *mtanh* fits

The pedestal profiles that are the subject of our analysis are functions fit to the measured raw profiles. The general form of the fits is

$$\text{mtanh}(x, h_{\text{Ped}}, s_{\text{Core}}) = \frac{h_{\text{Ped}}}{2} \left[\frac{(1 + x s_{\text{Core}})e^x - e^{-x}}{e^x + e^{-x}} + 1 \right], \quad (2.3)$$

where $x = (\psi_{\text{Pos}} - \psi_N)/w_{\text{Ped}}$ and all the free parameters will be defined below. Such *mtanh* fits are commonly used in many studies of H-mode physics and, thus, they offer the advantage of a generalised convention for pedestal profiles, while being differentiable fits with relatively few free parameters. These free parameters also represent formal definitions of several important properties of the pedestal. Namely, these are: the profile value at the top of the pedestal h_{Ped} ($= T_e^{\text{Ped}}$ and $= n_e^{\text{Ped}}$ for the temperature and density, respectively), the position of the middle of the pedestal ψ_{Pos} , the width of the pedestal w_{Ped} ($= w_{T_e, \text{Ped}}$ and $= w_{n_e, \text{Ped}}$), and the position of the pedestal top $\psi_{\text{Top}} = \psi_{\text{Pos}} - w_{\text{Ped}}/2$ ($= \psi_{T_e, \text{Top}}$ and $= \psi_{n_e, \text{Top}}$). The core slope to which the profile tends at $\psi_N \ll \psi_{\text{Top}}$ is determined by $h_{\text{Ped}}s_{\text{Core}}/2w_{\text{Ped}}$, parametrised by s_{Core} . The profiles derived using an *mtanh* fit, as per (2.3), will be referred to as *fit* profiles, and they represent the nominal plasma profiles used for the EUROfusion database (Frassinetti *et al.* 2020).

Importantly, the fits are performed taking into account the finite radial resolution of the HRTS diagnostic, as described by Frassinetti *et al.* (2012). Therefore, each *fit* profile is computed from the raw profile using deconvolved fits. For this, the *fit* electron-density profile is obtained by minimising the residual between the raw n_e values and the *fit* n_e profile, convolved with the spatial instrument function of the HRTS diagnostic. Then, the *fit* electron-temperature profile is obtained by minimising the residual between the raw T_e values and the *fit* T_e profile weighed by the *fit* n_e profile and convolved with the instrument function of the HRTS diagnostic. This is because the density profiles affect the inference of the temperature from the HRTS data. The instrument function of HRTS represents a ‘scattering kernel’ with a full-width-half-maximum of approximately 11 mm in the pedestal region. The scatter is caused by the HRTS diagnostic imaging a finite volume of the plasma around each measurement location, so the variations of the profile across this volume must be taken into account.

2.3. Gradients

Fitting the raw profiles to a differentiable function allows for gradients to be meaningfully calculated. The gradients are an essential ingredient for studying the local quasi-equilibrium that the plasma settles to during the inter-ELM period. They

must be related to the heat flux crossing the pedestal because they provide the drive for the turbulence that controls transport across the pedestal. It is important to acknowledge that the choice of an `mTanh` fit is a source of systematic bias in estimating the gradients of the profiles, however, this fit has been used to characterise the properties of pedestal shapes across devices for more than 20 years, which helps make our study easily generalisable to other databases.

The gradients of the T_e and n_e profiles are calculated with respect to the radial coordinate R . For any quantity $A(R)$, we define

$$L_A = A(R) \left(\frac{dA}{dR} \right)^{-1} \quad (2.4)$$

as the corresponding gradient scale length.

Figure 1 is an overview of pulse #90339: the plasma profiles (panels *a* and *b*), the value of the gyro-Bohm heat flux (shown in panel *c* and defined below), the profile gradients (panels *d* and *e*) and the gradient-length-scale ratio $\eta_e = L_{n_e}/L_{T_e}$ (panel *f*). The radial gradients dT_e/dR and dn_e/dR in figure 1(*d*) are normalised to the pedestal-top values T_e^{Ped} and n_e^{Ped} , respectively. This normalisation is done in order to show the qualitative behaviour of the raw gradients. For our analysis throughout this study, we will focus on the gradient length scales L_{T_e} and L_{n_e} , defined in (2.4). The most important trends manifest in figure 1 are the increase by approximately two orders of magnitude of the normalised gradients R/L_{T_e} and R/L_{n_e} towards the separatrix, and the typical values of $\eta_e = \mathcal{O}(1) - \mathcal{O}(10)$. The finite values of R/L_{T_e} and R/L_{n_e} outside the separatrix region is a result of the `mTanh` convention that all fit profiles tend to zero in the SOL. These tendencies of the gradients of pulse #90339 are qualitatively representative of the behaviour of most pulses across the database.

The gyro-Bohm heat flux is defined as

$$Q_{e,\text{gB}} = n_e T_e v_{\text{th},e} \frac{\rho_e^2}{R^2}, \quad (2.5)$$

where $v_{\text{th},e} = \sqrt{2T_e/m_e}$ is the electron thermal velocity, $\rho_e = m_e v_{\text{th},e}/eB$ is the electron gyroradius, m_e is the electron mass, e is the electron charge, and B is the magnitude of the magnetic field. The gyro-Bohm heat flux (2.5) maintains values of $\mathcal{O}(\text{MW m}^{-2})$ across the profiles inside of the half-way point between the density-pedestal top and and separatrix (defined as the ‘steep-gradient point’ in § 2.4); $Q_{e,\text{gB}}$ decreases rapidly towards the separatrix, approaching 0 together with the T_e and n_e profiles.

2.4. Pedestal regions and relevant locations

It is useful to define formally some important radial locations and regions for our profiles. Using the parameters of the `mTanh` fits in ψ_N coordinates, the pedestal-top positions of T_e and n_e are $\psi_{T_e,\text{Top}}$ and $\psi_{n_e,\text{Top}}$, respectively. These pedestal-top positions are defined as

$$\psi_{\text{Top}} = \psi_{\text{Pos}} - \frac{w_{\text{Ped}}}{2}. \quad (2.6)$$

We refer to the region $\psi_N < \psi_{T_e,\text{Top}}$ as the core of the plasma. This definition with respect to $\psi_{T_e,\text{Top}}$ holds across the database as a result of the relative shift between the tops of the n_e and T_e pedestals (Frassinetti *et al.* 2021): the top of the density

pedestal is radially further out than the top of the temperature pedestal. Hence, the low temperature gradients in the regions within $\psi_{T_e, \text{Top}}$ guarantee low electron-density gradients. We will also use this pedestal-shift property in § 2.5: the separatrix position will be conventionally placed at $\psi_N = \psi_{\text{Sep}} = 1.0$ such that $T_e = 100$ eV as per the normalisation described in § 2.1, and, consequently, the region of $\psi_N > 1.0$ is referred to as the SOL.

The region between the separatrix and $\psi_{T_e, \text{Top}}$ is the pedestal. Because most of the pulses obey $\psi_{T_e, \text{Top}} < \psi_{n_e, \text{Top}} < 1.0$, the interval between $\psi_{n_e, \text{Top}}$ and 1.0 will be referred to as the steep-gradient region, since both the electron-temperature and density gradients are very large there. In the interval between $\psi_{T_e, \text{Top}}$ and $\psi_{n_e, \text{Top}}$, the electron-density gradient is much smaller, which justifies making a qualitative distinction between these two neighbouring regions of the pedestal. We will identify this distinction again in § 5, and we will discuss the physical difference between the two regions in the sections thereafter.

The profile values at the top of the T_e pedestal are T_e^{Ped} and $n_e(\psi_{T_e, \text{Top}})$; at the top of the n_e pedestal they are $T_e(\psi_{n_e, \text{Top}})$ and n_e^{Ped} . These pedestal-top points (large dots in figure 1a,b) are the locations where the gradient scale lengths change from the high-gradient regime of the pedestal to the low-gradient regime characteristic of the core. As a representative location in the steep-gradient region, we define the ‘steep-gradient point’ location as

$$\psi_{\text{Steep}} = \frac{\psi_{n_e, \text{Top}} + 1}{2}, \quad (2.7)$$

i.e., half-way between the electron-density top and the separatrix. This location will be considered representative of the steepest region of the pedestal, as it will be far enough away from $\psi_{n_e, \text{Top}}$ and $\psi_{T_e, \text{Top}}$ to not be influenced by core turbulence. Further out from the steep-gradient point, towards (or beyond) the separatrix, the uncertainties of R/L_{n_e} and R/L_{T_e} become comparable to the values of the gradients themselves, as obtained by uncorrelated variation of the mtanh parameters within their confidence intervals (Field *et al.* 2023). This is because the normalisation of the gradients against the asymptotically vanishing values of n_e and T_e in the SOL results in diverging errors. We shall, therefore, use the steep-gradient point as indicative of the pedestal’s steep-gradient region.

2.5. Rescaled radial coordinate

In order to group together qualitatively similar regimes of the pedestal, we use the landmarks that we defined above, i.e., the separatrix position ψ_{Sep} , the density-pedestal-top position $\psi_{n_e, \text{Top}}$ and the temperature-pedestal-top position $\psi_{T_e, \text{Top}}$, to define a new ‘renormalised’ radial coordinate X_{Ren} . This will map each point in the density-top and steep-gradient regions to a value X_{Ren} that depends on its position relative to these landmarks.

Namely, a point whose flux label is $\psi_N \in [\min(\psi_{T_e, \text{Top}}, \psi_{n_e, \text{Top}}), 1]$ will have the radial coordinate $X_{\text{Ren}} \in [0, 2]$ defined by

$$X_{\text{Ren}}(\psi_N) = \begin{cases} \frac{\psi_N - 1}{\psi_{n_e, \text{Top}} - 1} & \text{if } 1 \geq \psi_N \geq \psi_{n_e, \text{Top}}, \\ 1 + \frac{\psi_N - \psi_{n_e, \text{Top}}}{\psi_{T_e, \text{Top}} - \psi_{n_e, \text{Top}}} & \text{if } \psi_{T_e, \text{Top}} \geq \psi_N > \psi_{n_e, \text{Top}}. \end{cases} \quad (2.8)$$

This is a piecewise linear transformation that maps ψ_{Sep} to 0, $\psi_{n_e, \text{Top}}$ to 1, $\psi_{T_e, \text{Top}}$ to 2, and ψ_{Steep} to 0.5. Because of the relative shift between the pedestal-top locations in the database (Frassinetti *et al.* 2020), all but 22 pedestals obey the pedestal-shift condition ($\psi_{n_e, \text{Top}} > \psi_{T_e, \text{Top}}$) in the subset relevant to this study, as per § 3.1. For the pedestals that do not obey this condition, the map (2.8) will place the entire pedestal profile in the $[0, 1]$ interval. The map (2.8) will be used in §§ 6 and 7 to uncover the overall structure of pedestal gradients.

More complex functional forms of X_{Ren} that maintain the separation of ψ_{Sep} , $\psi_{n_e, \text{Top}}$, and $\psi_{T_e, \text{Top}}$, but are also, e.g., differentiable at $X_{\text{Ren}} = 1$, offer little advantage while being more cumbersome.

3. Database and parameters

The database used for this study is the EUROfusion pedestal database described in Frassinetti *et al.* (2020). This database contains the parameters of over 2000 ELMy H-mode pulses from JET-C (carbon wall) and JET-ILW (ITER-like wall, with beryllium main chamber and tungsten divertor plates). The pulse parameters that we will use are the fit coefficients of mtanh functions as described in § 2.2. We narrow the scope of the database to a subset of 1251 pulses, based on considerations of the plasma equilibrium, as explained below.

3.1. Data selection and ELMs

In order to contextualise the data representing our H-mode pulses, it is useful to understand how the pedestals in the database are measured and how they relate to the stability of ELMs.

The pedestal profiles are constrained by the peeling–ballooning MHD stability limit, which depends on the pressure gradient and the current density (Wilson *et al.* 1999). Edge-localised modes are thought to be triggered when a plasma profile strays past this stability limit, and they result in a sudden reduction of the energy density on spatial scales larger than the size of the pedestal (Snyder *et al.* 2004). An ‘ELM cycle’ contains the ELM itself and the recovery period of the plasma profile until the occurrence of the next ELM. The classification of ELMs depends on their frequency, associated energy loss and on whether the pedestal is peeling-unstable or ballooning-unstable (Connor, Kirk & Wilson 2008; Murari *et al.* 2014). ‘Giant’ type-I ELMs are those occurring throughout the pulses in our database.

The EUROfusion pedestal database contains data from type-I ELMy H-mode pulses in JET, and the pedestal profiles are taken from measurements from the last 20% of each ELM cycle, as described in § 2. This choice ensures that the pedestal profiles are in near-steady state over the course of the measurement, and that they represent optimal operation of the plasma, just below the ELM stability threshold.

While the EUROfusion pedestal database contains a wide range of pulses, we will narrow the scope in order to limit the variation of the underlying MHD instability of the plasma. We will focus on deuterium-only plasmas (so the effective ion mass will be $M_{\text{eff}} = 2$ atomic units). Furthermore, methods of ELM control or mitigation affect the equilibrium of the plasma, meaning that plasma pressure is reduced to a subcritical, ELM-stable regime. Therefore, we chose pulses that had no methods of ELM control or mitigation applied, such as impurity seeding (added impurities producing a regime with high-frequency, low-amplitude ELMs), pellet injection (a method of plasma fuelling inside the ETB, triggering ELMs) or vertical kicks (vertical displacements of the plasma MHD equilibrium, also triggering ELMs).

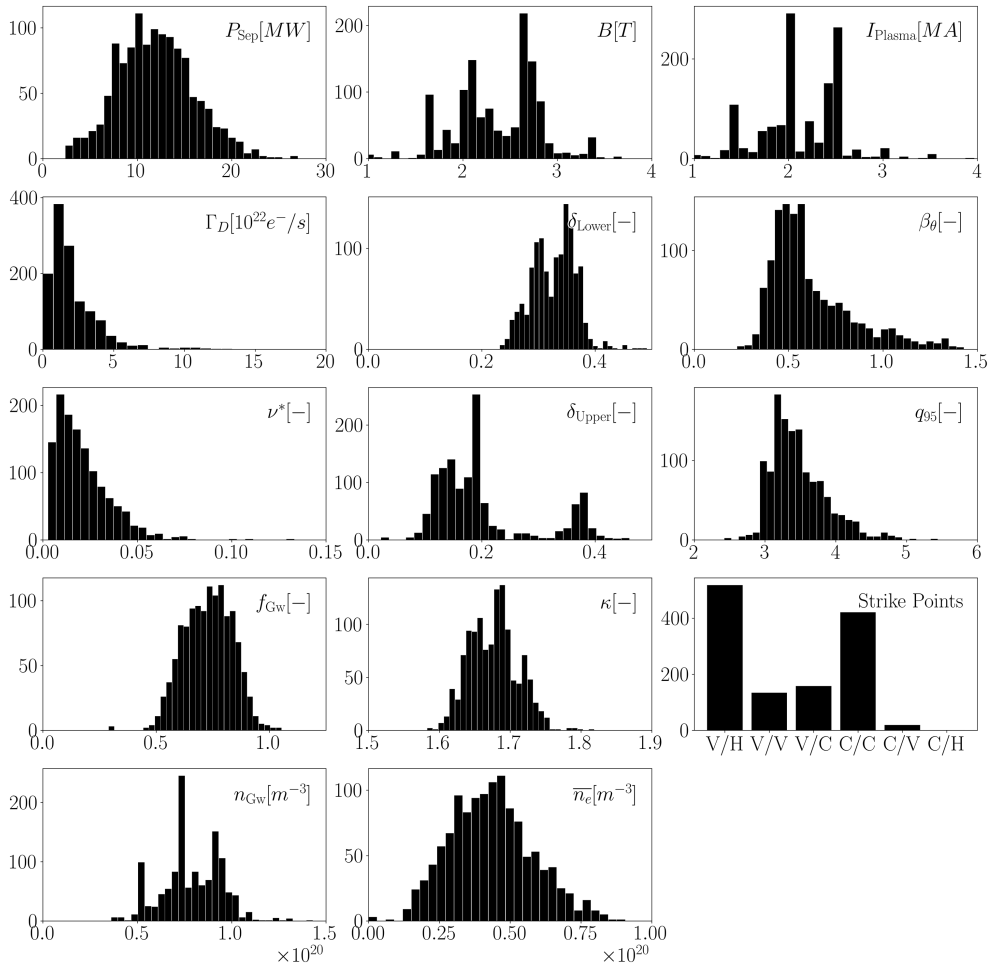


FIGURE 2. Histograms of the distributions of the engineering and equilibrium parameters described in [Appendix A.1](#). The number of pulses with each strike-point configuration is also shown. The strike points are formatted as ‘inner/outer’, as explained in [Appendix A.1](#).

This means that, for the current analysis, we will work with a subset of the database containing 1251 pulses.

3.2. Parameter values

The database also contains a wide range of experimental parameters, so we limit ourselves to a physically relevant subset of engineering parameters and magnetic-equilibrium-reconstruction parameters. The list of these parameters is provided in [Appendix A.1](#) and an overview of their distribution is given in [figure 2](#).

Because of the wide scatter in the values of these parameters, we anticipate the necessity for data augmentation ([Appendix B.2](#)) for effective training of T_e -predicting neural networks in [§ 4](#). We note that the poloidal beta, β_θ , and the electron-electron collisionality, ν^* , are physically important in determining the nature of local turbulence in the pedestal; however, the calculation of these parameters directly references T_e^{Ped} , as per [Frassinetti *et al.* \(2020\)](#). This becomes an issue as the neural networks of




FIGURE 3. Histograms of the distributions of relevant T_e and n_e profile values. Their definitions are given in § 2.2.

§ 4 manage to trivially invert the calculation of these parameters. This results in excellent T_e predictions with *a priori* knowledge of the T_e pedestal. The inclusion of such information into our neural networks defeats the point of our study, although they are worth mentioning amongst the parameters that are usually considered important for characterising the physical environment of a pedestal.

Some of these parameters will also be used in § 8 and Appendix D for predictions using heat-flux scalings and for the calculation of $Q_{e, \text{GB}}$.

For our analysis of the plasma profiles themselves, we will make use of the raw and fit profiles described in § 2. These profiles are not normally a part of the EUROfusion database, so we sourced them directly from processed pulse files (PPFs) for each pulse. We have access to the electron-density and temperature profiles and the corresponding R and ψ_N for each pulse, as described in § 2.

The distribution of `mtanh` parameters (defined in § 2.2) and other relevant values of the T_e and n_e profiles are shown in figure 3. The value of the temperature at the density-pedestal top, $T_e(\psi_{n_e, \text{Top}})$, will be an important measure of the accuracy of our temperature-pedestal reconstruction over the steep-gradient region. The density at the separatrix, $n_e^{\text{Sep}} = n_e(\psi_N = 1.0)$, and the ratio $n_e^{\text{Sep}}/n_e^{\text{Ped}}$ between the separatrix density and the pedestal-top density are strongly correlated with temperature-pedestal parameters, as was anticipated in Field *et al.* (2023) and will be confirmed in figure 22.

4. Neural-network predictions

The high-dimensional parameter space of the EUROfusion pedestal database together with the pulse profiles present an opportunity for ML to identify complex,

nonlinear correlations. Before attempting to use physics-based models, either stemming from theoretical considerations or from numerical experiments (§§ 5–8), we will leverage neural networks to predict T_e profiles using n_e profiles and selected database parameters.

This section will detail the prediction of T_e profiles using various input parameter groups, enabling us to quantify the contributions of specific engineering and equilibrium-reconstruction parameters to the accuracy of predictions. We will prove that such a prediction is indeed possible, and find the highest accuracy to which T_e profiles can be reconstructed using an ML interpolation method over our data, effectively establishing also the maximal standard of accuracy of such reconstruction by any method. We will also explore the difference between predictions of T_e carried out across the full-pedestal profiles versus those over narrow windows of the pedestal profile, the latter of which mimic predictions using some model of local turbulence. Finally, we will find the pulse parameters in our database that are most important for the correct prediction of such T_e profiles.

For our purposes, neural networks will act as a generalised fit to the experimental data, with the advantage that they allow for arbitrary complexity and easy inclusion of many parameters into such a fit. Simple analytical forms that have a physical prior will be fit to our data in § 8. However, we expect that any high-complexity or many-parameter analytical form resulting from such a dataset would not be any more transparent than a ML ‘black box’ that will be obtained below.

In [Appendix B.1](#), we explain the form of the data used in training our networks. Importantly, we will use T_e , n_e , and R/L_{n_e} profiles from the radial interval $R = [3.7, 3.85]$ m, sampled at evenly distributed locations in R . Once a new T_e profile is computed, a new `mtanh` fit is constructed in order to obtain the pedestal height, width, and top position. These are used as metrics of agreement with the experimental profile (for which they were derived from a deconvolved `mtanh` fit, as explained in § 2).

Two baseline requirements for using ML for such predictions are that the database parameter space is populated densely enough and that experimental uncertainties are low enough in order to be able to infer correct trends from the data. We deal with the issues of parameter-space sparsity and experimental uncertainty by using data augmentation in order to obtain robust methods to predict the T_e profiles. Methods and criteria for data augmentation are explained in [Appendix B.2](#).

Our choices of the model loss, the train-test split, and network architecture are described in [Appendices B.3–B.5](#). In particular, for all predictions, the database is split into a training set of 80 % and a testing set of 20 % of the database. The same split is maintained for consistency across all the presented results, and all our results will be shown for a testing set that is unseen during training.

Initial tests revealed that the electron collisionality ν^* and poloidal beta β_θ artificially improved accuracy of predictions because they included exact information about the T_e pedestal. These parameters were originally obtained using a specific reference to T_e^{Ped} (Frassinetti *et al.* 2020). Therefore, we have excluded these parameters to avoid circular reasoning and maintain the integrity of the predictive task. Similarly, ψ_N is omitted altogether because its normalisation explicitly demands that $T_e^{\text{Sep}} = 100$ eV.

Parameters	Full-pedestal groups			Local-value groups				
	All	No-EFIT	EFIT	None	All	No-R	No-EFIT	None
P_{Sep}	×	×			×	×	×	
Γ_D	×	×			×	×	×	
n_{Gw}	×	×			×	×	×	
\bar{n}_e	×	×			×	×	×	
I_{Plasma}	×	×	×		×	×	×	
B	×	×	×		×	×	×	
Strike points	×	×	×		×	×		
Upper δ	×		×		×	×		
Lower δ	×		×		×	×		
q_{95}	×		×		×	×		
κ	×		×		×	×		
R	×	×	×	×	×		×	×
n_e	×	×	×	×	×	×	×	×
R/L_{n_e}					×	×	×	×

TABLE 1. The parameter groups and the respective included parameters for the neural networks informed of the full n_e pedestal profile (§ 4.1) and of the local values of the n_e profile (§ 4.2).

4.1. Full-profile electron-temperature predictions

We first make predictions of the full temperature-pedestal profile using the information from the full density-pedestal profile and various database parameters. Randomly located radial cropping is used for the n_e profile during augmentation to reduce bias, ensuring that networks remain independent of specific radial data points while preserving overall profile features. The random crop keeps 80 % of the pedestal information, i.e., L is close to the full-pedestal extent, as detailed in figure 23 of Appendix B.2.

Four groups of parameters are defined for this comparison (found under the heading ‘Full-pedestal groups’ in table 1). Figure 4 illustrates the prediction results across these parameter groups, showing the mtanh parameters that describe the reconstructed T_e pedestals for the 251 pulses in the test set. The results in figure 4 correspond to the groups of parameters listed in table 1 as follows:

- (a-d) ‘None’ – only the radial coordinate and n_e profile, with no additional pulse parameters;
- (e-h) ‘No-EFIT’ – R and n_e profiles and engineering parameters only, excluding magnetic-equilibrium-reconstruction parameters;
- (i-l) ‘EFIT’ – R and n_e profiles and magnetic-equilibrium-reconstruction parameters, excluding engineering parameters;
- (m-p) ‘All Parameters’ – all available engineering and EFIT parameters alongside R and n_e profiles.

Figure 4 and similar diagnostic plots in what follows display the coefficient of determination

$$r^2 = 1 - \frac{S_{\text{res}}}{S_{\text{tot}}}, \quad (4.1)$$

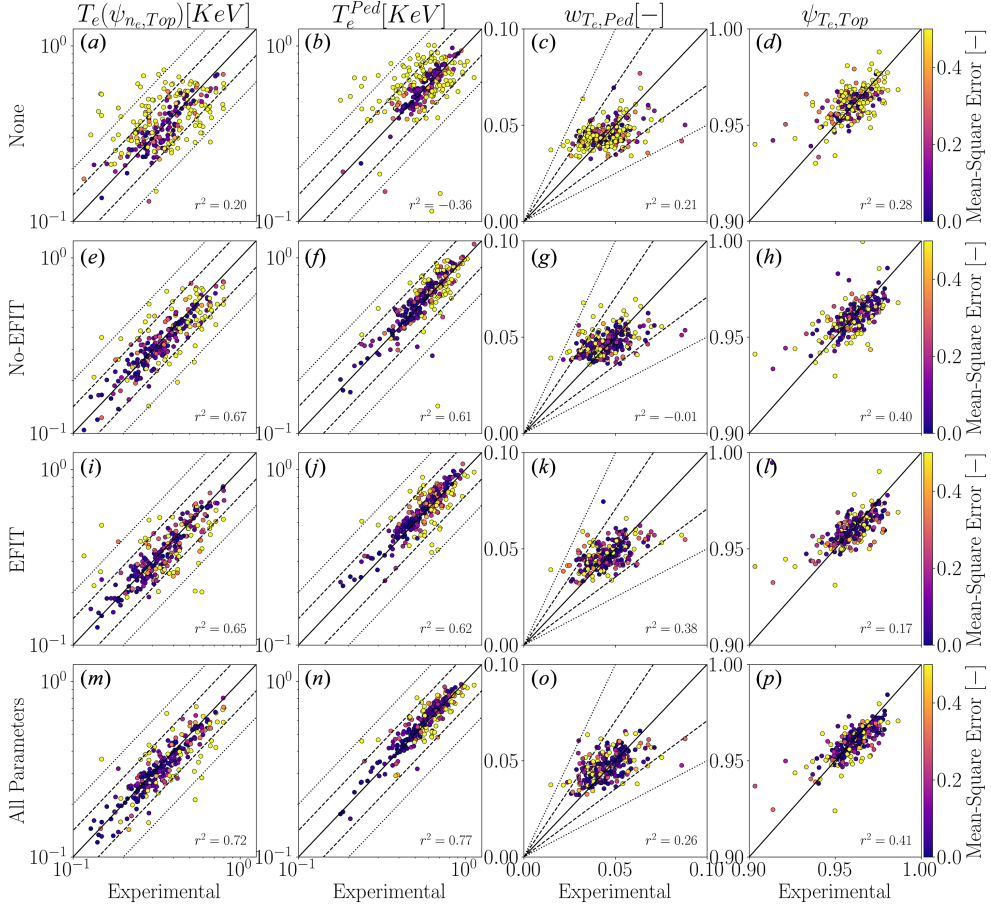


FIGURE 4. Database predictions using neural networks informed of the full density pedestal. These correspond to each of the four groups of parameters listed in table 1: (a–d) no database parameters used as inputs, with the exception of R and n_e ; (e–h) only engineering parameters alongside R and n_e ; (i–l) only the parameters relevant to the MHD-equilibrium reconstruction, alongside R and n_e ; (m–p) all of the aforementioned information used. The colour bar represents the mean-square error between the predicted T_e and the experimental fit profile. For each pedestal prediction, a new mTanh fit is performed in order to establish the T_e -pedestal parameters. Comparisons of the prediction accuracy of these mTanh parameters are provided (a,e,i,m) for T_e at the inner edge of the steep-gradient region, (b,f,j,n) for T_e at the top of the temperature pedestal, (c,g,k,o) for the temperature-pedestal widths, (d,h,l,p) for the temperature-pedestal locations. The solid lines represent the correct prediction, the dashed lines represent a discrepancy of a factor of $\sqrt{2}$, and the dotted lines represent a discrepancy of a factor of 2. The text on the bottom right of each plot signifies the respective r^2 coefficient of determination, with a value of 1 indicating a perfect prediction and lower values indicating less accurate predictions.

where, for a diagnostic $y \in \{T_e(\psi_{n_e, \text{Top}}), T_e^{\text{Ped}}, w_{T_e, \text{Ped}}, \psi_{T_e, \text{Top}}\}$, $S_{\text{res}} = \sum (y_{\text{meas}} - y_{\text{pred}})^2$ is the sum of the prediction residual, $S_{\text{tot}} = \sum (y_{\text{meas}} - \langle y_{\text{meas}} \rangle)^2$ is the total variance of the measured data, and $\langle y_{\text{meas}} \rangle = (1/N) \sum y_{\text{meas}}$ is the average experimental value. Here, y_{meas} is the experimentally measured value of a diagnostic y and y_{pred} is its predicted value. A value of r^2 close to 1 suggests an excellent reconstruction, low values suggest a poor agreement. These values serve as our measure

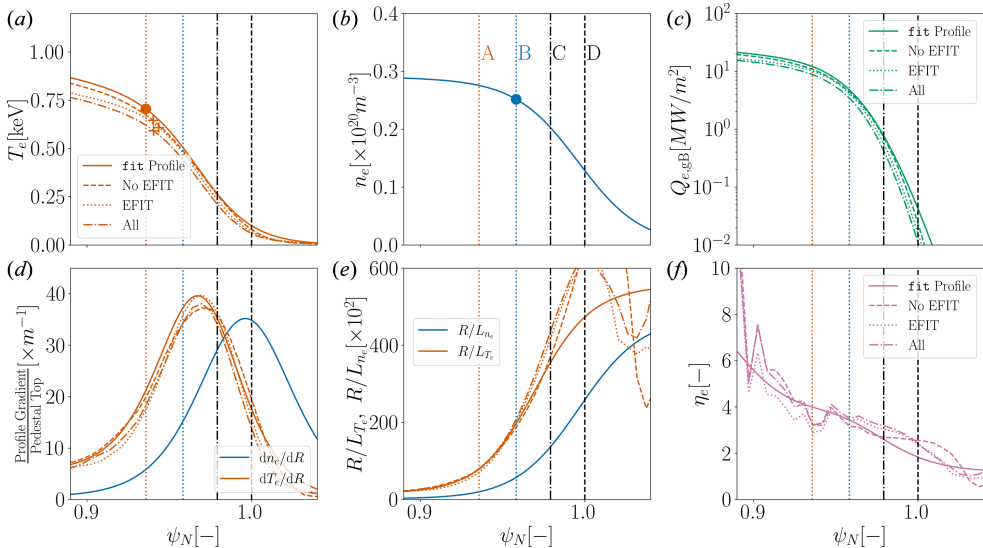


FIGURE 5. Pulse #90339 fit profiles compared with neural-network predictions of the T_e profile over the full pedestal. The layout is the same as in figure 1. The fit profiles are shown as solid lines, while the predictions using different parameter groups listed in table 1 are shown as dash-dotted (‘All Parameters’), dashed (‘No-EFIT’), and dotted (‘EFIT’) lines. The vertical lines and dots represent the landmarks of the fit profiles (see § 2.4 and figure 1), while the crosses represent the pedestal top of the T_e predictions.

of agreement between the experimental and predicted T_e profiles, and, in particular, they represent a direct measure of profile goodness-of-fit for $T_e(\psi_{n_e, \text{Top}})$ because $\psi_{n_e, \text{Top}}$ is common for both profiles.

The ‘None’ group yielded largely random predictions of the T_e -pedestal’s width $w_{T_e, \text{Ped}}$ and top T_e^{Ped} , demonstrating insufficient predictive power without parameter input.

The ‘No-EFIT’ group performed significantly better, with marked improvements in the predictions of $T_e(\psi_{n_e, \text{Top}})$, T_e^{Ped} , and $\psi_{T_e, \text{Top}}$, suggesting that engineering parameters are highly relevant to the T_e profile.

The ‘EFIT’ group showed enhanced accuracy of the prediction of the pedestal width $w_{T_e, \text{Ped}}$, and comparable performance in predicting $T_e(\psi_{n_e, \text{Top}})$ and T_e^{Ped} as the ‘No-EFIT’ group. However, this group predicts the pedestal-top location $\psi_{T_e, \text{Top}}$ more poorly compared with the ‘No-EFIT’ group.

The ‘All Parameters’ group provided the most accurate predictions, achieving close alignment with experimental profiles across all our metrics for pedestal amplitudes and position. This is confirmed by the high r^2 values for the $T_e(\psi_{n_e, \text{Top}})$, T_e^{Ped} , and $\psi_{T_e, \text{Top}}$ reconstructions. The prediction of $w_{T_e, \text{Ped}}$ also improved, compared with the ‘None’ and ‘No-EFIT’ groups, underscoring the complementary roles of EFIT and engineering parameters in predicting T_e profiles.

Examples of these predictions are shown in figure 5 for pulse #90339. The outputs of the models using the ‘All Parameters’, ‘No-EFIT’, and ‘EFIT’ groups closely follow the experimental T_e and R/L_{T_e} profiles.

4.2. Local electron-temperature predictions

We further explore the T_e predictions based on local values of the density n_e and normalised density gradient R/L_{n_e} within a narrow radial window. This approach emulates predictions using physical models that use local parameters within the pedestal. Here, neural networks are trained to predict T_e values over narrow radial intervals, using the density information at these positions. The radial extent used at each location is $L = 2.5\%$ of the sampled radial interval (Appendix B.2). By choosing many randomly positioned samples, the full T_e profile is then assembled from local predictions.

Four parameter groups are defined under the heading ‘Local-value groups’ in table 1. The resulting `mtanh` parameters of these reconstructions are given in figure 6, arranged as follows:

- (a-d) ‘None’ – only the radial coordinate R , n_e , and R/L_{n_e} profiles, with no additional pulse parameters;
- (e-h) ‘No-EFIT’ – engineering parameters and R , n_e , and R/L_{n_e} profiles, excluding magnetic-equilibrium-reconstruction parameters;
- (i-l) ‘No- R ’ – all parameters, n_e , and R/L_{n_e} , but excluding the radial coordinate R ;
- (m-p) ‘All Parameters’ – all available engineering and EFIT parameters, and profiles including the radial coordinate R .

The ‘None’ group yields near-constant pedestals, reflecting the failure of local n_e values alone in capturing pedestal variation over our database. This is a case where insufficient information is provided, resulting in the network yielding an output with little variation. The parameters for the `mtanh` fits have large negative values, symptomatic of this strategy’s inability to reconstruct the pedestal shape using local predictions with no added parameters.

The ‘No-EFIT’ group enhances predictions of $T_e(\psi_{n_e, \text{Top}})$, T_e^{Ped} , and $\psi_{T_e, \text{Top}}$ with a strong correlation with the experimental values, although the predictions of the pedestal width remain inaccurate.

The ‘No- R ’ group demonstrates the critical role of the radial position R for precise pedestal predictions; without R , T_e^{Ped} is over-predicted by a factor of about 1.5, and the pedestal shapes are predicted poorly, indicating that R could act as a proxy for larger-scale spatial effects. This means that the local shape of the density profile alone is not sufficient to predict how tall the T_e profile must be in that location.

The ‘All Parameters’ group achieves high accuracy across all metrics, comparable with the results using the full-pedestal profile in §4.1. The pedestal width, however, is reproduced with a large inaccuracy, suggesting that $w_{T_e, \text{Ped}}$ is linked to global properties of the n_e profile. This link was also noticed by Saarelma *et al.* (2023, 2024), who found $w_{n_e, \text{Ped}}$ and n_e^{Ped} to be related to the T_e profile based on the penetration of neutrals and their ionisation through the pedestal region.

Examples of these predictions for #90339 are presented in figure 7, showing the effectiveness of local predictions with full parameter inclusion, and a high scatter in the temperature gradient of the reassembled profiles. Overall, the local profile predictions provide slightly poorer pedestal reconstructions than the global ones of §4.1, and these results are generally far more easily degraded by omitting information.

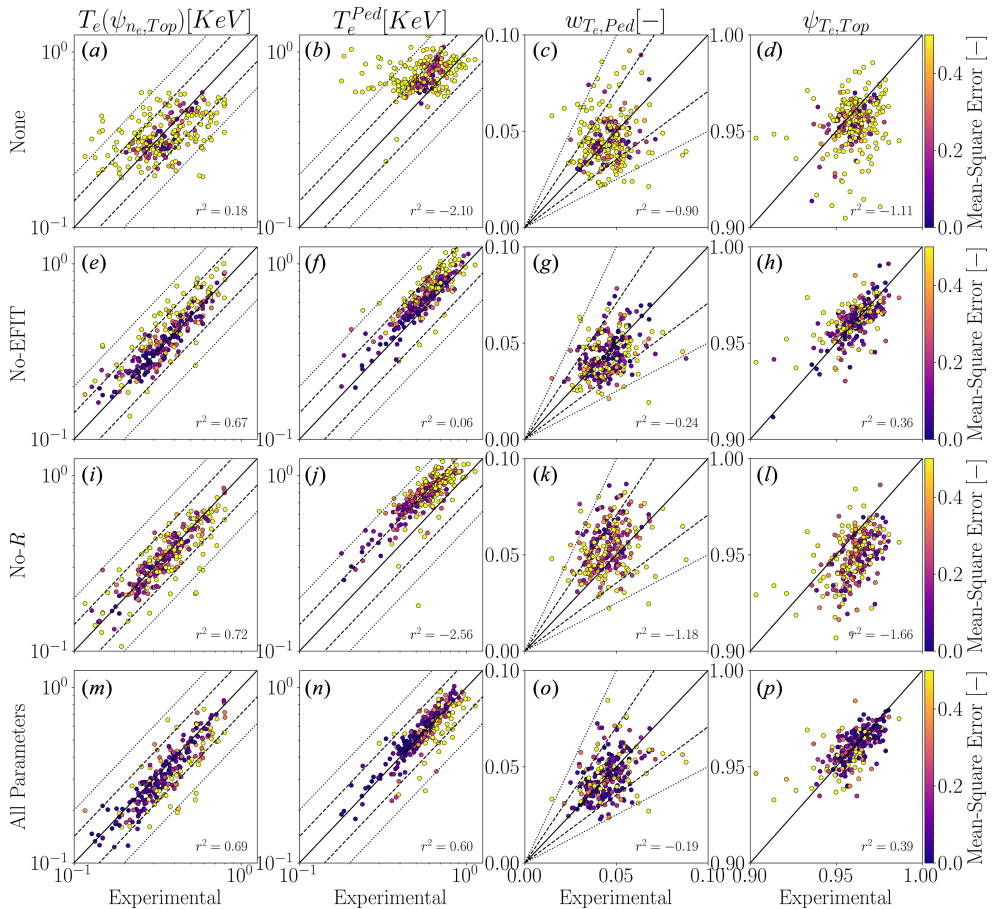


FIGURE 6. Database predictions using neural networks informed of the local values of n_e and R/L_{n_e} . These correspond to each of the four groups of parameters in table 1: (a–d) no database parameters used as inputs, with the exception of R and n_e ; (e–h) only engineering parameters alongside R and n_e ; (i–l) all available parameters and the n_e values, omitting the radial location R ; (m–p) all available parameters used, alongside R and n_e . The colour bar represents the mean-square error between the predicted T_e and the experimental fit profile. The plots follow the same layout as in figure 4.

4.3. Assessment of most important database parameters

For a quantitative assessment of parameter importance, we measure and compare the prediction accuracy of T_e profiles using various parameter combinations as inputs. We scan sequentially through combinations of N parameters from table 1 and identify those that correlate most with the success of the T_e -pedestal reconstruction.

Because of the difference in the sensitivity to parameter inclusion and in the performance of full-profile and local-value predictions, we scan parameters for each case. The results of these parameter scans are presented in figure 8, where we rank the performance of networks that use N parameters as inputs, with N indicated by different colours, as shown in the legends. The diagrams present the average performance for each parameter on the horizontal axis, when $N - 1$ other parameters are included. Our metric for the performance of each network is the total mean-square

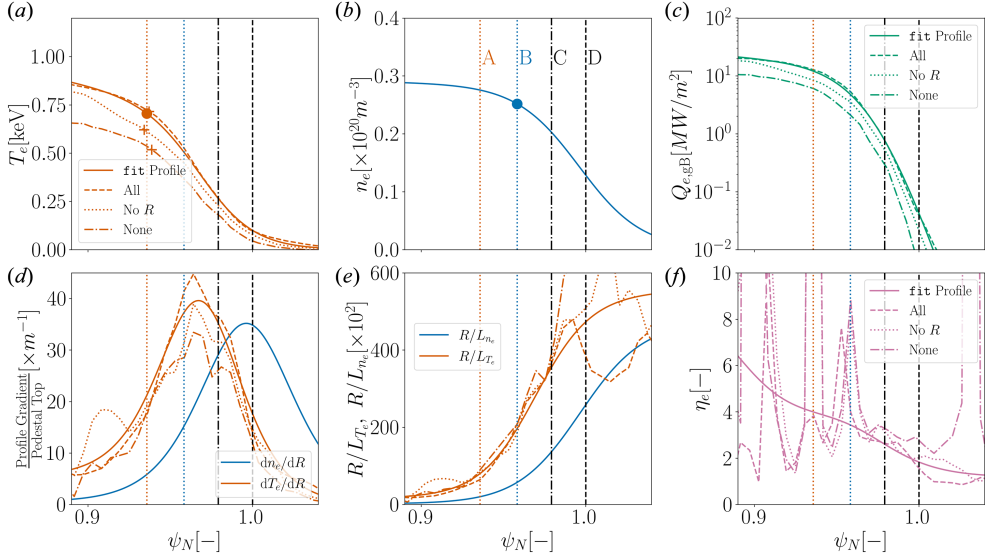


FIGURE 7. Pulse #90339 fit profiles, compared with neural-network local predictions of the T_e values. The layout is the same as in figures 1 and 5. The fit profiles are shown as solid lines, the predictions with different parameter groups listed in table 1 are shown as dash-dotted (‘None’), dashed (‘All Parameters’), and dotted (‘No- R ’) lines. The vertical lines and dots represent the landmarks of the fit profiles (see § 2.4 and figure 1), the crosses represent the pedestal top of the T_e predictions. The full predicted T_e profiles are a result of assembling many local predictions of T_e values at different locations.

error across all pulses, normalised to the error obtained from the respective ‘All Parameter’ groups from §§ 4.1 and 4.2. The error for any given network is defined as

$$\mathcal{E} = \sum_{\text{Pulses}} \sum_i [T_{\text{fit}}(R_i) - T_{\text{Pred}}(R_i)]^2, \quad (4.2)$$

where $\{R_i\}$ denote the evenly spaced radial locations where the profiles are sampled numerically, $T_{\text{fit}}(R_i)$ is the experimental fit profile, and $T_{\text{Pred}}(R_i)$ is the predicted value at each location. Bright colours represent the sum over the test pulses, pale colours represent the training pulses.

For full-profile predictions (figure 8a), for $N \geq 2$, the deuterium gas fuelling rate Γ_D becomes one of the highest-performance parameters in spite of its low importance in the $N = 1$ group. Overall, the inclusion of Γ_D , the separatrix loss power P_{Sep} , the plasma current I_{plasma} , and the strike-point configuration has proved essential for maximising the accuracy of neural networks with $N \geq 4$. Furthermore, the magnetic-equilibrium parameters have less impact than the engineering parameters, suggesting that there is redundancy in the system: the n_e profile contains sufficient information to infer some features of the MHD equilibrium.

The importance of the strike-point configuration and Γ_D in the full-profile models suggests that changing the magnitude and/or location of the neutral ionisation source (and also the efficiency of the pumping) and, hence, the particle flux through the pedestal, might affect the nature of the turbulence, and, consequently, the steady-state T_e profile that is consistent with this changed transport. One known pathway




FIGURE 8. Normalised mean-square error of neural-network predictions of (a) full T_e profiles and (b) local T_e profiles, averaged over combinations of N parameters that include each given parameter (horizontal axis). The normalisation is with respect to the mean-square error of the all-parameter networks, shown as the black dashed line at the value of 1 (the all-parameter networks are shown in figure 4*m-p* for full-profile predictions and in figure 6*m-p* for local-value predictions). The bright bars indicate the test-set errors, while the pale bars represent the training-set errors. The red bars show the error when only the single parameter marked on the horizontal axis is used. The pink bars display the average error when the parameter is used in two-parameter combinations, with further bars showing averages across combinations of more parameters. The parameters are arranged left to right in ascending order of the single-parameter test error, demonstrating the relative importance of each parameter for accurate full-profile predictions.

through which the information about T_e can be inferred lies in the ratio $n_e^{\text{Sep}}/n_e^{\text{Ped}}$, which is not contained in the n_e profiles: profiles with a higher $n_e^{\text{Sep}}/n_e^{\text{Ped}}$ have a higher mean η_e (figure 13) and are, therefore, more turbulent (§ 8), causing lower T_e^{Ped} values; in turn, $n_e^{\text{Sep}}/n_e^{\text{Ped}}$ is strongly related to Γ_D (Frassinetti *et al.* 2021) and to the strike-point configuration.

For local predictions (figure 8*b*), the radial coordinate R has proved to be the most important parameter across all metrics used, consistent with the findings of § 4.2. The inclusion of R and the magnitude of the magnetic field B on the axis has turned out to be the most effective two-parameter choice, however, for $N \geq 3$, P_{Sep} , Γ_D , I_{plasma} , and the strike-point configuration have proved to be the best-performing set. The importance of the strike-point configuration and Γ_D in this case is unsurprising, as they are essential for the formation of the n_e pedestal (Frassinetti *et al.* 2020; Saarelma *et al.* 2023, 2024), supplying information that would be contained in the full shape of the n_e .

Overall, these scans of parameter combinations suggest that the P_{Sep} , I_{Plasma} , Γ_D , and the strike points are the parameters strongest correlated to the correct prediction of the T_e -pedestal shapes. The difference between the ‘All Parameter’ groups in the full-profile and local predictions suggests that successful predictions rely on both global and local plasma characteristics. These global characteristics are contained in the full geometry of the n_e profile, and omitting this information degrades the quality of the predictions.

To conclude, the ML approach has proved quite successful in reconstructing the temperature pedestal given the density profile alongside other available parameters. Each of the eight neural networks presented in §§ 4.1 and 4.2 correspond to entries 39–46 in table 3, where their performances are compared with the other results in this study, stressing the success of such methods. Besides the practical importance of this success for modelling endeavours aiming to help design and operate future devices, the success of our local methods of pedestal reconstruction in § 4.2 also suggests that the search for a physics-based model of local transport is not a hopeless enterprise.

5. Temperature and density gradients

In this section, we will describe the qualitative behaviour of the temperature and density gradients across our database. These are expected to be the essential parameters of the local plasma equilibrium: local gradients represent a source of free energy, which causes microinstabilities. In the local picture, these microinstabilities saturate into a turbulent state that, in turn, determines the transport of heat (and particles) at each location. The object of theory and modelling is the relationship between the resulting transport coefficients, the underlying equilibrium gradients, and the consequent equilibrium profiles in a plasma carrying a known heat flux that is determined on the device scale.

Figure 9 shows the normalised gradients in the pedestal region for all pulses, separated based on their position relative to $\psi_{T_e, \text{Top}}$, $\psi_{n_e, \text{Top}}$, and ψ_{Sep} . This separation gives rise to a threshold that is effectively described by $R/L_{n_e} > \mathcal{O}(20)$ in the steep-gradient region; concomitantly, $R/L_{T_e} > \mathcal{O}(100)$ for the majority of pulses in the steep-gradient region. For both the density-top and the steep-gradient regions, it is apparent that the pulses reside in a regime roughly characterised by $L_{n_e}/L_{T_e} > 1$. In the density-top region, this difference in magnitude between R/L_{T_e} and R/L_{n_e} is consistent with the temperature pedestal’s inwards shift (Frassinetti *et al.* 2020), and is a contrivance of the split at $\psi_{n_e, \text{Top}}$. In contrast, in the steep-gradient region, it is believed to represent an important physical feature originating from the instability of electron-scale modes, as we will see below.

5.1. Instabilities and ETG turbulence

Besides the density and temperature gradients, which can trigger instabilities, the pedestal region features a large $\mathbf{E} \times \mathbf{B}$ flow shear. This has a suppressive effect on microinstabilities (Zhang & Mahajan 1992; Hahm & Burrell 1995; Terry 2000), quenching those of them that have growth rates smaller than the shear rate. Amongst the instabilities that are fast enough to overcome this suppression, the most relevant ones are modes with electron-scale radial extent, which cause heat transport in the electron channel (Told *et al.* 2008; Jenko *et al.* 2009; Hatch *et al.* 2015, 2016, 2017; Kotschenreuther *et al.* 2017, 2019). Namely, these are MTMs (see Drake & Lee 1977; Drake *et al.* 1980; Applegate *et al.* 2007; Guttenfelder *et al.* 2012;




FIGURE 9. Loci of the data from all pulses in terms of $(R/L_{n_e}, R/L_{T_e})$ coordinates: panels (a,c) contain the subsets of each profile spanning the steep-gradient region; panels (b,d) contain profiles between the T_e -pedestal top and the n_e -pedestal top; panels (a,b) contain all pulses considered by us, with their position in ψ_N used to colour code the points; the colour in panels (c,d) shows the number of pulses where profiles cross each bin in the $(R/L_{n_e}, R/L_{T_e})$ space. All data points lie above the slab-ETG stability threshold of $R/L_{T_e} = 0.8R/L_{n_e}$ (Jenko, Dorland & Hammett 2001), depicted as the lower blue dash-dotted line in (a,b). The upper blue line represents $R/L_{T_e} = 1.2R/L_{n_e}$, as explained in § 5.1.

Zocco *et al.* 2015; Larakers *et al.* 2020, 2021; Giacomini *et al.* 2023), which are electromagnetic instabilities relying on finite resistivity, and ETG modes (see Horton, Hong & Tang 1988; Farengo, Guzdar & Lee 1989; Jenko *et al.* 2000; Adkins *et al.* 2022, 2023), which are electrostatic modes. Both of these are driven by the equilibrium electron-temperature gradient.

While gyrokinetic simulations of pedestal-regime plasmas recover various contributions to heat transport from MTMs, ETG modes, or other instabilities (e.g., Guttenfelder *et al.* 2021; Chapman-Oplopoiou *et al.* 2022; Leppin *et al.* 2023), the consensus appears to be that ETG modes are often responsible for an experimentally significant proportion of the turbulent heat transport in the pedestal (see Ren *et al.* 2024 and references therein). A recent study of two JET pulses (one JET-C and

one JET-ILW) confirming the dominance of ETG modes was given by Hatch *et al.* (2019). Therefore, it makes sense to discuss the implications of ETG turbulence for the pulses in our database.

The ETG instability has two branches, one mediated by the magnetic-field curvature, called toroidal ETG, and one independent of the curvature and mediated by fluctuating parallel electron flow (current), called slab ETG. The slab-ETG instability is driven by the electron-temperature gradient but stabilised by the electron-density gradient. A measure of their competition is the gradient ratio

$$\eta_e(R) = \frac{L_{n_e}}{L_{T_e}}, \quad (5.1)$$

already introduced in § 1 as a key (local) physical parameter.

Fluid models (Adkins *et al.* 2022) and linear GENE simulations (Jenko *et al.* 2001) all result in linear stability thresholds characterised by a near-unity $\eta_{e,\text{crit}}$ such that, for $\eta_e > \eta_{e,\text{crit}}$, slab-ETG modes are unstable. The numerical stability scans performed by Jenko *et al.* (2001) yielded

$$\left(\frac{R}{L_{T_e}}\right)_{\text{crit}} = \max \left\{ 0.8 \frac{R}{L_{n_e}}, (1 + \tau)(1 - 1.5\varepsilon) \left(1.33 + 1.91 \frac{\hat{s}}{q} \right) \left(1 + 0.3\varepsilon \frac{d\kappa}{d\varepsilon} \right) \right\}, \quad (5.2)$$

where $\tau = Z_{\text{eff}}T_e/T_i$ is the temperature ratio, $\varepsilon = a/R$ is the aspect ratio, q is the magnetic safety factor, \hat{s} is the magnetic shear, and κ is the plasma elongation. This formula contains the linear thresholds for the slab-ETG branch (left) and the toroidal-ETG branch (right). Typically, the second term in (5.2) has order-unity values and is therefore much smaller than $0.8R/L_{n_e}$.

The linear gyrokinetic stability analyses of pedestal plasmas by Jenko *et al.* (2009) and Guttenfelder *et al.* (2021) find a value of $\eta_{e,\text{crit}}$ in the range 0.9–1.2. This variation in the critical threshold is attributed to the effect of magnetic geometry and toroidal modes on the instability (Jenko *et al.* 2009). The region of $\eta_{e,\text{crit}} \in [0.8, 1.2]$ is marked in figure 9, bounded by the two blue dash-dotted lines. These order-unity values of $\eta_{e,\text{crit}}$ are situated well below the experimental values of the gradients; only near-separatrix values considerably overlap with this region, albeit they are subject to large uncertainties. This means that, throughout the pedestal, the slab-ETG modes are linearly unstable, to the best of existing theoretical knowledge.

This discussion prepares the analysis of the gradient ratio η_e and the outcomes of reconstructing T_e profiles based on the assumption that the pedestal gradients are ‘pinned’ to a specific value of η_e in § 6. This is based on the possibility that the system simply pins itself locally to a marginal value of η_e (although whether that marginality is with respect to the traditional linear stability threshold is to be discussed). In contrast to this approach, in § 8, we will present the results of reconstructing pedestals using scalings of turbulent heat transport: we will relate the separatrix loss power P_{Sep} to the background gradients of temperature and density, allowing η_e to vary.

6. Gradient ratio η_e and pedestal reconstruction

A simple way to use a physics-inspired approach for pedestal reconstructions is to make use of $\eta_e = L_{n_e}/L_{T_e}$. First, we will understand the disposition of this gradient ratio in relation to R/L_{T_e} , R/L_{n_e} , and the pedestal location X_{Ren} . Once we will understand the trends of the experimental values of η_e in § 6.1, we use our own model fits to reconstruct the temperature-pedestal profiles in § 6.2.




FIGURE 10. Distributions of the gradient ratio η_e , defined in (5.1), across the database conditioned on (a) R/L_{T_e} , (b) R/L_{n_e} , (c) $Q/Q_{e,GB}$. All data are obtained from the pedestal region, between $\psi_{T_e, \text{Top}}$ and ψ_{Sep} . Panel (b) suggests that $R/L_{n_e} \propto \eta_e^\zeta$ describes the pedestal data well, as will be discussed in § 7. The orange lines in (a) and (b) represent three fit values of ζ to the maxima of the R/L_{n_e} distribution (dashed), the full R/L_{n_e} distribution (dotted) and the steep-gradient-region prediction across the database discussed in § 7.2 (solid).

6.1. Distributions of the gradient ratio η_e

In figures 10(a) and 10(b), the data from figure 9 is recast to show an important property of the pedestals: the steep-gradient region ($R/L_{n_e} \gtrsim 20$) displays a strong tendency of η_e to maintain values of around 2 across this region. It is also apparent that there is a power-law relationship between η_e and R/L_{n_e} , which will be discussed in § 7. There is no such relationship between η_e and R/L_{T_e} evident in figure 10(a), as there is far more variance in R/L_{T_e} for any given η_e . Furthermore, the gyro-Bohm-normalised heat flux $Q/Q_{e,GB}$ appears to increase as η_e approaches the stability threshold (or a different $\mathcal{O}(1)$ value). This confirms the intuition of Field *et al.* (2023) regarding the role of the linear threshold, although with the notable difference that the heat flux must decrease further away from it.

In (2.8), we defined a radial coordinate X_{Ren} in order to exploit the geometry of each pulse and to match ψ_{Sep} , $\psi_{n_e, \text{Top}}$, and $\psi_{T_e, \text{Top}}$ across the database. In figure 11 we provide the explicit dependence of η_e on the radial location relative to these pedestal landmarks. An important pedestal property highlighted by figure 11 is that the most probable values of η_e range between 1.5 and 3.5 across the steep-gradient region, increasing from the separatrix (figure 11e) towards $\psi_{n_e, \text{Top}}$ (figure 11c). Since the stability threshold of ETG modes is expected to be at $\eta_e \approx 1$ (see § 5), this implies a measurable deviation of the turbulence from the linear threshold. Consequently, the turbulence resides in an order-unity supercritical regime where η_e is closer to 2 (which has also been reported in several ASDEX experiments, e.g., Neuhauser *et al.* 2002). The distribution of η_e does overlap with the linear threshold in the vicinity of the separatrix, although only just.

We note that η_e increases rapidly once the profiles cross $X_{\text{Ren}} = 1$, which is a consequence of the inward shift of the electron-temperature pedestal: the temperature pedestal maintains large gradients irrespective of the reduction in the density gradient. This could be indicative of a difference between the saturated turbulent states in the steep-gradient and density-top regions. This qualitative difference between the two regions can be interpreted to imply that, in the steep-gradient region, the density

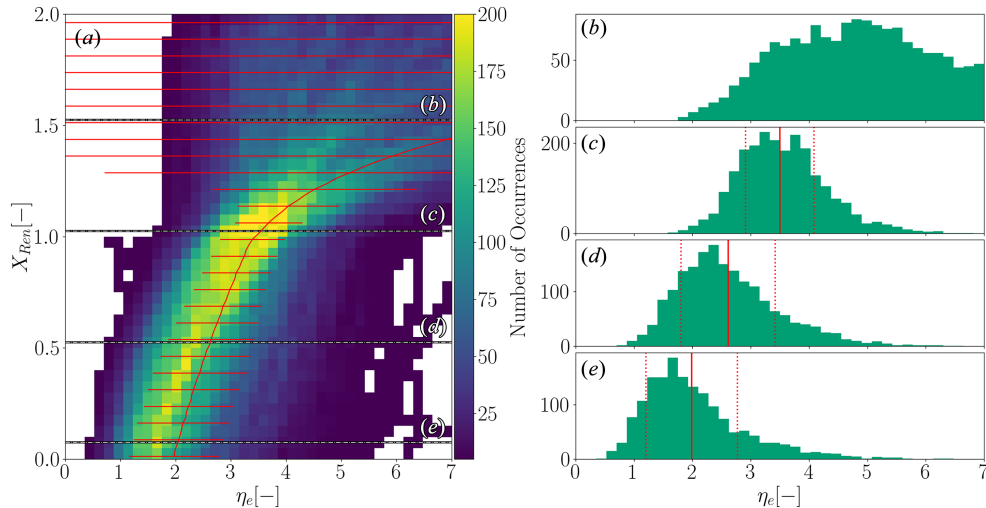


FIGURE 11. (a) Distribution of η_e (as defined in (5.1)) conditioned on the renormalised profile position X_{Ren} , defined in (2.8). The red lines represent the mean and variance of η_e at each X_{Ren} . (b–e) Cross-sections of the histogram at (b) $(\psi_{T_e, Top} + \psi_{n_e, Top})/2$, (c) $\psi_{n_e, Top}$, (d) ψ_{Steep} , and (e) ψ_{Sep} , corresponding to the white dash-dotted lines in (a). In each of these cross-sections, the continuous red lines represent the mean value of η_e , and the dotted red lines represent the values one standard deviation away from the mean. Note that these lines are outside of the range of η_e captured in (b).

gradient plays an essential role in characterising the turbulence, and the turbulent state resides in a critical regime where the temperature gradient is (almost) uniquely determined by the density gradient. In contrast, inside the density-pedestal top, the density gradient does not uniquely determine the temperature gradient, and R/L_{T_e} depends more strongly on other plasma parameters, as can be inferred from the comparison of figures 10(a) and 10(b). This difference could be attributed to the physically expected prevalence of slab-ETG modes in the steep-gradient region, with the density gradient being the dominant stabilising effect on these modes. At small R/L_{n_e} , toroidal-ETG modes become unstable before the slab-ETG ones (Jenko *et al.* 2001), and the addition of such subdominant modes can alter the heat transport (Parisi *et al.* 2022). This might be the relevant physics in the density-top region. The inclusion of electromagnetic effects into this picture complicates matters further, as we briefly discussed in § 5.1.

6.2. Electron-temperature reconstruction using constant gradient ratio

Using the tendency of η_e to have order-unity values in the pedestal region, we will proceed to predict the electron-temperature profiles assuming there is a simple functional relationship between the normalised gradients R/L_{T_e} and R/L_{n_e} . The simplest such relationship, proposed by Field *et al.* (2020), would be a set value of η_e describing the turbulence across the whole profile. This is, in itself, an assumption of infinitely ‘stiff’ transport limited by a critical η_e , maintained across the pedestal: the heat flux diverges as soon as this value of η_e is crossed, reducing the local temperature gradient and consequently bringing the heat flux back to a finite value. Based on the results of § 6.1, we expect this imposed value of η_e to be in the vicinity of 2. We note that the experimental value of $\eta_e \approx 2$ across the pedestal was

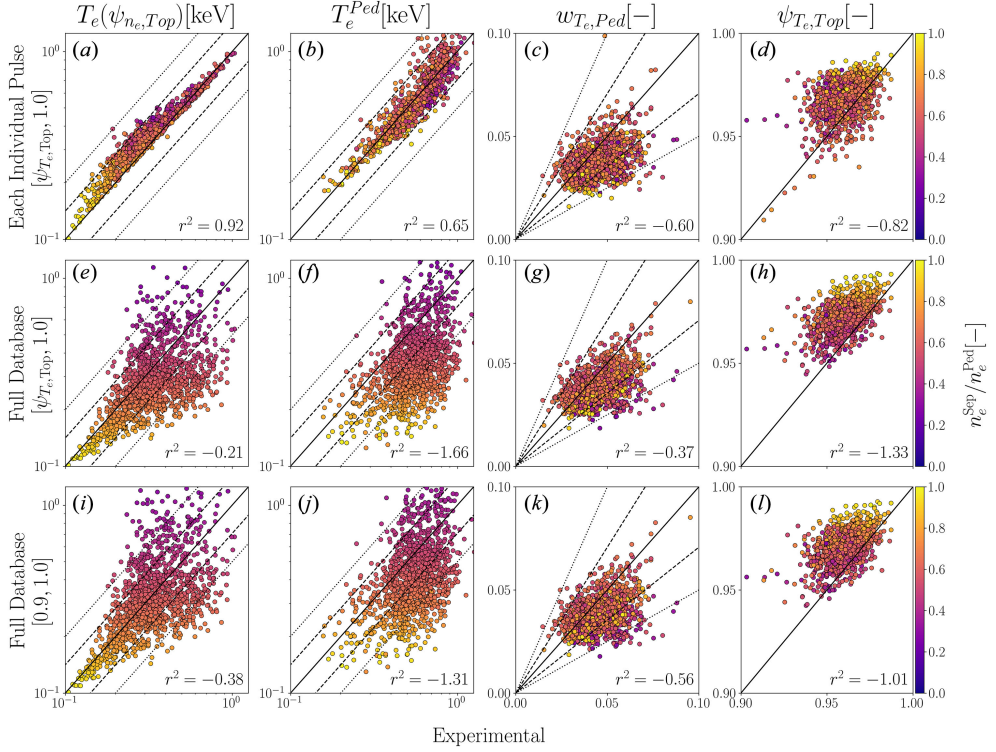


FIGURE 12. Database predictions obtained by integrating $1/L_{T_e} = \eta_e/L_{n_e}$ inwards from the separatrix, (a–d) with the best-fit η_e for each pulse over the pedestal region, (e–h) with the best-fit $\eta_e = 1.90$ across the database over the pedestal region, and (i–l) with the best-fit $\eta_e = 2.19$ across the database over the region $\psi_N \in [0.9, 1.0]$. The colour bars represent the ratio between the separatrix density n_e^{Sep} and the density-pedestal height n_e^{Ped} . The plots follow the same layout as in figure 4.

initially measured by Neuhauser *et al.* (2002) on ASDEX, but has been observed to characterise H-mode pedestals correctly on other devices as well (Ren *et al.* 2024). Horton *et al.* (2005) used the assumption of $\eta_e = 2$ to reproduce correctly an electron-temperature-pedestal profile on ASDEX Upgrade.

In order to find the values of η_e that best describe the pulses in our database, we reconstruct the electron-temperature profiles by first assuming that there is a constant η_e in (5.1), and then integrating inwards from the boundary at R^{Sep} (as detailed in Appendix C). We then identify the best ‘nominal’ η_e such that the difference between the fit and the reconstructed pedestal is minimised (as explained in Appendix C.1).

In figure 12, we present three sets of results obtained using this method of integration and parameter optimisation over different domains. These three database reconstructions are calculated by maximising the accuracy of the T_e profile across:

- (a–d) the pedestal region, $\psi_N \in [\psi_{T_e, \text{Top}}, 1.0]$, with a different η_e for each pulse;
- (e–h) the pedestal region, $\psi_N \in [\psi_{T_e, \text{Top}}, 1.0]$, with the best-fit $\eta_e = 1.90$ across the database;
- (i–l) the region including a part of the core, $\psi_N \in [0.9, 1.0]$, with the best-fit $\eta_e = 2.19$ across the database.




FIGURE 13. Values of η_e that best predict the T_e profiles across the pedestal region. (a) The distribution of values of fit η_e . (b) Comparison between the fit values of η_e and the database values of η_e at ψ_{Steep} (as shown in figure 11d).

From the pedestal parameter comparison in figure 12, it is apparent that reconstructed temperature-pedestal profiles have temperature-top locations pushed further outwards than the experimental location. The inflexibility of mandating a constant value of η_e yields predicted values of $\psi_{T_e, \text{Top}}$ that follow $\psi_{n_e, \text{Top}}$, irrespective of whether η_e is fixed across the database or allowed a different value for each pulse. Furthermore, for the database-wide η_e predictions in figure 12(e–h), we find that the temperature-pedestal height, width, and top locations are strongly correlated to the ratio between the separatrix density n_e^{Sep} and the pedestal density n_e^{Ped} , supporting this observation. Overall, while it is reassuring that we recover $\eta_e \approx 2$ in our database, setting a constant value of $\eta_e \approx 2$ produces poor predictions of the pedestal-top regions. Furthermore, the negative values of r^2 in the case of $T_e(\psi_{n_e, \text{Top}})$ indicate that a fixed value of η_e is unable to reproduce the steep-gradient region, as it gives a poorer reconstruction than asserting that the predicted T_e pedestal is the average of the pedestals across our database.

The performance of the models using database-wide η_e are presented in entries 31–34 of table 3. In particular, those presented in figures 12(e–h) and 12(i–l) correspond to entries 32 and 34, respectively.

Thus, allowing the value of η_e to vary across the database, i.e., assuming that the turbulence in each individual pedestal can be described by a single parameter that is constant across the profile, yields reasonably good predictions for the steep-gradient region, while over-predicting T_e at $\psi_{n_e, \text{Top}}$ and under-predicting T_e^{Ped} within a factor of $\sqrt{2}$ of the experimental value. The distribution of the best-fit η_e across the steep-gradient region is shown in figure 13. We find a considerable overlap between the best-fit values of η_e from the pedestal predictions of figure 12(a–d) and the experimental values of $\eta_e(\psi_{\text{Steep}})$ at the steep-gradient point (shown in figure 11d). The largest discrepancies between the experimental $\eta_e(\psi_{\text{Steep}})$ and the fit η_e are correlated with the ratio $n_e^{\text{Sep}}/n_e^{\text{Ped}}$: outwards shifted n_e pedestals require a higher optimal η_e than the experimental η_e in order for the T_e pedestal to rise enough. In general, pedestals with a significant shift have a poor prediction accuracy because of the inability of the constant- η_e models to reproduce large R/L_{T_e} beyond the density-top point.

Applying symbolic regression models (Cranmer 2023) to estimate the optimised η_e for each pulse as a function of database parameters shows that the engineering and magnetic-field parameters lack any correlation with the obtained η_e . The inclusion of $n_e^{\text{Sep}}/n_e^{\text{Ped}}$ alongside these parameters is essential for the correct estimation of η_e , but the scalings fit to the optimised η_e are extremely complex and unlikely to have physical meaning. The importance of $n_e^{\text{Sep}}/n_e^{\text{Ped}}$ could not have been elicited in § 4 as the location of the separatrix (and, hence, the value of n_e^{Sep}) was omitted from the training information of the neural networks. This importance of $n_e^{\text{Sep}}/n_e^{\text{Ped}}$ must complement the relationship this ratio bears with the pedestal pressure (Frassinetti *et al.* 2021).

6.3. Physical implications of η_e results

The distribution of experimental values of η_e (figure 10) over the steep-gradient region and their agreement with the best-fit η_e values over the pedestal region in figure 13(b) represent an important finding. These best-fit values of η_e are nominally higher than 2, with a considerable spread.

One can interpret the results of this section as providing physical insight into the steep-gradient-region turbulence in several alternative ways.

- (i) The turbulence can be described by a quasilinear model strongly dominated by the linear picture of electrostatic and electromagnetic modes, but the true linear threshold $\eta_{e,\text{crit}}$ differs from the solution of Jenko *et al.* (2001) by an order-unity value.
- (ii) The turbulence saturates in a nonlinear ‘Dimits’-like state (Dimits *et al.* 2000; Colyer *et al.* 2017; Ivanov, Schekochihin & Dorland 2022), and the transport level increases abruptly once some nonlinear threshold $\eta_{e,\text{Dimits}}$ is surpassed; based on our results, we expect that $\eta_{e,\text{Dimits}}$ would be higher than the linear $\eta_{e,\text{crit}}$ by an $\mathcal{O}(1)$ amount.
- (iii) The turbulence causes a moderately (rather than near-infinitely) ‘stiff’ heat transport (Field *et al.* 2023) and, given the separatrix loss power P_{Sep} in JET-ILW pulses, it results in a value of η_e that is greater than $\eta_{e,\text{crit}}$ by an $\mathcal{O}(1)$ amount.

In § 6.1, we also identified and discussed the stark difference between the steep-gradient and density-top regions. Explaining this difference demands a theoretical model that accounts for qualitative changes in the turbulent behaviour with the change of the density gradient R/L_{n_e} . Such changes could be the result of the contribution of toroidal-ETG modes to the turbulent transport, as recently found by Chapman-Oplopoiou *et al.* (2025) and Krutkin *et al.* (2025). We will discuss gyrokinetics-inspired models (Guttenfelder *et al.* 2021; Chapman-Oplopoiou *et al.* 2022; Hatch *et al.* 2022) that use the turbulent heat flux to reconstruct the T_e profiles in § 8, but before that, in § 7, we will construct and assess a simpler scaling prescription that allows for η_e to adjust over the pedestal region.

7. Scaling exponent α and pedestal reconstruction

The relationship between η_e and R/L_{n_e} shown by the dashed orange line in figure 10(b) suggests that a power-law scaling between the gradients can effectively describe the pedestal. Furthermore, a clear characteristic of the plasma gradients in figure 9 is that R/L_{T_e} monotonically increases as R/L_{n_e} does. This increase can




FIGURE 14. Distribution of the exponent α , defined in (7.2), across the database as a function of (a) R/L_{T_e} , (b) R/L_{n_e} , (c) η_e . All data are obtained from the pedestal region, between $\psi_{T_e, \text{Top}}$ and ψ_{Sep} .

perhaps be described by a power-law relationship of the form

$$\frac{R}{L_{T_e}} = A \left(\frac{R}{L_{n_e}} \right)^\alpha, \quad (7.1)$$

where α and A are parameters to be measured (and theorised about). To analyse the behaviour of the exponent α at different locations in the pedestal, we formally define

$$\alpha(R) = \frac{d \ln (R/L_{T_e})}{d \ln (R/L_{n_e})}, \quad (7.2)$$

where R is the radial location at which the gradients are evaluated. Following the R/L_{T_e} vs R/L_{n_e} loci for any given pulse in figure 9, we find that $\alpha(R)$ decreases from $\psi_{T_e, \text{Top}}$ (low-gradient region of figure 9a) to ψ_{Sep} (dark-red lines in figure 9b).

In § 7.1, this local $\alpha(R)$ will be calculated for each pulse in the database. Then, in § 7.2, overall, the range of value, we will fit α -based models to the database and reconstruct the electron-temperature pedestals, similar to our analysis of η_e in § 6.

7.1. Distributions of the scaling exponent α

In figure 14, we show the distribution of the exponent α conditioned on the gradients R/L_{T_e} , R/L_{n_e} , and their ratio η_e . The values of α at the separatrix are clustered around ≈ 0.33 (as seen in figure 14b at large R/L_{n_e}). Overall, the range of values taken by α is narrow, indicating that a power-law relationship between R/L_{T_e} and R/L_{n_e} with a constant exponent might indeed represent the profiles in the steep-gradient region correctly. The relationship between α and η_e is shown in figure 14(c), and it implies a clear correlation between increasing η_e and increasing α .

Rearranging the values of α using the renormalised radial coordinate X_{Ren} given by (2.8) gives us figure 15. Its most important feature is the convergence of all profiles to $\alpha \approx 0.33$ at the separatrix (X_{Ren} approaching 0 in figure 15a). This coincides with the highest values of the normalised gradients R/L_{T_e} (figure 14a) and R/L_{n_e} (figure 14b). The exponent α commonly lies in the range of 0.3–0.5 in the steep-gradient region, with a large spread towards $\psi_{n_e, \text{Top}}$. The values of α remain order unity in the density-top region.




FIGURE 15. (a) Distribution of α , as defined in (7.2), conditioned on X_{Ren} . The red lines represent the mean and variance of α at each X_{Ren} . (b–e) Cross-sections of the histogram at (b) $(\psi_{T_e, Top} + \psi_{n_e, Top})/2$, (c) $\psi_{n_e, Top}$, (d) ψ_{Step} , and (e) ψ_{Sep} , corresponding to white dash-dotted lines in (a). In each of these cross-sections, the continuous red lines represent the mean value of α , and the dotted red lines represent the values one standard deviation away from the mean.

7.2. Electron-temperature reconstructions using constant exponent α

Because of the small variance of α in the pedestal region, it makes sense to explore the consequences of asserting that α is constant across the pedestal. This is conceptually similar to the analysis carried out in § 6.2, where $\alpha = 1$ and $A \equiv \eta_e$, but now $\alpha \neq 1$ allows the physically relevant parameter η_e to adjust over the pedestal (as a function of either of the gradients). The case for such an adjustment of η_e across the pedestal is clear in figure 10(a,b), where the orange lines represent fits of α to the distributions of η_e vs R/L_{n_e} – in the figure caption, these fits are presented as $R/L_{n_e} \propto \eta_e^\zeta$, where $\zeta = (\alpha - 1)^{-1}$. The dotted line with $\alpha \approx 0.1$ is the fit to the whole distribution of η_e in figure 10(b), while the dashed line with $\alpha = 0.41$ (obstructed by the continuous orange line with $\alpha = 0.39$ obtained below as a best fit for database reconstructions) is the fit to the highest-likelihood η_e values in the histogram. It is clear that $\alpha \approx 0.1$ fails to reproduce the trend of pedestal profiles, whereas $\alpha = 0.41$ meaningfully matches the relationship between η_e and R/L_{n_e} . Hence, using a constant α to describe a pedestal should lend enough flexibility to the relationship between R/L_{T_e} and R/L_{n_e} .

Thus, similar to our approach in § 6.2, we attempt to reproduce the pedestals using (7.1) as a fitting formula, seeking the best values of A and α . Figure 16 shows two database reconstructions that were calculated using the scaling (7.1). These reconstructions maximise the accuracy of the T_e profile across the pedestal region, $\psi_N \in [\psi_{T_e, Top}, \psi_{Sep} = 1.0]$, with:

(a–d) A and α fit individually for each pulse;

(e–h) $A = 50.6$ and $\alpha = 0.39$ fit across the database.

For these reconstructions in figure 16, the values of A and α are the best-fit values obtained using the optimisation method described in Appendix C.1. Note that the

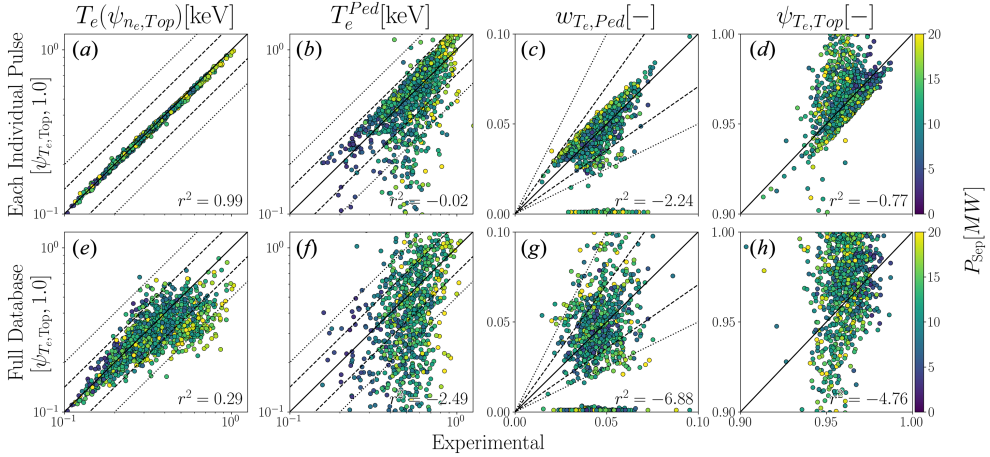


FIGURE 16. Database predictions obtained by integrating $R/L_{T_e} = A(R/L_{n_e})^\alpha$ inwards from the separatrix, with (a–d) the best-fit A and α for each pulse, and (e–h) the best-fit $A = 50.6$ and $\alpha = 0.39$ across the database. The fits are done over the pedestal region. The colour bar represents the separatrix loss power P_{Sep} . The layout of the figure is the same as in figure 4.

value of $\alpha = 0.39$ in this method agrees with the fit to the highest-likelihood values of η_e of figure 10(b).

The predictions obtained by allowing α and A to be fit individually for each pulse (figure 16a–d) are very accurate in reproducing the pedestal region overall, as seen in figure 16(a). Furthermore, the pedestal shapes are reconstructed well for a large proportion of the database, although for some pulses, the T_e profiles continue increasing exponentially towards the core, so mtanh parameters do not describe these profiles well. This can be seen in figure 16(b–d), where a subset of the mtanh parameters is very far from the correct prediction.

This issue of unbounded T_e profiles is more pronounced for the reproductions using a single combination of $A = 50.6$ and $\alpha = 0.39$ across the database. Since many reconstructed profiles do not resemble an mtanh , the corresponding pedestal parameters are obtained poorly (figures 16f–h). However, the accuracy of these fits across the steep-gradient region (figure 16e) is much better than the best-fitting η_e predictions of § 6.2 (figure 12e,i), meaning that such an approach is more effective for the steep-gradient region than the methods in § 6. This combination of A and α is shown in figure 10(a,b) as the continuous orange lines, recovering the fit performed directly to the distribution of R/L_{n_e} and η_e shown by the dashed line.

If A and α are instead optimised over the interval $\psi_N \in [0.9, 1.0]$ for a database-wide fit, we obtain $A = 31.7$ and $\alpha = 0.46$ (not shown). This combination produces slightly degraded pedestals over the $\psi_N \in [\psi_{T_e, \text{Top}}, 1.0]$ interval when compared with the combination in figure 16(e–h), and it does not solve the general issue of the exponential divergence of T_e profiles. The combinations of A and α optimised over the intervals $[\psi_{T_e, \text{Top}}, 1.0]$ and $[0.9, 1.0]$ correspond to entries 36 and 38 of table 3 in Appendix A, where their accuracies over different ranges are compared with our other attempts at T_e -pedestal reconstructions in this study.

Examples of the profiles that constitute figure 16 are given in figure 17 for the pulse #90339. The database-optimised fit (dotted lines) has a lower value of α than the optimal value for pulse #90339, causing R/L_{T_e} to vary less with R/L_{n_e} . Because

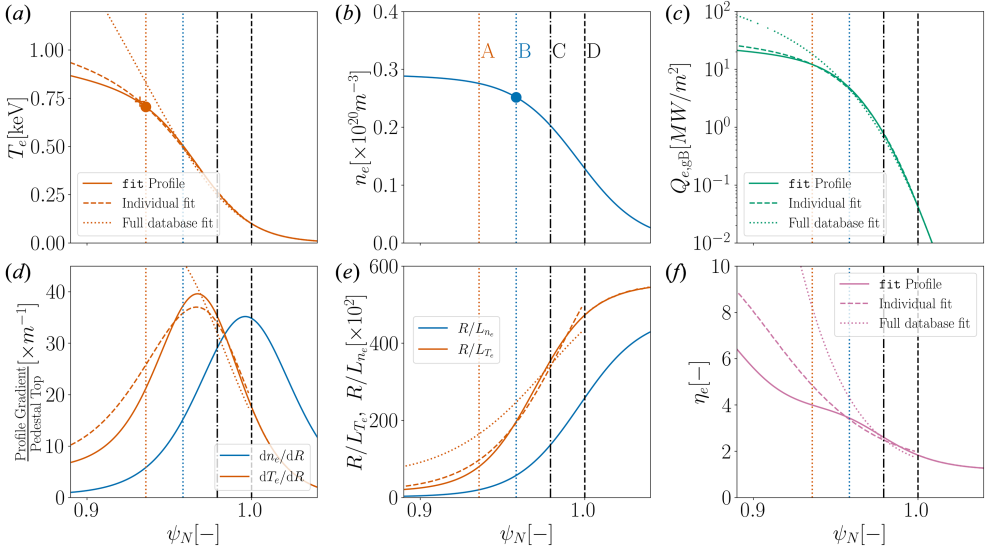


FIGURE 17. Pulse #90339 fit profiles compared with predictions using the scaling (7.1). The layout is the same as in figure 1. The fit profiles are shown as solid lines, while the predictions using the A and α optimised individually for pulse #90339 are shown as dashed lines, and predictions using the A and α optimised for the whole database are shown as dotted lines. Both fits are done over the interval $[\psi_{T_e, \text{Top}}, \psi_{\text{Sep}}]$. The vertical lines and dots represent the landmarks of the fit profiles (see § 2.4 and figure 1), while the crosses represent the pedestal top of the T_e predictions.

of this, the database-optimised values give a large R/L_{T_e} beyond the density-pedestal top and a diverging T_e profile without a $\psi_{T_e, \text{Top}}$ within the bounds shown in figure 17. The optimal combination of A and α for this individual pulse (dashed lines) results in a good T_e pedestal and reproduces η_e correctly between $\psi_{T_e, \text{Top}}$ and ψ_{Sep} .

The distributions of A and α fit for each pedestal (corresponding to figure 16a–d) are shown in figure 18. The key finding, manifest in figure 18(b), is that, while we found that a two-parameter fit (7.1) is much better than the one-parameter fit ($\alpha = 1$, $A = \eta_e$ in § 6), the two parameters appear to be highly correlated. Their relationship is given by $\ln A = 5.64 \pm 0.05 - (4.42 \pm 0.08)\alpha$, which allows for (7.1) to be recast as

$$\frac{R}{L_{T_e}} = (281 \pm 15) \left(\frac{1}{84 \pm 7} \frac{R}{L_{n_e}} \right)^\alpha. \quad (7.3)$$

The ranges for the numerical constants represent a 95 % confidence interval. This means that the fits performed in this section in fact intrinsically have only one degree of freedom. We have found that a similar locus of fit parameters A and α , given by $\ln A = 5.73 \pm 0.09 - (5.18 \pm 0.12)\alpha$, describes a subset of 273 JET-C pulses from the EUROfusion pedestal database, selected based on criteria similar to those described in § 3.1. These JET-C pulses are not otherwise included in this study.

Finding a relationship between α and physical parameters would represent a major success in modelling the pedestal steep-gradient region. However, we anticipate that a model like (7.3) might be much harder to justify physically than other alternatives




FIGURE 18. Values of α and A that best predict the T_e profiles individually across the pedestal region: (a) the distribution of values of fit α , (b) the fit values of α vs A . The black cross represents the A and α fit for the whole database (shown in figure 16e–h). The linear fit between $\ln A$ and α is shown as a blue line.

used in §§ 6 and 8. We also note that we have failed to find any correlation between the distribution of α and database parameters.

8. Heat-flux models

As we saw throughout the parameter-combination scans of § 4.3, the separatrix loss power P_{Sep} is strongly correlated with the T_e pedestals. As also highlighted by the correlation matrix in figure 22, the values of the T_e -pedestal profile at the pedestal-top locations T_e^{Ped} and $T_e(\psi_{n_e, \text{Top}})$ are positively correlated with P_{Sep} . We will now attempt to exploit this dependence of T_e on P_{Sep} using models of heat transport to reconstruct T_e pedestals.

This section will show the outcomes of T_e -pedestal-profile reconstruction using models that relate the gradients R/L_{n_e} and R/L_{T_e} to the turbulent heat flux transported in the electron channel. First, in § 8.1, we will build upon the discussion of linear instabilities in the pedestal given in § 5.1 and present the properties of the turbulent transport caused by ETG modes and other instabilities. This will provide context to the rest of our analysis. Section 8.2 will describe the calculation of the turbulent heat flux and its normalisation using database parameters. Using the models explained in § 8.1, in § 8.3, we will reconstruct electron-temperature pedestals and assess the quality of these predictions. Our reconstructions will use the parameter-optimisation methods described in Appendix C, and the results in this section will represent an illustrative subset of the full set of models we assessed in Appendix D.

8.1. Electron-scale turbulence

The nonlinear dynamics of ETG-driven turbulence in the steep-gradient region of pedestals are strongly influenced by the complex interactions between microinstabilities that emerge as a result of the underlying plasma equilibria. Here, we briefly review the key findings from nonlinear gyrokinetic simulations and their implications for pedestal turbulence.

Chapman-Oplopoiou *et al.* (2022) carried out gyrokinetic simulations using the nominal gradients R/L_{T_e} and R/L_{n_e} evaluated at ψ_{Steep} for four JET-ILW pedestals, which are also present in this study. Varying these gradients within the experimental confidence intervals showed that the electron-channel turbulent heat flux $Q_{e,\text{Turb}}$ increased with increasing R/L_{T_e} and decreased with increasing R/L_{n_e} . This is consistent with the linear picture of the slab-ETG instability, where electron-temperature gradients drive the instability, while electron-density gradients suppress it. Chapman-Oplopoiou *et al.* (2022) showed that a good fit to their results was the relationship

$$\frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = A \left(\frac{R}{L_{T_e}} \right)^2 (\eta_e - \eta_{e,\text{crit}})^\beta, \quad (8.1)$$

where, as defined in § 2.3, $Q_{e,\text{gB}} = n_e T_e v_{\text{th},e} \rho_e^2 / R^2$, and the best-fit parameters were $A = 0.85$, $\eta_{e,\text{crit}} = 1.28$, and $\beta = 1.32$. The simulated heat fluxes were also consistent with the fit found in simulations of DIII-D pedestals by Guttenfelder *et al.* (2021): $A = 1.5$, $\eta_{e,\text{crit}} = 1.4$, and $\beta = 1$. Models of this form can also reproduce results such as those from the R/L_{T_e} scans in Parisi *et al.* (2022): keeping η_e constant yields $Q_{e,\text{Turb}}/Q_{e,\text{gB}} \propto (R/L_{T_e})^2$, whereas keeping R/L_{n_e} constant gives $Q_{e,\text{Turb}}/Q_{e,\text{gB}} \propto (R/L_{T_e})^4$. This behaviour would be matched by (8.1) with $\beta = 2$ when $\eta_e \gg \eta_{e,\text{crit}}$.

The behaviour of the turbulent heat flux $Q_{e,\text{Turb}}$ parametrised by (8.1) can be justified by the physical picture proposed by Field *et al.* (2023). In their model, the $(R/L_{T_e})^2$ term is justified by replacing the radial size R of the tokamak with the locally relevant length scale L_{T_e} in the gyro-Bohm normalisation of the heat flux:

$$Q_{e,\text{gB}} = n_e T_e v_{\text{th},e} \frac{\rho_e^2}{R^2} \rightarrow Q_{e,\text{gB}}^* = n_e T_e v_{\text{th},e} \frac{\rho_e^2}{L_{T_e}^2}, \quad (8.2)$$

i.e., the relevant scale associated with free-energy injection is L_{T_e} , rather than R . This is a natural choice if turbulent transport is controlled by the slab-ETG instability. With the redefinition (8.2), (8.1) becomes

$$\frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}^*} = A (\eta_e - \eta_{e,\text{crit}})^\beta, \quad (8.3)$$

and, therefore, for $\beta > 0$, the normalised turbulent heat flux is a function of how much the slab-ETG drive (expressed through η_e) exceeds the (nonlinear) threshold $\eta_{e,\text{crit}}$. This is an essential characteristic that sets this model of turbulent transport apart from the models using a fixed η_e (§ 6) or a power law for the relationship between R/L_{T_e} and R/L_{n_e} (§ 7). Those models amounted to assuming that the turbulence was marginal and, therefore, ‘pinned’ to a specific value of $\eta_{e,\text{crit}}$, constant across the pedestal or varying in a prescribed manner. Model (8.1) allows the profiles to reach beyond this marginal regime.

This type of modelling of turbulent transport is common in the context of ion-temperature-gradient mode (ITG) turbulence in core plasmas and has been reported to reproduce measurements correctly (e.g., Garbet *et al.* 2004; Mantica *et al.* 2009). In the case of ITG turbulence, the critical value of the ion-temperature gradient is usually expected to be larger than the linear stability threshold because the nonlinear saturation just above the linear threshold is in the ‘Dimits’ regime, in which strong, nonlinearly saturated zonal flows suppress heat transport (Dimits *et al.* 2000;

Ivanov *et al.* 2022). A similar behaviour is sometimes found in simulations of ETG turbulence, where the long-term onset of zonal flows suppresses the heat transport caused by ETG (Colyer *et al.* 2017). Thus, it is conceivable (although, at this stage, far from certain) that $\eta_{e,\text{crit}}$ is a nonlinear threshold.

Note also that in the limit of strong drive ($\eta_e \gg \eta_{e,\text{crit}}$), (8.1) with $\beta = 1$ leads to the scaling $Q/Q_{e,\text{gB}} \propto (R/L_{Te})^3$, which is in accord with a simple scaling theory (Barnes, Parra & Schekochihin 2011; Adkins, Ivanov & Schekochihin 2023).¹

Electron-temperature-gradient turbulence in the pedestal is subject to further complexity (see the review by Ren *et al.* 2024). The structure of toroidal-ETG modes along the field lines (Jenko *et al.* 2009; Parisi *et al.* 2020) leads to nonlinear coupling between ion-scale ($k_y \rho_i \sim 1$) and electron-scale ($k_y \rho_e \sim 1$) modes. This can lead to quenching of ETG turbulence under certain conditions (Parisi *et al.* 2022). Other studies of the effects of toroidal-ETG modes on pedestal turbulence suggest that the contribution of toroidal-ETG modes amplifies transport (Chapman-Opoloioiu *et al.* 2025; Krutkin *et al.* 2025).

The global study by Leppin *et al.* (2023) of pedestal transport over a radial interval wider than the steep-gradient region suggests that the modes that constitute the turbulence are different between the pedestal-top and the steep-gradient region. At the top, ion-scale trapped-electron modes dominate the transport, whereas in the steep-gradient region, they are suppressed as a result of high $\mathbf{E} \times \mathbf{B}$ shear (Hahn & Burrell 1995). Consistent with this, the study of a JET pedestal by Leppin *et al.* (2024) concluded that the steep-gradient region was dominated by transport through the electron channel, with a large ion-scale heat flux at the top of the pedestal.

In spite of the above, ETG-centric models are unlikely to be a complete description of the turbulence in the pedestal. It has been shown by Hatch *et al.* (2016, 2021) that MTMs are responsible for magnetic fluctuations that are measured in pedestals in JET-C and JET-ILW. Large transport through MTMs was also found in the multi-machine study by Kotschenreuther *et al.* (2019), and these modes were recovered as non-negligible contributors to the turbulent heat flux in cases dominated by ETG (Chapman-Opoloioiu *et al.* 2022). Electromagnetic effects were also found to be important for the saturation of pedestal turbulence in the comparison carried out by Dicorato *et al.* (2024). Still, they found that turbulent transport had a much higher contribution from ETG modes for the type-I ELMy pedestal when compared with the small-ELM pedestals.

Aside from electron-scale turbulent transport, ion-scale effects can be relevant transport contributors through ion-scale turbulence and ion neoclassical transport (Trinczek *et al.* 2023). Predebon *et al.* (2023) found, in a study of the isotope effect, the electron-scale heat transport to be smaller than the ion neoclassical heat flux and the ion-scale turbulent heat flux close to the SOL, although the comparison was reversed close to the pedestal top. Furthermore, earlier work by Hatch *et al.* (2019) showed that, for a discharge in JET-C and one in JET-ILW, ion-scale and ion neoclassical transport could be larger than that originating from ETG turbulence.

In yet another body of work, the pedestal is considered to be limited by the instability of the kinetic ballooning mode (KBM; see Tang, Connor & Hastie 1980; Snyder 1999; Snyder & Hammett 2001; Pueschel, Kammerer & Jenko 2008; Waltz 2010). This mode is a source of high particle transport and effectively acts as a stiff transport threshold that limits the pressure gradient of the pedestal. This threshold is

¹Barnes *et al.* (2011) proposed this scaling for ITG turbulence, but its validity has recently been questioned by Nies *et al.* (2024). It is not clear whether a similar objection exists for ETG turbulence.

used in predictive models such as EPED (Snyder *et al.* 2009; Saarelma *et al.* 2019), and pedestals in various devices were consistently found by, e.g., Kotschenreuther *et al.* (2019) to be marginally stable to KBMs, with other forms of instabilities causing transport, such as MTMs.

Keeping in mind the foregoing discussion and acknowledging the potentially diverse nature of pedestal transport, we will make the simplifying assumption that the heat transport in the steep-gradient region is dominated by electron-scale turbulence. It is arguably not entirely unreasonable to claim that the ETG-centric model (8.1) can crudely represent pedestal turbulence, but we also extend our considerations to other heat-flux models that have been obtained by analysing a large dataset of gyrokinetic simulations by Hatch *et al.* (2022). Their table 1 contains several symbolic-regression fits to the heat transport in those simulations. These relationships are explicitly stated (see D2 and D3) and tested in Appendix D.1. Optimised versions of these relationships, including (8.1), will also be used in our analysis below in the forms of (8.12) and (8.13).

Further extrapolations obtained from simulation datasets are those found by Farçaş *et al.* (2024),

$$\frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} \propto \sqrt{\frac{m_e}{m_i}} \left(\frac{R}{L_{T_e}} \right)^{1.40} (\eta_e - 1)^{0.79} q^{-0.51} \beta_e^{-0.87} \lambda_D^{-0.51}, \quad (8.4)$$

and by Hatch *et al.* (2024),

$$\frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} \propto \sqrt{\frac{m_e}{m_i}} \left(\frac{d \ln T_e}{d\psi} \right)^2 (\eta_e - 1) \eta_e^{1.57} \tau_e^{-0.5} \lambda_D^{-0.4}. \quad (8.5)$$

Here λ_D is the electron Debye length normalised with respect to the electron gyro-radius ρ_e in Farçaş *et al.* (2024) and to the ion sound radius $\rho_s = \sqrt{m_i T_e} / eB$ in Hatch *et al.* (2024). These are relationships where the properties of the magnetic field appear via the safety factor q , the electron beta β_e , and the normalised temperature gradient with respect to the poloidal flux, $1/\tilde{L}_{T_e} \equiv d \ln T_e / d\psi$. Because we do not have a good magnetic-equilibrium reconstruction for our pedestals, we omit these models from our study.

8.2. Assumptions for the turbulent heat flux

The nonlinear dynamics of ETG turbulence provide a pathway to connect the separatrix loss power P_{Sep} with pedestal profiles. Using the scalings derived from gyrokinetic simulations, T_e -pedestal profiles will be reconstructed in § 8.3 by assuming that the heat flux is dominated by the electron-channel turbulent transport, $Q(R) = Q_{e,\text{Turb}}$, and that Q matches the experimentally determined separatrix loss power P_{Sep} , i.e.,

$$Q(R^{\text{Sep}}) = \frac{P_{\text{Sep}}}{S_{\text{Sep}}}, \quad (8.6)$$

where $S_{\text{Sep}} \approx 135 \text{ m}^2$ is the surface area of the separatrix. Further to this, we argue that, given the small width of the pedestal $w_{\text{ped}} = \mathcal{O}(\text{cm})$ compared with the device length scales R or $a = \mathcal{O}(\text{m})$, the variation in the area of the flux surface as a function of R is negligible. We also assume that the plasma profiles are in steady state (see § 2 and Frassinetti *et al.* 2020), so the rate of change of the thermal energy over the

small extent of the pedestal is small. Therefore, we approximate the turbulent heat flux as constant over the pedestal:

$$Q_{e,\text{Turb}}(R) \approx \frac{P_{\text{Sep}}}{S_{\text{Sep}}}. \quad (8.7)$$

To calculate the gyro-Bohm normalisation $Q_{e,\text{gB}}$ of the turbulent heat flux (see 8.2), we use a locally approximated value of the magnetic field, i.e.,

$$B(R) = \sqrt{B_\theta(R)^2 + B_\phi(R)^2} \approx B_{\text{vac}} \frac{R_0}{R} \sqrt{1 - \frac{1}{q_{95}^2}} \approx 0.66 B_{\text{vac}}, \quad (8.8)$$

where B_θ and B_ϕ are, respectively, the poloidal and toroidal components of the magnetic field, B_{vac} is the vacuum magnetic field on the axis of JET-ILW (shown as B in figure 2 and table 2), q_{95} is the safety factor evaluated at $\psi_N = 0.95$, and $R_0 = 0.92$ m is the radius of the axis of JET-ILW. This should remain a good approximation over the extent of the pedestal, irrespective of the divergence of the safety factor at the separatrix, $q \rightarrow \infty$ as $R \rightarrow R^{\text{Sep}}$. This approximation also agrees well with what one would obtain by using I_{plasma} and Ampere's law to estimate $B_\theta(R)^2$. These approximations fall short of the precision of the magnetic field calculated via the MHD equilibrium, but they are good enough for the purpose of gyro-Bohm normalisation.

Because of the approximations that we use, uncertainties remain as to the exact contribution of additional transport mechanisms such as neoclassical transport (Guttenfelder *et al.* 2021; Trinczek *et al.* 2023) and ELM-induced losses (Field *et al.* 2020). Order-unity errors of the heat flux do not significantly affect the reconstruction of pedestals because of the 'stiffness' of pedestal transport described in § 8.1. This property means that small changes in R/L_{T_e} in the models introduced below, (8.9)–(8.13), can lead to large changes in heat flux. Conversely, a large inaccuracy in the heat flux results in a comparatively small change in R/L_{T_e} , meaning that pedestal reconstructions are robust. This is illustrated in figure 25 (Appendix D.3), where an order-unity change in $Q_{e,\text{Turb}}/Q_{e,\text{gB}}$ (represented in figure 25d by a change in A) has only a small effect on R/L_{T_e} and the reconstructed T_e profile.

8.3. Electron-temperature reconstruction using heat-flux models

Using the methods outlined in Appendix C, we optimise the free parameters in (8.1) and the general forms of the heat-flux models from Hatch *et al.* (2022) given in (D2) and (D3). With these, we then reconstruct the T_e profiles.

Figure 19 shows the mtanh parameters that describe such reconstructions using models optimised over the pedestal region $\psi_N \in [\psi_{T_e,\text{Top}}, 1]$ (figure 19a–p) and the steep-gradient region $\psi_N \in [\psi_{n_e,\text{Top}}, 1]$ (figure 19q–t). The exact models that we used correspond to the panels of figure 19 as follows:

- (a–d) model (8.1) with fixed $\beta = 1$ and $\eta_{e,\text{crit}} = 1.4$, and the value of $A = 0.95$ optimised over the pedestal region $\psi_N \in [\psi_{T_e,\text{Top}}, 1]$; this corresponds to entry 8 in table 3:

$$\text{Model 1: } \frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = 0.95 \left(\frac{R}{L_{T_e}} \right)^2 (\eta_e - 1.4); \quad (8.9)$$

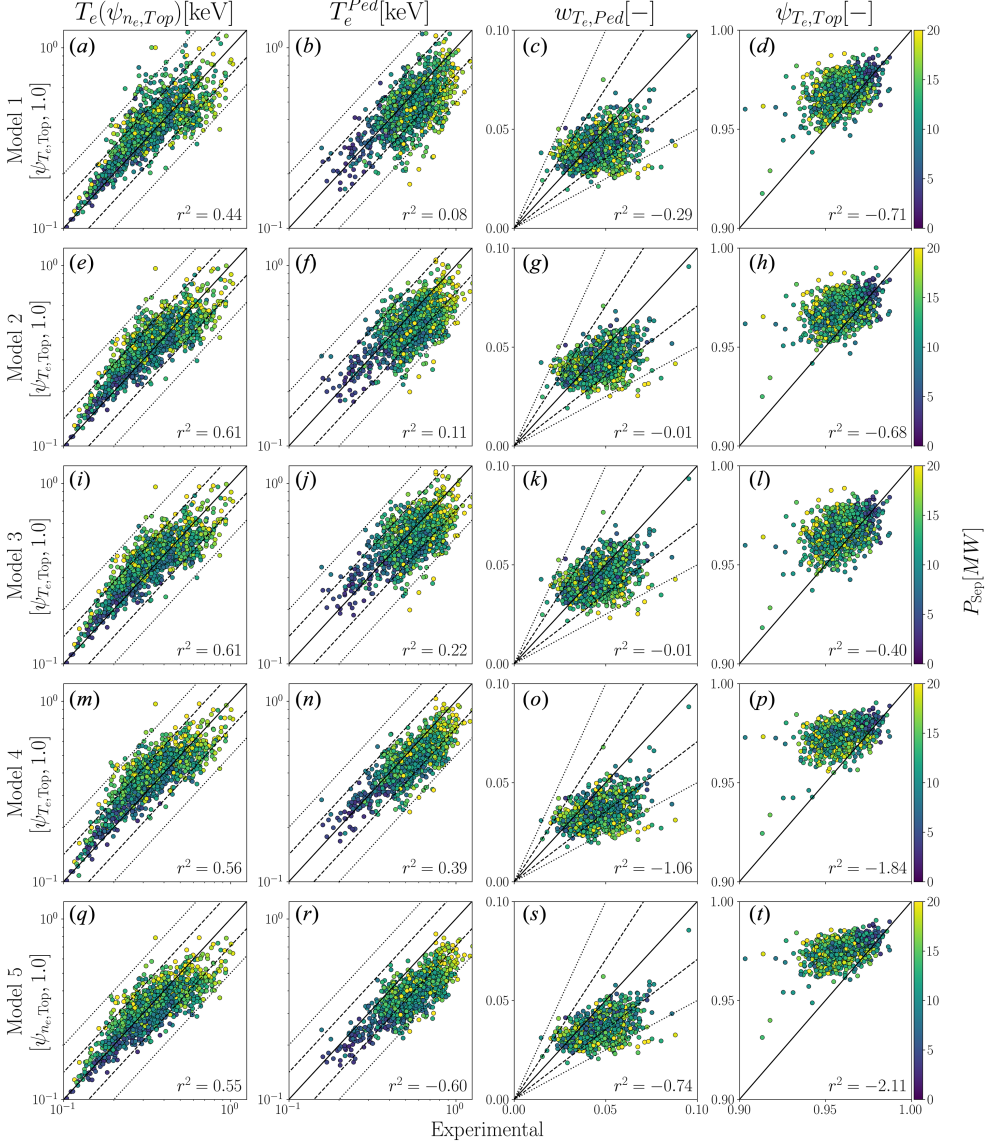


FIGURE 19. Database predictions obtained by inverting heat-flux models for R/L_{T_e} and integrating inwards from the separatrix. The models use best-fit parameters over the pedestal region (a–p) (models (8.9)–(8.12)) and over the steep-gradient region (q–t) (model (8.13)). The colour bar represents the separatrix loss power P_{Sep} . The plots follow the same layout as in figure 4.

(e–h) model (8.1) with fixed $\beta = 1$ and the values of ‘ $\eta_{e,\text{crit}}$ ’ = -0.55 and $A = 0.29$ optimised over the pedestal region $\psi_N \in [\psi_{T_e, \text{Top}}, 1]$; this corresponds to entry 17 in table 3:

$$\text{Model 2: } \frac{Q_{e, \text{Turb}}}{Q_{e, \text{gB}}} = 0.29 \left(\frac{R}{L_{T_e}} \right)^2 (\eta_e + 0.55); \quad (8.10)$$

(i-l) model (8.1) with the values of ‘ $\eta_{e,\text{crit}}$ ’ = -8.6 , $\beta = 3.75$, and $A = 10^{-3}$ optimised over the pedestal region $\psi_N \in [\psi_{T_e,\text{Top}}, 1]$; this corresponds to entry 14 in table 3:

$$\text{Model 3: } \frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = 10^{-3} \left(\frac{R}{L_{T_e}} \right)^2 (\eta_e + 8.6)^{3.75}; \quad (8.11)$$

(m-p) model (D2) from Hatch *et al.* (2022) with the values of the numerical constants optimised over the pedestal region $\psi_N \in [\psi_{T_e,\text{Top}}, 1]$; this corresponds to entry 22 in table 3:

$$\text{Model 4: } \frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = 9.64 \frac{R}{L_{T_e}} (\eta_e^{2.92} + 16.02); \quad (8.12)$$

(q-t) model (D3) from Hatch *et al.* (2022) with the values of the numerical constants optimised over the steep-gradient region $\psi_N \in [\psi_{n_e,\text{Top}}, 1]$; this corresponds to entry 30 in table 3:

$$\text{Model 5: } \frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = 0.01 \frac{R}{L_{T_e}} (\eta_e + 8.86)^{4.77}. \quad (8.13)$$

A more comprehensive list of accuracy metrics for various heat-flux models is given in table 3 of Appendix D. There, we present the performance of the exact models from Chapman-Oploupoiou *et al.* (2022) and Hatch *et al.* (2022): those given by (8.9)–(8.13) and their variations that minimise the mean-squared difference between the reconstructed and the fit T_e profiles. We also compare them with the results of §§ 4, 6, and 7.

We note that the optimised prefactor $A = 0.95$ of (8.1) over the steep-gradient region is in good agreement with the value of $A = 0.85$ proposed by Chapman-Oploupoiou *et al.* (2022).

Model (8.9) maintains the physical intuition of ETG-based turbulent transport outlined in § 8.1 and manages to capture the shapes of pedestals much better than the simplistic models of §§ 6 and 7. Model (8.10) is a further improvement to all metrics of pedestal shape, although the optimisation of $\eta_{e,\text{crit}}$ yields a negative value over the ranges of the pedestal or steep-gradient region (see entries 12, 14, 16, and 17 in table 3). The improvement in reconstruction accuracy by setting this value to be negative casts doubt on the idea of a turbulent threshold $\eta_{e,\text{crit}}$. Models (8.9) and (8.10) share two systematic limitations: R/L_{T_e} is over-predicted close to the separatrix and the reconstructed $\psi_{T_e,\text{Top}}$ tracks the experimental $\psi_{n_e,\text{Top}}$. The inaccuracy in R/L_{T_e} comes as a result of the following process, illustrated in figure 20: the large value of $Q_{e,\text{Turb}}/Q_{e,\text{gB}}$ near the separatrix yields a large R/L_{T_e} (panel d), and this then quickly increases the value of the T_e profile as the integration goes on, reducing the value of $Q_{e,\text{Turb}}/Q_{e,\text{gB}}$ faster than the experimental profile suggests (panel c). This means that in the steep-gradient region of the T_e pedestal, the temperature gradient is always over-predicted, with an over-estimated R/L_{T_e} at the separatrix. Furthermore, the increase of n_e away from the separatrix increases $Q_{e,\text{gB}}$, which reduces $Q_{e,\text{Turb}}/Q_{e,\text{gB}}$. This produces the shape of the T_e pedestal, by switching to low values of R/L_{T_e} away from the separatrix, and, in the process, causing $\psi_{T_e,\text{Top}}$ to match $\psi_{n_e,\text{Top}}$ (see figure 19d,h,l and 20a).

Seeking optimal values of the exponent β in (8.1) yields large exponents for the term containing the gradient ratio η_e (entries 13 and 14 in table 3). In spite of this,

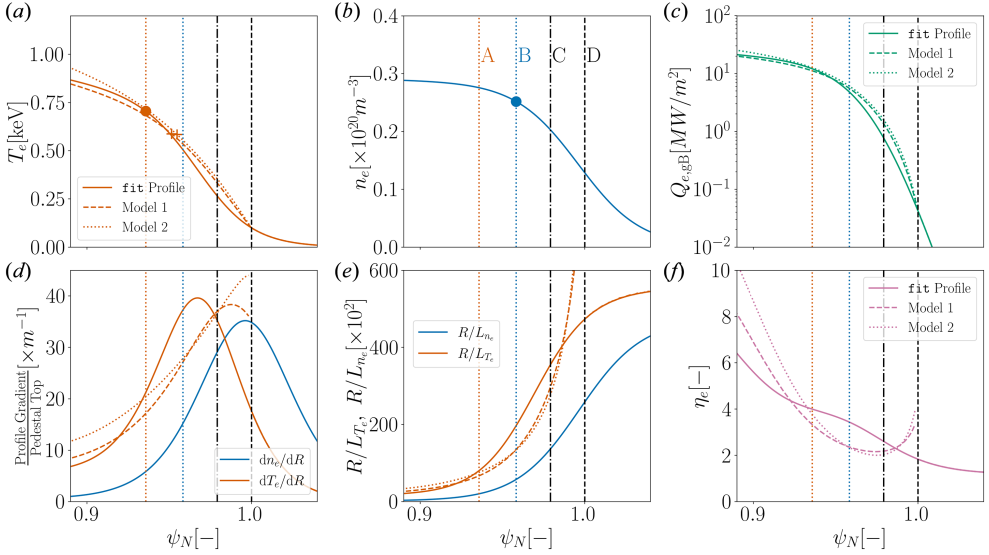


FIGURE 20. Pulse #90339 fit profiles compared with predictions using heat-flux models 1 (8.9) and 2 (8.10). The layout is the same as in figure 1. The fit profiles are shown as solid lines, with the predictions using the heat-flux models shown as dashed (model 8.11), and dotted (model 8.9) lines. The fit is done over the interval $[\psi_{T_e, \text{Top}}, \psi_{\text{Sep}}]$. The vertical lines and dots represent the landmarks of the fit profiles (see § 2.4 and figure 1), while the crosses represent the pedestal top of the T_e predictions.

the parameter scan shown in figure 25 (Appendix D.3) suggests that reconstructions using (8.1) are not very sensitive to the exponent β , and that the corresponding reduction of the predicted values of R/L_{T_e} are compensated by the decrease in the prefactor A to 10^{-3} (effectively increasing the turbulent heat flux). This insensitivity of the qualitative shape of predictions to the exponent β is also present for models (8.12) and (8.13).

The predictions using models (8.12) and (8.13) (derived from Hatch *et al.* 2022) give improved r^2 values of the predictions of $T_e(\psi_{n_e, \text{Top}})$ and T_e^{Ped} (see figure 19 m, n, q, r) when compared with those obtained from model (8.9). However, the shape of pedestals suggests that there is little variation in $w_{T_e, \text{Ped}}$ and $\psi_{T_e, \text{Top}}$ (see figure 19 o, p, s, t), and that the reconstructed pedestals are very thin and pushed towards the separatrix. This signals the importance of a larger exponent of the temperature gradient in such models (such as $R^2/L_{T_e}^2$ in 8.1), since the use of $(R/L_{T_e})^1$ in (8.12) and (8.13) (corresponding to models (D2) and (D3) further discussed in Appendix D) demands a large value of the temperature gradient in order to balance the magnitude of $Q_{e, \text{gB}}/Q_{e, \text{Turb}}$ near the separatrix. This temperature gradient produces outwards shifted T_e pedestals, as illustrated in figure 21. The shear magnitude of R/L_{T_e} near the separatrix also means that there is little sensitivity of the m and n parameters to the n_e profile for models (8.12) and (8.13).

Another important common characteristic of our fit parameters is the embarrassing tendency of $\eta_{e, \text{crit}}$ to become negative. This is manifest in (8.10)–(8.13) and in the majority of optimisations in table 3, and we discuss this at length in Appendix D. While the sign of $\eta_{e, \text{crit}}$ does not have a dramatic effect on the reconstructed T_e profiles (see figures 19 a – h and 20), we interpret this result as an indication that either

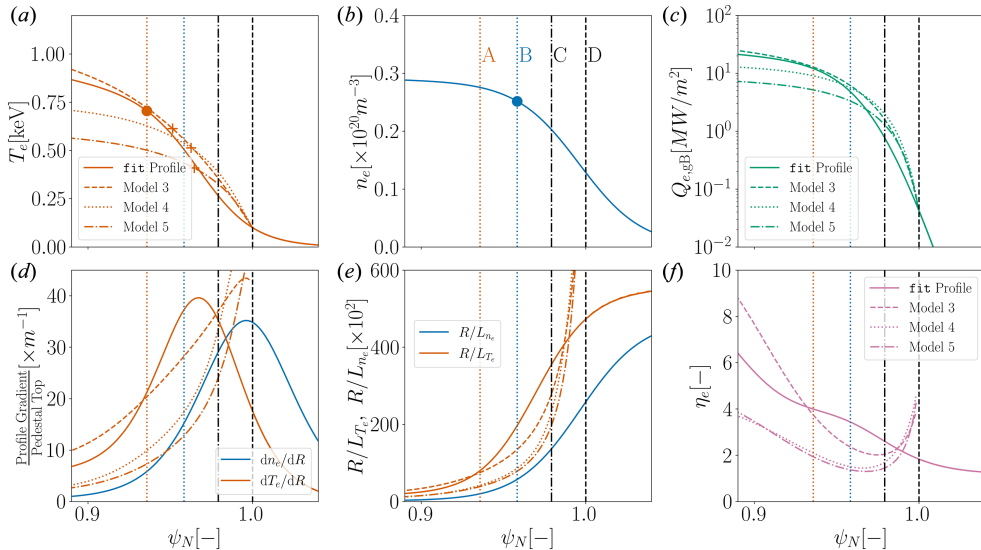


FIGURE 21. Same as figure 20, but for predictions using heat-flux models 3 (8.11), 4 (8.12), and 5 (8.13). The fit profiles are shown as solid lines, while the predictions using different heat-flux models are shown as dashed (model (8.10)), dotted (model (8.12)), and dash-dotted (model (8.13)) lines. The fit is done over the intervals $[\psi_{T_e, \text{Top}}, \psi_{\text{Sep}}]$ for models (8.11) and (8.12), and $[\psi_{n_e, \text{Top}}, \psi_{\text{Sep}}]$ for model (8.13).

the forms of the models for $Q_{e, \text{Turb}}$ are inaccurate for the data, or that the physical argument for $Q_{e, \text{Turb}} \propto \eta_e - \eta_{e, \text{crit}}$ for some physical threshold $\eta_{e, \text{crit}}$ is not supported by the data.

8.4. Discussion of heat-flux model results

In this section, we have obtained proof that the models of Guttenfelder *et al.* (2021), Chapman-Oplopoiou *et al.* (2022), and Hatch *et al.* (2022) perform well in reproducing the mtanh parameters of pedestals (figure 19), alongside matching the experimental profile well over the regions with steep gradients (table 3) when compared with our simple database-wide models from §§ 6 and 7.

A thorough analysis of these gyrokinetics-inspired turbulent transport models has confirmed the necessity of R/L_{T_e} appearing with an exponent larger than unity in the expression for $Q_{e, \text{Turb}}/Q_{e, \text{gB}}$, as also recognised by models (8.4) and (8.5). We have also found a unanimously negative sign of the constant $\eta_{e, \text{crit}}$, however, this physically dubious feature is not demanded strongly by the predictions: forcing $\eta_{e, \text{crit}}$ to be positive does not severely impact prediction performance.

We note that all of our pedestal reconstructions using optimised heat-flux models generally produce similar mtanh parameters. The minimisation of the sum of squared discrepancies between prediction and experiment also yields comparable reconstruction accuracy between fits done across the same range of the pedestal (entries 7–30 in table 3). These two outcomes mean that a significant part of our assessment of which transport models are best comes from the qualitative shape of the profiles such as those in figures 20 and 21 (as we discuss at the end of § 8.3).

Overall, the inclusion of the turbulent heat flux into our pedestal-reconstruction methods addresses to a noticeable degree the shortcomings of the ‘marginal’ models

that used database-wide parameters in §§ 6 and 7. This said, the fact that the reconstructed T_e pedestals inherit the shape of the n_e pedestals is a sign that such models require more complexity in order to capture the density-top region correctly, and hence, the shift between the T_e and n_e pedestals (Frassinetti *et al.* 2020). This is also supported by the significantly better performance of our neural-network models in § 4.2, which successfully learned this difference in turbulent particle and energy transport.

It is also worth noting that the absence in the database of the exact profiles of the magnetic field, e.g., $q(R)$, $B(R)$, $\hat{s}(R)$, limits the range of models that can be tested without including the results of numerical magnetic-field reconstruction.

9. Summary and conclusions

In this paper, we have analysed the JET-ILW pedestal database (Frassinetti *et al.* 2020) and reconstructed pedestal temperature profiles using density profiles and other engineering parameters.

In § 2, we described the pedestal profiles used in this study and defined the pedestal parameters via `mtanh` fits. We showed the distributions and correlations of database parameters in § 3 and Appendix A. This overview helped frame § 4, which navigated this parameter space using ML methods. Our objective was to use neural networks to predict T_e profiles. Different parameter sets were tested and, in § 4.1, we found that using the full n_e profiles yielded excellent results for reconstructing T_e pedestals, in particular when all of the available information was included. This was shown both in terms of `mtanh` metrics such as pedestal height and location (figure 4), and in terms of the mean-squared difference between prediction and experiment over the steep-gradient region (table 3). In § 4.2, local predictions of the T_e profile showed good results for the pedestal heights, although there were inaccuracies in predicting the pedestal width and location (figure 6).

The ability to reconstruct pedestals with excellent quality using neural networks suggests that predicting T_e pedestals using n_e profiles and other database parameters is a well-posed problem, both because of the underlying physics being governed by reproducible phenomena, and because the dataset is statistically robust. The success of local predictions indicates that this task can be achieved using models that take into account the local amplitudes and gradients of the pedestal. This said, we have observed that the inclusion of the radial location R of the local prediction resulted in improved performance. This could be explained by the fact that, during training, the networks developed a general model of the large-scale structure of pedestals, and the inclusion of the radial location indirectly gave an indication of global physics, improving the prediction. In § 4.3, we explicitly identified the most important parameters for accurate T_e prediction in the global and local cases: the separatrix loss power P_{Sep} , the plasma current I_{plasma} , the fuelling rate Γ_D and the strike-point configuration.

In § 5, we discussed the importance of the normalised temperature and density gradients R/L_{T_e} and R/L_{n_e} in the context of known linear ETG mode stability thresholds (figure 9). We found that all steep-gradient region pedestals were in an order-unity-supercritical regime, which suggested that the gradient ratio η_e was an important parameter. In § 6, our analysis of η_e showed that it is not a good descriptor of the T_e and n_e inside of the density-top location because the values of η_e range over several orders of magnitude. We found that $\eta_e \approx 2$ across the steep-gradient region (figure 11), confirming across a large database the previous experimental

observations of values of η_e in the pedestal. This result could be interpreted as pointing to the existence of a nonlinear threshold $\eta_{e,\text{crit}}$ in pedestals, or as an indication that the turbulence saturates in an order-unity-supercritical regime to match the transport necessary to maintain a steady state.

Next, we assumed a simple prescription for the relationship between the pedestal gradients R/L_{T_e} and R/L_{n_e} : a constant η_e across the pedestal. This was empirically motivated by the clustering of η_e around 2 in the steep-gradient region, and is equivalent to assuming that turbulence is marginal, close to an (in general, nonlinearly) established threshold. Reconstructing T_e pedestals using a given η_e value gave reasonable but not very accurate pedestal predictions (figure 12). Using an optimisation algorithm to find a nominal value of η_e across the database yielded values between 1.9 and 2.1. Pulse-by-pulse fits of η_e performed much better at reproducing the pedestal region (figure 12a–d), although T_e pedestals were shifted outwards towards the separatrix, in line with the n_e pedestals. The pulse-by-pulse optimisation yielded a distribution of fit values of η_e that matched the experimental η_e in the middle of the steep-gradient region (figure 13), suggesting that the value of η_e is a reliable indicator of the transport processes in the pedestal.

The results shown in figure 10 hinted at a power-law relationship between R/L_{T_e} and R/L_{n_e} of the form $R/L_{T_e} = A(R/L_{n_e})^\alpha$, a possibility that was explored further in §7. There, we examined the distribution of the exponent α of such a hypothetical power-law relationship between the temperature and density gradients. While this lacked physical grounding, it was a simple model for our data. In §7.2 we found that the best database-wide value of the exponent was $\alpha \approx 0.4$ for the pedestal region (figure 16e–h). Fitting α and A independently for each pulse gave excellent agreement over the steep-gradient region (figure 16a). The resulting pairs of A and α all turned out to lie on the similarity line $\ln A = 5.64 - 4.42\alpha$ (figure 18), giving the relationship (7.3) to describe all pedestals of our JET-ILW subset of the database. Remarkably, a large set of JET-C pulses also obeyed a similar linear relationship between the gradients. This was one of the clearest modelling insights that we obtained in this study, even though we could not find a relationship between the pedestal wide α and other database parameters. Constructing a model for α would complete the predictive model of the steep-gradient region, but the lack of physical intuition for such a power-law relationship makes this task difficult.

Finally, in §8, we explored the role of the separatrix loss power P_{sep} in relating temperature and density gradients. We focused our analysis on transport models inspired by gyrokinetic simulations (Guttenfelder *et al.* 2021; Chapman-Oplopoiou *et al.* 2022; Hatch *et al.* 2022). In §8.3 and Appendix D, we computed the best-fit parameters for these models and discussed the quality of the resulting T_e -profile reconstructions.

The models of the form (8.1) (Guttenfelder *et al.* 2021; Chapman-Oplopoiou *et al.* 2022) produced good reconstructions of the T_e profile without major adjustments. Such models have a theoretical root in slab-ETG-driven turbulent transport, as described in §8.2. Reconstructions of pedestal profiles using parameters optimised for these models and those found by Hatch *et al.* (2022) provided excellent values for the T_e -profile values at the top of the n_e and T_e pedestals as a result of our optimisation process, although the pedestal shapes were subject to systematic inaccuracies (figure 19). The temperature gradients were wildly over-estimated near the separatrix when the formulae used for the heat flux included a prefactor of $(R/L_{T_e})^1$ (figure 21), whereas models with $(R/L_{T_e})^2$ resulted in much better reconstructions (figure 20). This is a confirmation that the locally relevant length scale L_{T_e} must be

used for the gyro-Bohm estimate of the heat flux $Q_{e,\text{gB}}$. Optimising for the critical value of the gradient ratio $\eta_{e,\text{crit}}$ almost always resulted in negative values (table 3), undermining the physical picture that we described in §8.2. We interpret this as a sign that the models that we analysed are missing key elements of the underlying physics of the pedestal.

In spite of the shortcomings of our heat-flux models, the results of §8 confirm that we have successfully accounted for the separatrix loss power in order to find the electron-temperature gradients. These results were a major improvement over the predictions that used database-wide fitting parameters in §§6 and 7. However, the much better quality of the local-neural-network reconstructions of §4.2 suggests that the upper bound of prediction accuracy has not yet been reached by the ‘physics-based’ models. It appears that, in order to obtain a more accurate model of pedestal transport, a clearer theoretical understanding of the underlying gyrokinetic turbulence is required, and, also, more clarity regarding whether the pedestal-shape parameters, such as the width and position of the pedestal, are determined by local turbulence or by device-scale physics. Extending our investigation to other devices could give the key for creating a simple generalised model for pedestal transport.

Our analysis aims to reconstruct the pedestal profiles of electron temperature in the JET-ILW pedestal database. While a complete local-transport model would predict the height, width, and position of a pedestal – essential parameters for tasks such as core-edge integrated modelling – such a model would also describe the structure of the plasma on scales smaller than the pedestal width. Because of this focus on the small scales and transport models, this is a different approach from that of existing pedestal models based on MHD stability such as EPED (Snyder *et al.* 2009). In fact, local-transport models could perhaps remedy the occasional inaccuracy of EPED predictions. Finally, an ELM-free and transport-limited pedestal (Parisi *et al.* 2025) would be governed by self-consistent transport prescriptions such as those we have tested in this study.

Acknowledgements

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 number 101052200 – EUROfusion) and by the EPSRC Energy Programme (grant number EP/W006839/1). Views and opinions expressed herein are, however, those of the authors 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. The authors would like to thank the referees for their attention to detail and constructive criticism.

Editor Nuno Loureiro thanks the referees for their advice in evaluating this article.

Funding

The work of L.-P. T. was supported by an EPSRC CASE studentship in partnership with UKAEA. The work of A. A. S. was supported in part by EPSRC (grant number EP/R034737/1), the STFC (grant number ST/W000903/1), and by the Simons Foundation via a Simons Investigator Award. The work of L. F. was supported by Vetenskapsrådet (grant number 2023-04895). L.-P. T. would like to thank M. Abazorius, G. Acton, B. Chapman-Oploupoiou, W. A. Clarke, R. Greif, D. R. Hatch, P. G. Ivanov, S. Newton, P. Reichherzer, J. Ruiz Ruiz, M. Senstius, and F. C. Spangle for helpful advice and discussions.

Data availability statement

The EUROfusion pedestal database is available through IMAS (the ITER Integrated Modelling and Analysis Suite; Imbeaux *et al.* 2015) on the EUROfusion Gateway Infrastructure, accessible at <https://wiki.eufus.eu/doku.php>. Access to the pedestal database and JET PPFs is restricted to members of EUROfusion. The neural networks obtained in this study directly represent the data and are, therefore, subject to the same access restriction. The code used for data analysis is available upon reasonable request by contacting the authors at leonard-petru.turica@physics.ox.ac.uk.

Declaration of interests

The authors report no conflict of interest.

Appendix A. Pulse parameters and correlations

In this appendix, we will list the definitions of each plasma parameter included in our study of the pedestal database. The distributions of these parameters across the database are shown in [figure 2](#). Their measurement units, as well as their minimum, maximum, mean, and median values are given in [table 2](#). This is relevant context for carrying out profile predictions using neural networks in §4. We loosely distinguish two groups of parameters, based on their origin: engineering parameters and magnetic-equilibrium parameters.

In [table 2](#) and [figure 3](#), we also present the distributions of the mtanh fit parameters, which describe the T_e and n_e profiles in our databases, alongside other relevant parameters such as the separatrix density n_e^{Sep} , the temperature at the density-top location $T_e(\psi_{n_e, \text{Top}})$, and the ratio between the separatrix density and the pedestal-top density $n_e^{\text{Sep}}/n_e^{\text{Ped}}$. The correlation matrix for all the parameters presented in [table 2](#) is presented in [figure 22](#). This is similar to [figure 1](#) of Kit *et al.* (2023), adapted to our dataset of 1251 pulses.

A.1 Pulse-parameter definitions

The engineering parameters used in this study are:

- the separatrix loss power P_{Sep} ;
- the deuterium gas fuelling rate Γ_D ;
- the line-integrated electron density following one of the interferometer chord diagnostics of JET through the plasma edge \bar{n}_e ;
- the plasma current I_{Plasma} ;
- the strike-point configuration, defined by which of the divertor plates the footprints of the open field lines rest upon, with several basic configurations described below.

The separatrix loss power is computed from the power balance of the plasma, *viz.*,

$$P_{\text{Sep}} = P_{\text{Ohm}} + P_{\text{NBI}} + P_{\text{ICRH}} - P_{\text{Shine}} - P_{\text{Rad}}, \quad (\text{A1})$$

where the different power components are the Ohmic heating power P_{Ohm} , the neutral-beam injection power P_{NBI} , the ion-cyclotron-resonance heating power P_{ICRH} ,

Parameter	Units	Min	Max	Mean	Median
P_{Sep}	[MW]	2.37	27.0	11.9	11.7
Γ_D	[e ⁻ /s · 10 ²²]	0.0	22.3	2.08	1.58
I_{Plasma}	[MA]	0.972	3.97	2.12	1.99
B	[T]	0.967	3.68	2.37	2.38
\bar{n}_e	[m ⁻² · 10 ²⁰]	0.0	0.907	0.434	0.430
κ	[-]	1.58	1.82	1.68	1.68
$\langle \delta \rangle$	[-]	0.153	0.456	0.265	0.262
δ_{Lower}	[-]	0.232	0.492	0.327	0.332
δ_{Upper}	[-]	0.021	0.461	0.2	0.182
q_{95}	[-]	2.42	5.44	3.46	3.37
β_θ	[-]	0.226	1.41	0.615	0.551
ν^*	[-]	0.002	0.133	0.021	0.017
n_{Gw}	[m ⁻³ · 10 ²⁰]	0.361	1.43	0.787	0.768
f_{Gw}	[-]	0.29	1.05	0.729	0.734
T_e^{Ped}	[keV]	0.157	1.38	0.573	0.556
$w_{T_e, \text{Ped}}$	[-]	0.014	0.144	0.046	0.045
$\psi_{T_e, \text{Top}}$	[-]	0.899	0.987	0.96	0.961
n_e^{Ped}	[m ⁻³ · 10 ²⁰]	0.179	1.08	0.457	0.44
$\psi_{n_e, \text{Top}}$	[-]	0.923	0.999	0.979	0.98
$w_{n_e, \text{Ped}}$	[-]	0.016	0.231	0.045	0.041
$T_e(\psi_{n_e, \text{Top}})$	[keV]	0.104	1.032	0.366	0.325
n_e^{Sep}	[m ⁻³ · 10 ²⁰]	0.058	0.658	0.273	0.254
$n_e^{\text{Sep}}/n_e^{\text{Ped}}$	[-]	0.212	0.988	0.596	0.597

TABLE 2. Relevant engineering parameters, magnetic-equilibrium parameters and profile values in the database, also depicted in figures 2 and 3. Each group of parameters is separated by a horizontal bar. The measurement units, ranges, mean values and median values of the parameters are displayed for the subset of the database analysed in this paper.

the shine-through power P_{Shine} (the power injected but not absorbed by the plasma), and the power radiated by the bulk of the plasma P_{Rad} .

The strike-point configurations are formatted as the ‘inner strike point/outer strike point’ in figure 2. The different ways to contact the divertor plates are: V for vertical, C for corner, and H for horizontal. A depiction of these strike-point configurations in JET-ILW can be found in figure 5 of Frassinetti *et al.* (2020). For the purpose of computing parameter correlations in figure 22, the strike-point configurations (denoted as SPC in the figure) are encoded by integer values: V/H → 0, C/H → 1, V/V → 2, C/V → 3, V/C → 4, and C/C → 5.

The magnetic-equilibrium parameters used in this study are:

- the magnetic field on the axis B ;
- the safety factor $q_{95} = d\Phi/d\psi$ evaluated at $\psi_N = 0.95$ (here, Φ and ψ are the toroidal and poloidal magnetic fluxes, respectively);
- the δ_{Upper} and δ_{Lower} plasma triangularities, used in the Miller equilibrium parametrisation (Miller *et al.* 1998);
- the plasma elongation, Miller equilibrium index κ ;



FIGURE 22. Matrix of Pearson correlation coefficients computed for engineering parameters, magnetic-equilibrium parameters, and relevant profile parameters across the database. Here SPC refers to the strike-point configurations, as denoted in [Appendix A.1](#). The upper triangle contains the numerical values of the correlation coefficients, the lower triangle represents the same values pictorially. Values that are smaller in magnitude than 0.10 are omitted.

- the kinetic pressure normalised by the poloidal magnetic energy density, $\beta_\theta = 2\mu_0 p / B_\theta^2$;
- the electron-electron collisionality at the pedestal top ν^* , normalised to the thermal ion bounce frequency, as per (7) of Frassinetti *et al.* (2020);
- the Greenwald density limit n_{Gw} , a measure of the maximum line-averaged density at the edge of the plasma, defined as $n_{\text{Gw}} = I_{\text{plasma}} / \pi a^2$, where a is the minor radius of the device;
- the Greenwald fraction f_{Gw} , the ratio between the measured line-averaged density and n_{Gw} .

Appendix B. Neural networks and methods

This appendix details the ML methods used to reconstruct T_e in our database. We will describe our methods of handling data, including selection, processing, and augmentation, alongside our choice of training strategies and network architecture.

B.1 Data preprocessing

The database that we use includes parameters associated with each pulse, which can have discrete values (e.g., the strike-point configuration) or continuous ones (e.g., the database parameters or the profiles). The aim of our data preprocessing is to translate this data into a series of scalars that can be assembled into structured vectors of labels (inputs) and features (outputs) to describe each pulse. These sets of labels and features will then be used to reconstruct T_e profiles using a neural network.

In order to be able to make use of the strike-point configurations through our numerical method, the discrete-valued strike-point configurations (which can take six possible values) are ‘one-hot’ encoded to a six-dimensional vector. Consequently, any strike-point configuration is represented by a vector that has all components equal to 0, with the exception of one component that is 1. The location of the non-zero component uniquely maps to each strike-point configuration.

The profiles of n_e , T_e , R/L_{n_e} , and R are sampled from `fit` profiles at 128 (§§ 4.1 and 4.2) or 64 (§ 4.3) equally spaced radial locations. The radial locations are chosen in an interval between $R = 3.7$ m and $R = 3.85$ m. This range encompasses all points located between $\psi_{T_e, \text{Top}}$ and ψ_{Sep} . The use of R as the radial coordinate (as opposed to ψ_N or X_{Ren}) is made in order to avoid the inclusion of data that contains *a priori* information about the T_e pedestal – as ψ_N contains information obtained from the equilibrium reconstruction of the plasma profiles, and X_{Ren} makes use of the pedestal-top locations.

Where available, experimental uncertainties are included. These can be associated with pulse parameters or profiles. For `fit` profiles, we include the `raw` uncertainties. The experimental errors included in `raw` profiles are sampled at the same 128 or 64 radial locations using the following procedure: the `raw` values of the experimental uncertainties of the n_e and T_e profiles are convolved with a Gaussian kernel of full-width-half-maximum matching the instrument function of the HRTS diagnostic (Frassinetti *et al.* 2012). Then, a low-pass filter is applied to obtain a smooth numerical interpolation of these convolved errors. The convolved and filtered experimental errors offer an excellent estimate for the mean uncertainties at each location in the pedestal. One such set of uncertainties is displayed in figure 23 for pulse #90339. These values will be used in data augmentation, as will be explained next.

B.2 Augmentation

Augmentation is essential for the stability and convergence of neural networks, and plays an important role in transforming our limited database of 1251 pulses into a statistically meaningful training set. For robust interpolation, it is also important to introduce awareness of the experimental uncertainties into the neural network. Therefore, we indirectly weight each pulse based on how sparse the neighbouring parameter space is. We will describe below our measure of parameter-space sparsity and how we use it to decide the amount of augmentation (duplication) that each piece of data must be subjected to. We also add random Gaussian noise to the augmented data in order to model the experimental uncertainties.




FIGURE 23. The n_e (blue) and T_e (orange) profiles for pulse #90339, with 64 numerical samples at equally spaced radial locations. The error bars represent the experimental uncertainties, obtained as described in [Appendix B.1](#). The green box represents a subset of L consecutive numerical values used in the training of neural networks. The beginning position of the subset of length L is chosen randomly as part of the augmentation process described in [Appendix B.2](#).

For the purpose of assembling the labels, each separate set of them consists of N different parameters, not including profiles. These pulse parameters are therefore treated as populating an N -dimensional space, for which we first calculate the parameter-space density. Consider one of the possible parameters j , whose value for each pulse is $j(\text{Pulse})$. We calculate the parameter-space density using a Gaussian-kernel density estimator as

$$\rho_j(x) \propto \sum_{\text{Pulse}} \exp \left\{ -\frac{[j(\text{Pulse}) - x]^2}{2h^2} \right\}, \quad (\text{B1})$$

where x is the value of j where the parameter-space density is evaluated and h is the full-width half-maximum of the Gaussian kernel. We set $h = [\max(j) - \min(j)]/10$ in order to discern distributions that are not unimodal. Any overall prefactor in (B1) is irrelevant in this method, as will be seen below.

Using the single-parameter density evaluated for each individual pulse, $\rho_j(\text{Pulse}) \equiv \rho_j[j(\text{Pulse})]$, we calculate the number of times each pulse must be duplicated (augmented) for the dataset. First we define the ‘augmentation score’ for each pulse,

$$S(\text{Pulse}) = \sum_j \exp \left[\frac{d + 1}{d + \rho_j(\text{Pulse})/\max(\rho_j)} \right], \quad (\text{B2})$$

where d is a positive constant and the sum is over all parameters. Definition (B2) assigns the highest scores to the least represented pulses in parameter space. Using a value of d of around 0.2–0.4 results in an approximately uniform distribution of S . The augmentation number of each pulse, defined as the number of times this pulse will appear in the training set as a result of the augmentation process, is a linear rescaling of the augmentation score, i.e.,

$$A(\text{Pulse}) = A_{\min} + (A_{\max} - A_{\min}) \frac{S(\text{Pulse}) - \min(S)}{\max(S) - \min(S)}, \quad (\text{B3})$$

where A_{\min} and A_{\max} are chosen as follows: 1200 and 2500, respectively, for full-profile predictions (§ 4.1); 2000 and 3500 for local predictions (§ 4.2); 500 and 5000

for the parameter scans (§ 4.3). This means that, for the pulses that sample high-density regions of the parameter space, $A(\text{Pulse})$ will be smaller than for those in sparser regions. Consequently, the pulses in dense regions will appear fewer times in the training set of our neural network to compensate for the density. This effectively balances our dataset.

The choice of the augmentation score (B2) was made so as to avoid the behaviour of some alternatives, e.g., $S(\text{Pulse}) = [\prod_j \rho_j(\text{Pulse})]^{-1/N}$, that caused $> 95\%$ of pulses to end up with the maximum augmentation score $A(\text{Pulse}) \approx A_{\max}$ when N was large, which would have defeated the purpose of the augmentation procedure.

Our augmentation procedure includes a random cropping of the profiles of T_e , n_e , and R/L_{n_e} , if they are included in the information used in the network. The same cropping is also applied to the values of the radial location R where the profiles are sampled numerically. This size reduction also includes selecting randomly located subsets of L consecutive values from each profile, as illustrated in figure 23, and discarding the rest of the profile. For full-profile predictions, this value of L is set to 80% of the number of numerical samples of each profile so as to allow the data to be subject to radial ‘jitter’. For local-value predictions, $L = 3$ for each profile, effectively selecting an extremely narrow (one-point) sample. The random positioning of the subsets helps the network match the features of the profiles to the prediction, avoiding an inflexible memorisation based on the exact radial location of each numerical value. The inclusion of R in the input data is used to test the importance of the location for the prediction.

So far, each pulse is represented in the dataset by $A(\text{Pulse})$ vectors of labels and $A(\text{Pulse})$ vectors of features. Each of the label vectors has length $N + ML$, being composed of the N pulse parameters and M cropped profiles of length L . The feature vectors have length L , representing the cropped T_e profile.

All resulting entries in the augmented dataset (parameters and profiles) also include Gaussian random noise matching the experimental uncertainties in each individual value. This noise is added to each pulse parameter, i.e.,

$$j(\text{Pulse}) \rightarrow j(\text{Pulse}) + \mathcal{F}[\sigma_j(\text{Pulse})], \quad (\text{B4})$$

where $\sigma_j(\text{Pulse})$ is the experimental uncertainty of the parameter j in a given pulse, and $\mathcal{F}[\sigma]$ is a random sample from a normal distribution with standard deviation σ and zero mean. Similarly, each point in the cropped profiles of T_e , n_e , and R/L_{n_e} is subjected to random noise corresponding to the experimental uncertainty evaluated at the point’s location.

B.3 Normalisation and loss

The augmentation procedure described above results in a dataset comprised of vectors of labels of length $N + ML$ and vectors of features of length L , containing the concatenated parameters and profile values. For numerical stability, all of the values in these vectors are normalised after augmentation. Depending on the case, we map each label and each feature independently to a distribution with mean zero and standard deviation 1. These have no effect on the results, but they are sometimes necessary to avoid divergence during training.

In order to train the neural network, a mean-squared loss function between the target features and the predicted features is minimised. This loss function is defined to be

$$\mathcal{L} = \frac{1}{LA_{\text{total}}} \sum_{k=1}^{A_{\text{total}}} \sum_{i=1}^L [(y_{k,i})_{\text{true}} - (y_{k,i})_{\text{pred}}]^2, \quad (\text{B5})$$

where $A_{\text{total}} = \sum_{\text{Pulse}} A(\text{Pulse})$ is the total number of label-feature pairs of vectors, $(y)_{\text{true}}$ are the target values (in our case, the normalised and cropped values of T_e), and $(y)_{\text{pred}}$ are the features predicted by the neural network. We denote by $y_{k,i}$ the i 'th value of the k 'th feature vector.

Additionally, L_2 regularisation is used in the optimisation of the model, i.e., each linear transformation layer in the architecture of the network (explained in [Appendix B.5](#)) contributes proportionally to its L_2 norm to the loss function. This means that the optimiser causes the linear transformations in the neural network to have few singular values in their singular-value decompositions, reducing over-fitting.

B.4 Train-test split

We use a randomised 80 %–20 % train-test split, i.e., 1000 pulses are used for the training of the neural networks and 251 pulses are used for testing the accuracy of predictions, without being seen during training. For the augmentation process, parameter-space densities and normalisations are calculated using only the 1000 training pulses. These parameter-space densities and normalisations are then applied in the augmentation of both the training 1000 pulses and the 251 test pulses.

A prediction of the T_e profile for a pulse is made by using the trained network to make predictions using all the corresponding augmented entries in the dataset. Then, the final T_e profile is computed by averaging over all predictions corresponding to augmented entries.

B.5 Neural-network architecture

Our predictions are carried out using neural networks that consist of several transformations that are applied sequentially. We refer to these transformations as layers, and because of the sequential application of the layers (as per the network architecture shown in [figure 24](#)), these networks are called ‘feed-forward’. Our feed-forward neural-network architecture consists of a sequence of six linear layers, interposed with five rectified-linear-unit (ReLU) layers. If appropriate, dropout layers are also included before the ReLU layers. These will be explained below.

Each linear layer represents an affine transformation of a vector \mathbf{x} , of the form $\mathbf{x} \mapsto W\mathbf{x} + \mathbf{b}$, where W is a weight matrix and \mathbf{b} is a bias vector, both of which are ‘learned’ during training (i.e., they are subject to the optimisation algorithm that reduces the loss and regularisation terms, introduced in [Appendix B.3](#)). The internal dimensions of the affine transformations are set to 64, whereas the input and output dimensions are matched to the profiles and parameter combinations in the labels and features, respectively, as described in [Appendix B.2](#).

The ReLU layers apply a piecewise linear transformation to every component of a vector \mathbf{x} , viz., $\text{ReLU}(x_i) = \max(0, x_i)$. This introduces nonlinearity into the network and allows for complex behaviour. Without such a source of nonlinearity, the network would simply amount to a single affine transformation.

To improve generalisability and stability, particularly in models with small L for local predictions, where overfitting may be more pronounced, we optionally include dropout layers. These layers perform a transformation on a vector \mathbf{x} that randomly chooses a subset of components x_i and sets them to 0 when applied during training. Outside of training, the probability of ‘dropping out’ any x_i is 0, and the network is fully used. This method reduces over-reliance on specific labels and helps prevent overfitting.



FIGURE 24. Feed-forward neural-network architecture described in Appendix B.5. The layers are shown in the order in which they are applied to the labels. The internal dimensions of the linear layers are 64×64 , whereas the first and last layers match the dimensions of the label and feature vectors, respectively. Dropout layers are included only for local-value predictions. The ‘other layers’ block represents a sequence of 64×64 linear layers, dropout layers, and ReLU layers.

Further details regarding the numerical methods described in this appendix can be found in the documentation of the PyTorch package² (Paszke *et al.* 2019), which was used to obtain our results.

Appendix C. Method for integrating temperature profiles

The results obtained in §§ 6.2, 7.2, and 8 began, in general, by stating an explicit or implicit functional relationship between R/L_{T_e} and R/L_{n_e} . Therefore, the starting point of our profile predictions is a differential equation of the form $R/L_{T_e} = f(*)$, or, equivalently,

$$\frac{dT_e}{dR} = -\frac{T_e}{R} f(*), \quad (\text{C1})$$

where the $*$ represents the database parameters and the numerical constants that are used for numerical integration, including R .

Following Simpson *et al.* (2019), the boundary condition for the integration is set as $T_e = 100$ eV at R^{Sep} , the radial location where $\psi_N = 1.0$ (§ 2.1). Thus, we match the starting point of the integration to where the `fit` T_e profile has the value of 100 eV. The numerical integration is then carried out inwards from R^{Sep} , spanning the extent of the pedestal. Once a predicted profile has been computed for each pulse, a new `mtanh` profile is fit to the prediction, and the resulting parameters are compared with the parameters that are contained in the original database, i.e., to the `mtanh` parameters of the `fit` profile.

C.1 Finding best-fit parameters

Determination of the best-fit parameters for a certain scaling, e.g., identifying an optimal η_e in (5.1), is carried out by minimising the absolute difference between the `fit` T_e profile and the reconstructed T_e profile over some radial domain \mathcal{D} . The optimisation and integration are done in real space, but, for the consistency of exposition, we use intervals of ψ_N or X_{Ren} to denote locations. The domain \mathcal{D} can represent a full-pedestal interval encompassing the core slope $\psi_N \in [0.85, 1.0]$, or any region in the pedestals matching some subinterval of $X_{\text{Ren}} \in [0, 2]$. This means that we can find the optimal parameters that match the steep-gradient region, the pedestal region or a larger part of the pedestal.

²The documentation of the PyTorch package can be found at <https://pytorch.org/docs/>.

Minimising the residual between the fit T_e profile and the reconstructed T_e profile,

$$\Delta = \int_{\mathcal{D}} dR [T_{e,\text{fit}}(R) - T_{e,\text{Rec}}(R)]^2 \rightarrow \min \quad (\text{C2})$$

is equivalent to minimising a ‘weighted’ difference between the fit R/L_{T_e} and the reconstructed R/L_{T_e} . This weighting arises because any profile is given by

$$T_e(R) = T_e^{\text{Sep}} \exp\left(-\int_{R^{\text{Sep}}}^R dr \frac{1}{L_{T_e}(r)}\right), \quad (\text{C3})$$

and, therefore, an inaccuracy in R/L_{T_e} close to the separatrix is compounded across the whole $T_{e,\text{Rec}}$ profile. While it is true that if $(R/L_{T_e})_{\text{fit}} = (R/L_{T_e})_{\text{Rec}}$, then $T_{e,\text{fit}} = T_{e,\text{Rec}}$, we aim for a correct prediction of $T_{e,\text{Rec}}$ and, therefore, we allow for more leeway for the prediction of $(R/L_{T_e})_{\text{Rec}}$.

Appendix D. Heat-flux-model optimisation and comparison of T_e reconstructions

In this appendix, we give a general overview of the equations used to obtain the results of § 8 and the outcomes of numerical optimisation of the constants determining these equations. In order to obtain these results, we optimised the free parameters of the transport models listed in Appendix D.1, originally proposed by Chapman-Oplopoiou *et al.* (2022) and Hatch *et al.* (2022). These optimised transport models are used in § 8 to reconstruct T_e profiles. A comprehensive set of optimised parameters, including those from § 8, and their accuracy, are documented in table 3, where we also give a comparison of the quality of predictions performed in §§ 4.1, 4.2, 6.2, and 7.2. The metrics of comparison listed in table 3 are explained in Appendix D.2.

D.1 Explicit expressions for heat-flux models

For all of the heat-flux models described here, we maintain the usage of three free dimensionless parameters: A , $\eta_{e,\text{crit}}$, and β . The two models proposed by Chapman-Oplopoiou *et al.* (2022) are of the form

$$\frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = A \left(\frac{R}{L_{T_e}}\right)^2 (\eta_e - \eta_{e,\text{crit}})^\beta, \quad (\text{D1})$$

where the major distinction between the two is in the value of β . The case with a general value of β is referred to in table 3 as ‘Chap 1 models’, ‘Chap 2 models’ have $\beta = 1$, as originally posited by Guttenfelder *et al.* (2021). Their physical foundation is discussed by Field *et al.* (2023): the positive value of $\eta_{e,\text{crit}}$ in (D1) has the role of a threshold value of η_e , as described in § 8.1.

The five models obtained by fitting to a database of gyrokinetic simulations in table 1 of Hatch *et al.* (2022) (referred to there as Q_1 to Q_5) are of three types. First,

$$\frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = A \frac{R}{L_{T_e}} (\eta_e^\beta - \eta_{e,\text{crit}}), \quad (\text{D2})$$

where $\beta = 2$ for Q_1 (‘Hatch 1’) and $\beta = 4$ for Q_2 and Q_3 (‘Hatch 2’ and ‘Hatch 3’). The distinction between the Hatch 2 and Hatch 3 models is in the constant A and the inclusion of the temperature ratio $\tau = ZT_e/T_i$ as a parameter. Here,

Model		Parameters					Fit metrics					
		$\eta_{e,\text{crit}}$	β	A	M_{Steep}	r_{Steep}^2	M_{Ped}	r_{Ped}^2	M_{Top}	r_{Top}^2	M_{Full}	r_{Full}^2
1	Chap 1 (D1)	1.28	1.43	0.85	0.60	0.55	1.11	0.64	1.88	0.33	237.73	-16.94
2	Chap 2 (D1)	1.40	1.00	1.50	0.49	0.64	1.05	0.66	1.95	0.30	207.90	-14.69
3	Hatch 1 (D2)	1.80	2.00	6.73	16.13	-11.05	28.26	-8.09	43.30	-14.51	2590.88	-194.49
4	Hatch 2 (D2)	-2.88	4.00	0.50	2.35	-0.75	5.50	-0.77	9.47	-2.38	68.59	-4.18
5	Hatch 3 (D2)	-5.12	4.00	0.63	1.14	0.15	2.72	0.13	4.67	-0.66	47.81	-2.61
6	Hatch 5 (D2)	-0.41	4.00	0.31	11.13	-7.32	21.03	-5.76	32.25	-10.55	913.32	-67.91
7	Chap 1 (D1)	1.28	1.43	0.82	0.56	0.58	1.06	0.66	1.81	0.35	1.17	0.91
8	Chap 2 (D1)	1.40	1.00	0.95	0.54	0.60	0.98	0.68	1.67	0.40	0.92	0.93
9	Hatch 1 (D2)	1.80	2.00	116.41	0.81	0.39	1.38	0.56	2.19	0.18	2.46	0.81
10	Hatch 2 (D2)	-2.88	4.00	3.09	0.47	0.65	1.10	0.65	1.97	0.30	2.16	0.84
11	Hatch 3 (D2)	-5.12	4.00	2.10	0.50	0.63	1.14	0.63	2.02	0.28	2.05	0.85
12	Hatch 5 (D2)	-0.41	4.00	6.40	0.74	0.45	1.20	0.61	1.79	0.30	2.39	0.82
13	Chap 1 (D1)	-9.96	3.71	1e-3	0.33	0.75	0.72	0.77	1.36	0.48	0.97	0.93
14	Chap 1 (D1)	-8.60	3.75	1e-3	0.42	0.68	0.70	0.78	1.13	0.57	0.75	0.94
15	Chap 1 (D1)	-0.24	1.44	0.20	0.47	0.65	0.75	0.76	1.19	0.55	0.72	0.95
16	Chap 2 (D1)	-1.15	1.00	0.33	0.34	0.75	0.73	0.77	1.30	0.48	0.85	0.94
17	Chap 2 (D1)	-0.55	1.00	0.29	0.43	0.68	0.71	0.77	1.09	0.57	0.81	0.94
18	Chap 2 (D1)	0.63	1.00	0.50	0.45	0.66	0.75	0.76	1.22	0.54	0.73	0.94
19	Hatch 1 (D2)	-4.98	2.00	50.32	0.44	0.67	1.06	0.66	1.27	0.23	3.30	0.75
20	Hatch 1 (D2)	-4.91	2.00	28.34	0.82	0.39	0.88	0.72	0.70	0.59	1.70	0.87
21	Hatch 1 (D2)	-7.17	2.00	12.64	2.19	-0.63	2.07	0.33	1.18	0.24	0.85	0.94
22	Hatch 1 (D2)	-16.02	2.92	9.64	0.77	0.42	0.86	0.72	0.76	0.61	1.71	0.87
23	Hatch 2 (D2)	-1.04	4.00	12.41	0.44	0.67	1.31	0.58	2.59	0.07	3.64	0.73
24	Hatch 2 (D2)	-1.31	4.00	5.05	0.48	0.64	1.08	0.65	1.92	0.32	2.28	0.83
25	Hatch 2 (D2)	-1.22	4.00	0.54				0.26	0.91	0.60	0.95	
26	Hatch 2 (D2)	-1.73	4.00	1.75	0.92	0.31	1.81	0.42	2.88	-0.03	1.43	0.89
27	Hatch 5 (D2)	-7.12	4.00	0.07	0.44	0.67	1.07	0.66	1.27	0.22	3.31	0.75
28	Hatch 5 (D2)	-6.67	4.00	0.05	0.82	0.38	0.89	0.71	0.71	0.59	1.70	0.87
29	Hatch 5 (D2)	-8.81	4.00	0.01	2.21	-0.65	2.09	0.33	1.10	0.27	0.86	0.94
30	Hatch 5 (D2)	-8.86	4.77	0.01	0.44	0.67	1.06	0.66	1.28	0.22	3.30	0.75
31	§ 6 (5.1)		$\eta_e = \mathbf{1.73}$		0.88	0.34	2.82	0.09	5.28	-0.94	6.76	0.49
32	§ 6 (5.1)		$\eta_e = \mathbf{1.90}$		0.91	0.32	2.67	0.14	5.14	-0.85	5.90	0.56
33	§ 6 (5.1)		$\eta_e = \mathbf{3.10}$					0.53	0.81	1.29	0.90	
34	§ 6 (5.1)		$\eta_e = \mathbf{2.19}$		1.33	0.01	3.47	-0.12	6.47	-1.32	5.27	0.60
35	§ 7 (7.1)	$\alpha = \mathbf{0.32}$, $A = \mathbf{81.21}$			0.45	0.67	1.80	0.42	1.19	-0.04	69.61	-4.25
36	§ 7 (7.1)	$\alpha = \mathbf{0.39}$, $A = \mathbf{50.67}$			0.53	0.60	1.37	0.56	1.13	0.25	7.59	0.43
37	§ 7 (7.1)	$\alpha = \mathbf{0.42}$, $A = \mathbf{29.33}$						0.40	0.81	1.08	0.92	
38	§ 7 (7.1)	$\alpha = \mathbf{0.46}$, $A = \mathbf{31.77}$			0.69	0.48	1.80	0.42	2.66	-0.19	1.81	0.86
39	Local-value None NN				0.83	0.38	1.71	0.45	2.55	0.09	1.81	0.86
40	Local-value No-EFIT NN				0.27	0.80	0.43	0.86	0.55	0.80	0.27	0.98
41	Local-value No-R NN				0.18	0.86	0.29	0.91	0.38	0.86	0.22	0.98
42	Local-value All NN				0.17	0.87	0.25	0.92	0.34	0.88	0.15	0.99
43	Full-pedestal None NN				0.48	0.64	0.79	0.74	1.19	0.58	0.74	0.94
44	Full-pedestal No-EFIT NN				0.22	0.83	0.32	0.90	0.45	0.84	0.21	0.98
45	Full-pedestal EFIT NN				0.16	0.88	0.22	0.93	0.30	0.89	0.20	0.99
46	Full-pedestal All NN				0.12	0.91	0.17	0.94	0.26	0.91	0.11	0.99

TABLE 3. Parameters for the models defined in Appendix D.1 and metrics assessing the quality of T_e reconstructions as detailed in Appendix D.2. Bold values in the ‘Parameters’ columns denote optimised parameters, while, in the ‘Fit metrics’ columns they indicate the corresponding optimisation ranges. Reconstructions using neural-network (NN) models discussed in § 4 are included for comparison. ‘Chap’ is an abbreviation referring to models from Chapman-Oploupiou *et al.* (2022), and ‘Hatch’ to models from Hatch *et al.* (2022).

we take our deuterium plasmas to have the temperature ratio $\tau = 1$ throughout the pedestal profile as an extrapolation from measurements that suggest $T_e = T_i$ (available only for a small subset of the database; see Frassinetti *et al.* 2020). The second model type is

$$\frac{Q_{e,\text{Turb}}}{Q_{e,\text{gB}}} = A \frac{R}{L_{T_e}} (\eta_e - \eta_{e,\text{crit}})^\beta, \quad (\text{D3})$$

where $\beta = 4$ for Q_5 ('Hatch 5'). Finally, Q_4 of Hatch *et al.* (2022) is the $\beta = 1$ version of (D1) ('Chap 2').

We will now assess the relative performance of all these models in predicting correct temperature profiles and also compare them with the models introduced in §§ 6 and 7. In this process, we will also adjust the parameters of (D1) to (D3) for an optimal T_e prediction. All of this information is contained in table 3.

D.2 Comparison of models

All of the numerical optimisations in this study are performed using the method described in Appendix C.1. Given that most of these heat-flux models were inspired by gyrokinetic simulations of a large-gradient and high-shear region of the pedestal, we limit the radial ranges of optimisation to the physically relevant regions:

- the steep-gradient region, 'Steep', bounded by ψ_{Sep} and $\psi_{n_e,\text{Top}}$;
- the full-pedestal region, 'Ped', bounded by ψ_{Sep} and $\psi_{T_e,\text{Top}}$;
- the density-top region, 'Top', bounded by $\psi_{n_e,\text{Top}}$ and $\psi_{T_e,\text{Top}}$.

These definitions are consistent with our approach across this study, and the pedestal locations are as defined in § 2.4. We also include an extended range:

- the pedestal profile including part of the core, 'Full', bounded by ψ_{Sep} and $\psi_N = 0.85$.

Table 3 presents several sets of parameters and metrics of the quality of the pedestal reconstruction corresponding to each set. The reconstructions are carried out using (5.1), (7.1), and (D2) to (D3). Table 3 also includes the quality metrics of neural-network reconstructions carried out in §§ 4.1 and 4.2.

Our metrics of reconstruction quality are all defined over a certain 'Range' \in {Steep, Ped, Top, Full}. The first of them is M_{Range} , the averaged sum of squared discrepancies between the experimental $T_{e,\text{fit}}(R_i)$ and the reconstructed $T_{e,\text{Rec}}(R_i)$, sampled at 100 equally spaced radial locations R_i . Similarly to (C2), it is defined as

$$M_{\text{Range}} = \frac{1}{1251} \sum_{\text{Pulse}} \sum_{i=1}^{100} \left[\frac{T_{e,\text{fit}}(R_i) - T_{e,\text{Rec}}(R_i)}{\max_j T_{e,\text{fit}}(R_j)} \right]^2. \quad (\text{D4})$$

Here M_{Range} is normalised for each pulse to the maximum value of the `fit` profile over this range. The second metric is given the coefficient of determination between $T_{e,\text{fit}}(R_i)$ and $T_{e,\text{Rec}}(R_i)$, defined the usual way, i.e.,

$$r_{\text{Range}}^2 = 1 - \frac{\sum_{\text{Pulse}} \sum_{i=1}^{100} [T_{e,\text{fit}}(R_i) - T_{e,\text{Rec}}(R_i)]^2}{\sum_{\text{Pulse}} \sum_{i=1}^{100} [T_{e,\text{fit}}(R_i) - \langle T_{e,\text{fit}} \rangle]^2}, \quad (\text{D5})$$

where $\langle T_{e,\text{fit}} \rangle = \sum_{\text{Pulse}} \sum_{i=1}^{100} T_{e,\text{fit}}(R_i) / 125100$ is the average value of T_e across all data in the considered 'Range'.

The bold values in the table indicate which parameters are optimised and which metrics correspond to the range over which the optimisation is carried out. Some optimisations do not converge when the exponent β is used as a fit parameter, and they are therefore omitted.

Entries 1–6 in [table 3](#) suggest that the exact scalings from Chapman-Oplolopoiou *et al.* (2022) result in generally passable reconstructions over the pedestal regions, while those from Hatch *et al.* (2022) perform more poorly. Entries 7–12, which have the prefactor A fit over the ‘Ped’ region, result in vast improvements in all metrics for the Hatch models, but, notably, the change in the Chapman 1 model is minimal, whereas the Chapman 2 model requires an adjustment of A of less than 3%. The Hatch 2 and 3 models optimised in this fashion reproduce the ‘Steep’ region better than the Chapman models. It is worth noting that our renditions of the Hatch models in (D2) and (D3) omit a square-root-mass-ratio factor $\sqrt{m_i/m_e}$ compared with the original equations (Q_1) to (Q_5) of table 1 in Hatch *et al.* (2022). However, the adjustment of A that we obtain in entries 9–12 of [table 3](#) cannot be explained by this omission.

Further optimisations (entries 13–30) reveal that all of these models are superior to the constant- η_e and the constant- α reconstructions of §§ 6 and 7 (entries 31–34 for 5.1 and 35–38 for 7.1). The neural-network reconstructions (entries 39–46, presented in [figures 4](#) and [6](#)) perform far better than all other methods, with the exception of the local model that uses no input parameters besides n_e (entry 41 and [figure 6a–d](#)).

A somewhat problematic result of the $\eta_{e,\text{crit}}$ optimisations is that the Chapman models yield $\eta_{e,\text{crit}} < 0$ over the ‘Steep’ and ‘Ped’ regions, clashing with the theoretical picture where $\eta_{e,\text{crit}}$ is either the linear stability threshold or a larger nonlinear one, and should certainly be positive. The values of $\eta_{e,\text{crit}}$ that we obtain for the Hatch models are consistent with this (entries 19–30), and, consequently, all of our optimisations indicate that, in general, terms of the form $f(\eta_e) + C$ are required for the correct modelling of pedestal heat transport, with C a positive constant and f a strictly increasing and positive function.

Finally, all optimisations over the ‘Full’ range result in excellent r^2 coefficients, but the resulting pedestal parameters are very poorly recovered. This conclusion, in conjunction with the other results in [table 3](#), seems to favour mixed models that have different fitting parameters in the qualitatively different regions of the pedestal.

D.3 Sensitivity of ETG-heat-flux models to parameters

In § 8, we describe the results of adopting various heat-flux models motivated by gyrokinetic simulations of turbulence and, to some extent, physical considerations. In particular, we find (8.1) (equivalently, (D1)) to be a promising candidate for representing the turbulent transport in pedestals, giving good pedestal reconstructions. Optimising the parameters of this model yields two interesting results: the exponent β becomes large if we attempt to match the region between $\psi_{T_e,\text{Top}}$ and ψ_{Sep} , and $\eta_{e,\text{crit}}$ takes negative values in this context. The latter is a physically important conclusion, however, the results of these optimisations are also a reflection of the numerical properties of model (8.1).

[Figure 25](#) shows parameter scans for model (8.1). We take the ‘nominal’ parameters to be $\eta_{e,\text{crit}} = 1.28$, $\beta = 1.43$, and $A = 0.85$ (corresponding to entry 1 in [table 3](#)), and vary each of them individually over a range wide enough to encompass the results of all of our optimisations in [Appendix D](#). Alongside $\eta_{e,\text{crit}}$, β , and A , we also vary the implicit parameter T_e^{Sep} (which is everywhere else set to 0.1 KeV) in order to show the effect that it has on the reconstructed T_e profile. Here R^{Sep} is consistently

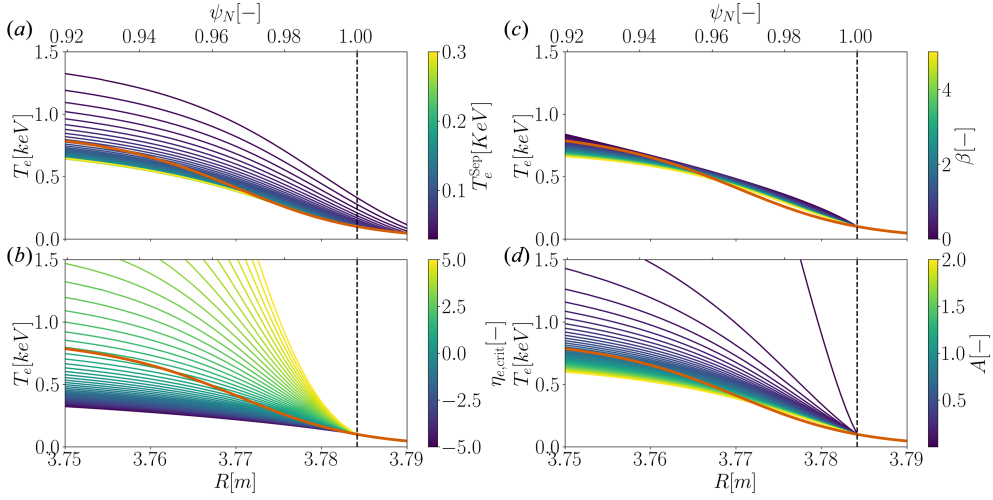


FIGURE 25. The fit T_e profiles for pulse #90339 are shown in orange. Reconstructed profiles using model (D1) are shown in blue–yellow. These reconstructions are integrated from an initial electron temperature T_e^{Sep} . Each panel shows the effect of varying one of the parameters of the model: (a) T_e^{Sep} , (b) $\eta_{e,\text{crit}}$, (c) β , (d) A ; the other parameters are kept fixed at the values given in entry 1 in table 3. The vertical dashed lines represent the conventional separatrix location at $T_e = 0.1$ KeV.

varied with T_e^{Sep} so that the starting point of the integration lies on the fit profile of #90339.

The predictions using models (8.9)–(8.11) are very similar for the whole database. In order to understand the robustness of the prediction, it is therefore important to discuss the sensitivity of the prediction to the exact value of T_e^{Sep} and the values of the heat-flux scaling parameters $\eta_{e,\text{crit}}$, β , and A . Figure 25 displays the results of scans in each of the free parameters. These models are very insensitive to changes of β , and A requires order-of-magnitude variations to impact the resulting profile. Changing A is equivalent to changing the local heat flux (or its gyro-Bohm normalisation), affecting the overall size of R/L_{T_e} . Changes in β slightly adjust the position of the pedestal top, alongside affecting the profile’s behaviour beyond it. High values of β are equivalent to ‘pinning’ R/L_{T_e} to $\eta_{e,\text{crit}} \cdot R/L_{n_e}$, whereas low values of β reduce the importance of $\eta_e - \eta_{e,\text{crit}}$. The profiles in this scan are insensitive to increases of T_e^{Sep} as a result of the fact that, at nominal values of $\eta_{e,\text{crit}}$, β , and A , the prediction approximates the fit profile well. Large reductions of T_e^{Sep} result in large displacements of R^{Sep} , causing high differences in the prediction. The value of $\eta_{e,\text{crit}}$ changes the predictions most significantly. It introduces an effective minimal limit on R/L_{T_e} in order to keep the term $\eta_e - \eta_{e,\text{crit}}$ positive, or, physically, if $\eta_{e,\text{crit}} > 0$, to maintain the turbulence in a supercritical regime.

REFERENCES

- ADKINS, T., IVANOV, P.G. & SCHEKOCIHIN, A.A. 2023 Scale invariance and critical balance in electrostatic drift-kinetic turbulence. *J. Plasma Phys.* **89**, 905890406.
- ADKINS, T., SCHEKOCIHIN, A.A., IVANOV, P.G. & ROACH, C.M. 2022 Electromagnetic instabilities and plasma turbulence driven by electron-temperature gradient. *J. Plasma Phys.* **88**, 905880410.

- APPLEGATE, D.J., ROACH, C.M., CONNOR, J.W., COWLEY, S.C., DORLAND, W., HASTIE, R.J. & JOINER, N. 2007 Micro-tearing modes in the mega ampere spherical tokamak. *Plasma Phys. Control Fusion* **49**, 1113.
- BARNES, M., PARRA, F.I. & SCHEKOCIHIN, A.A. 2011 Critically balanced ion temperature gradient turbulence in fusion plasmas. *Phys. Rev. Lett.* **107**, 115003.
- CHAPMAN-OPLOPOIOU, B., *et al.* 2022 The role of ETG modes in JET-ILW pedestals with varying levels of power and fuelling. *Nucl. Fusion* **62**, 086028.
- CHAPMAN-OPLOPOIOU, B., WALKER, J., HATCH, D.R., GORLER, T. & CONTRIBUTORS, J.E.T. 2025 Composition of electron temperature gradient driven plasma turbulence in JET-ILW tokamak plasmas. *Phys. Rev. Res.* **7**, L012004.
- COLYER, G.J., SCHEKOCIHIN, A.A., PARRA, F.I., ROACH, C.M., BARNES, M.A., GHIM, Y.-C. & DORLAND, W. 2017 Collisionality scaling of the electron heat flux in ETG turbulence. *Plasma Phys. Control. Fusion* **59**, 055002.
- CONNOR, J.W., KIRK, A. & WILSON, H.R. 2008 Edge localised modes (ELMs): experiments and theory. *AIP Conf. Proc.* **1013**, 174.
- CRANMER, M. 2023 Interpretable machine learning for science with PySR and SymbolicRegression.jl. arXiv:2305.01582
- DICORATO, M., *et al.* 2024 Turbulent transport mechanisms and their impact on the pedestal top of JET plasmas with small-ELMs. *Plasma Phys. Control. Fusion* **66**, 125002.
- DIMITS, A.M., COHEN, B.I., MATTOR, N., NEVINS, W.M., SHUMAKER, D.E., PARKER, S.E. & KIM, C. 2000 Simulation of ion temperature gradient turbulence in tokamaks. *Nucl. Fusion* **40**, 661.
- DRAKE, J.F., GLADD, N.T., LIU, C.S. & CHANG, C.L. 1980 Microtearing modes and anomalous transport in tokamaks. *Phys. Rev. Lett.* **44**, 994.
- DRAKE, J.F. & LEE, Y.C. 1977 Nonlinear evolution of collisionless and semicollisional tearing modes. *Phys. Rev. Lett.* **39**, 453.
- DUNNE, M.G., *et al.* 2016 The role of the density profile in the ASDEX upgrade pedestal structure. *Plasma Phys. Control. Fusion* **59**, 014017.
- FARÇAŞ, I.-G., MERLO, G. & JENKO, F. 2024 Advanced surrogate model for electron-scale turbulence in tokamak pedestals. *J. Plasma Phys.* **90**, 905900510.
- FARENGO, R., GUZDAR, P.N. & LEE, Y.C. 1989 Collisionless electron temperature gradient-driven instability in field-reversed configurations. *Phys. Fluids B* **1**, 2181.
- FIELD, A.R., *et al.* 2020 The dependence of exhaust power components on edge gradients in JET-C and JET-ILW H-mode plasmas. *Plasma Phys. Control. Fusion* **62**, 055010.
- FIELD, A.R., CHAPMAN-OPLOPOIOU, B., CONNOR, J., FRASSINETTI, L., HATCH, D., ROACH, C. & SAARELMA, S. 2023 Comparing pedestal structure in JET-ILW H-mode plasmas with a model for stiff ETG turbulent heat transport. *Phil. Trans. A R. Soc. Lond.* **381**, 20210228.
- FRASSINETTI, L., BEURSKENS, M.N.A., SCANNELL, R., OSBORNE, T.H., FLANAGAN, J., KEMPENAARS, M., MASLOV, M., PASQUALOTTO, R., WALSH, M. & JET-EFDA Contributors 2012 Spatial resolution of the JET Thomson scattering system. *Rev. Sci. Instr.* **83**, 013506.
- FRASSINETTI, L., *et al.* 2020 Pedestal structure, stability and scalings in JET-ILW: the EUROfusion JET-ILW pedestal database. *Nucl. Fusion* **61**, 016001.
- FRASSINETTI, L., *et al.* 2021 Role of the separatrix density in the pedestal performance in deuterium low triangularity JET-ILW plasmas and comparison with JET-C. *Nucl. Fusion* **61**, 126054.
- GARBET, X., *et al.* 2004 Physics of transport in tokamaks. *Plasma Phys. Control. Fusion* **46**, B557.
- GIACOMIN, M., DICKINSON, D., KENNEDY, D., PATEL, B.S. & ROACH, C.M. 2023 Nonlinear microtearing modes in MAST and their stochastic layer formation. *Plasma Phys. Control. Fusion* **65**, 095019.
- GUTTENFELDER, W., *et al.* 2012 Simulation of microtearing turbulence in national spherical torus experiment. *Phys. Plasmas* **19**, 056119.
- GUTTENFELDER, W., GROEBNER, R.J., CANIK, J.M., GRIERSON, B.A., BELLI, E.A. & CANDY, J. 2021 Testing predictions of electron scale turbulent pedestal transport in two DIII-D ELMy H-modes. *Nucl. Fusion* **61**, 056005.
- HAHM, T.S. & BURRELL, K.H. 1995 Flow shear induced fluctuation suppression in finite aspect ratio shaped tokamak plasma. *Phys. Plasmas* **2**, 1648.

- HATCH, D.R., KOTSCHENREUTHER, M., MAHAJAN, S., VALANJU, P., JENKO, F., TOLD, D., GÖRLER, T. & SAARELMA, S. 2016 Microtearing turbulence limiting the JET-ILW pedestal. *Nucl. Fusion* **56**, 104003.
- HATCH, D.R., KOTSCHENREUTHER, M., MAHAJAN, S., VALANJU, P. & LIU, X. 2017 A gyrokinetic perspective on the JET-ILW pedestal. *Nucl. Fusion* **57**, 036020.
- HATCH, D.R., *et al.* 2019 Direct gyrokinetic comparison of pedestal transport in JET with carbon and ITER-like walls. *Nucl. Fusion* **59**, 086056.
- HATCH, D.R., *et al.* 2021 Microtearing modes as the source of magnetic fluctuations in the JET pedestal. *Nucl. Fusion* **61**, 036015.
- HATCH, D.R., KOTSCHENREUTHER, M.T., LI, P.-Y., CHAPMAN-OPLOPOIOU, B., PARISI, J., MAHAJAN, S.M. & GROEBNER, R. 2024 Modeling electron temperature profiles in the pedestal with simple formulas for ETG transport. *Nucl. Fusion* **64**, 066007.
- HATCH, D.R., *et al.* 2022 Reduced models for ETG transport in the tokamak pedestal. *Phys. Plasmas* **29**, 062501.
- HATCH, D.R., TOLD, D., JENKO, F., DOERK, H., DUNNE, M.G., WOLFRUM, E., VIEZZER, E., PUESCHEL, M.J. & The ASDEX Upgrade Team 2015 Gyrokinetic study of ASDEX Upgrade inter-ELM pedestal profile evolution. *Nucl. Fusion* **55**, 063028.
- HORTON, L.D., *et al.* 2005 Characterization of the H-mode edge barrier at ASDEX Upgrade. *Nucl. Fusion* **45**, 856.
- HORTON, W., HONG, B.G. & TANG, W.M. 1988 Toroidal electron temperature gradient driven drift modes. *Phys. Fluids B* **31**, 2971.
- HUYSMANS, G.T.A., *et al.* 1999 MHD stability of optimized shear discharges in JET. *Nucl. Fusion* **39**, 1489.
- IMBEAUX, F., *et al.* 2015 Design and first applications of the ITER integrated modelling & analysis suite. *Nucl. Fusion* **55**, 123006.
- IVANOV, P.G., SCHEKOCIHIN, A.A. & DORLAND, W. 2022 Dimits transition in three-dimensional ion-temperature-gradient turbulence. *J. Plasma Phys.* **88**, 905880506.
- JENKO, F., DORLAND, W. & HAMMETT, G.W. 2001 Critical gradient formula for toroidal electron temperature gradient modes. *Phys. Plasmas* **8**, 4096.
- JENKO, F., DORLAND, W., KOTSCHENREUTHER, M. & ROGERS, B.N. 2000 Electron temperature gradient driven turbulence. *Phys. Plasmas* **7**, 1904.
- JENKO, F., TOLD, D., XANTHOPOULOS, P., MERZ, F. & HORTON, L.D. 2009 Gyrokinetic turbulence under near-separatrix or nonaxisymmetric conditions. *Phys. Plasmas* **16**, 055901.
- JOFFRIN, E.H., WISCHMEIER, M., BARUZZO, M., HAKOLA, A., KAPPATOU, A., KEELING, D., LABIT, B., TSITRONE, E. & VIANELLO, N., *et al.* 2024 Overview of the EUROfusion tokamak exploitation programme in support of ITER and DEMO. *Nucl. Fusion* **64**, 112019.
- KIT, A., JÄRVINEN, A., FRASSINETTI, L. & WIESEN, S. 2023 Supervised learning approaches to modeling the pedestal density. *Plasma Phys. Control. Fusion* **65**, 045003.
- KOTSCHENREUTHER, M., HATCH, D.R., MAHAJAN, S., VALANJU, P., ZHENG, L. & LIU, X. 2017 Pedestal transport in H-mode plasmas for fusion gain. *Nucl. Fusion* **57**, 064001.
- KOTSCHENREUTHER, M., *et al.* 2019 Gyrokinetic analysis and simulation of pedestals to identify the culprits for energy losses using ‘fingerprints’. *Nucl. Fusion* **59**, 096001.
- KRUTKIN, O., CHAPMAN-OPLOPOIOU, B., FRASSINETTI, L., BRUNNER, S., CODA, S. & LABIT, B. 2025 Modelling of ETG turbulent transport in the TCv pedestal. *Plasma Phys. Control. Fusion* **67**, 025029.
- LAO, L.L., JOHN, H.ST, STAMBAUGH, R.D., KELLMAN, A.G. & PFEIFFER, W. 1985 Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nucl. Fusion* **25**, 1611.
- LARAKERS, J.L., CURIE, M., HATCH, D.R., HAZELTINE, R.D. & MAHAJAN, S.M. 2021 Global theory of microtearing modes in the tokamak pedestal. *Phys. Rev. Lett.* **126**, 225001.
- LARAKERS, J.L., HAZELTINE, R.D. & MAHAJAN, S.M. 2020 A comprehensive conductivity model for drift and micro-tearing modes. *Phys. Plasmas* **27**, 062503.
- LEPPIN, L.A., GÖRLER, T., CAVEDON, M., DUNNE, M.G., WOLFRUM, E. & JENKO, F. the ASDEX Upgrade Team 2023 Complex structure of turbulence across the ASDEX upgrade pedestal. *J. Plasma Phys.* **89**, 905890605.

- LEPPIN, L.A., GÖRLER, T., FRASSINETTI, L., SAARELMA, S., HOBIRK, J., JENKO, F. & CONTRIBUTORS, J.E.T. 2024 The JET hybrid H-mode scenario from a pedestal turbulence perspective. arXiv:2405.10668.
- MAILLOUX, J., *et al.* 2022 Overview of JET results for optimising ITER operation. *Nucl. Fusion* **62**, 042026.
- MANTICA, P., *et al.* 2009 Experimental study of the ion critical-gradient length and stiffness level and the impact of rotation in the JET Tokamak. *Phys. Rev. Lett.* **102**, 175002.
- MILLER, R.L., CHU, M.S., GREENE, J.M., LIN-LIU, Y.R. & WALTZ, R.E. 1998 Noncircular, finite aspect ratio, local equilibrium model. *Phys. Plasmas* **5**, 973.
- MURARI, A., PISANO, F., VEGA, J., CANNAS, B., FANNI, A., GONZALEZ, S., GELFUSA, M., GROSSO, M. & EFDA-JET Contributors 2014 Extensive statistical analysis of ELMs on JET with a carbon wall. *Plasma Phys. Control. Fusion* **56**, 114007.
- NEUHAUSER, J., *et al.* Transport into and across the scrape-off layer in the ASDEX Upgrade divertor tokamak. *Plasma Phys. Control. Fusion* **44**, 855.
- NIES, R., PARRA, F., BARNES, M., MANDELL, N. & DORLAND, W. 2024 Saturation of magnetised plasma turbulence by propagating zonal flows. arXiv:2409.02283
- PARISI, J.F., HATCH, D.R., LI, P.Y., BERKERY, J.W., NELSON, A.O., KAYE, S.M., IMADA, K. & LAMPERT, M. 2025. Turbulent transport-limited pedestals in tokamaks. arXiv:2505.09101.
- PARISI, J.F., *et al.* 2020 Toroidal and slab ETG instability dominance in the linear spectrum of JET-ILW pedestals. *Nucl. Fusion* **60**, 126045.
- PARISI, J.F., *et al.* 2022 Three-dimensional inhomogeneity of electron-temperature-gradient turbulence in the edge of tokamak plasmas. *Nucl. Fusion* **62**, 086045.
- PASQUALOTTO, R., NIELSEN, P., GOWERS, C., BEURSKENS, M., KEMPENAARS, M., CARLSTROM, T. & JOHNSON, D.JET-EFDA Contributors 2004 High resolution Thomson scattering for Joint European Torus (JET). *Rev. Sci. Instr.* **75**, 3891.
- PASZKE, A., *et al.* 2019 PyTorch: An imperative style, high-performance deep learning library. arXiv:1912.01703.
- PREDEBON, I., HATCH, D.R., FRASSINETTI, L., HORVATH, L., SAARELMA, S., CHAPMAN-OPLOPOIOU, B., GÖRLER, T. & MAGGI, C.F. 2023 Isotope mass dependence of pedestal transport in JET H-mode plasmas. *Nucl. Fusion* **63**, 036010.
- PUESCHEL, M.J., KAMMERER, M. & JENKO, F. 2008 Gyrokinetic turbulence simulations at high plasma beta. *Phys. Plasmas* **15**, 102310.
- REN, Y., GUTTENFELDER, W., KAYE, S.M. & WANG, W.X. 2024 Transport from electron-scale turbulence in toroidal magnetic confinement devices. *Rev. Mod. Plasma Phys.* **8**, 5.
- SAARELMA, S., *et al.* 2024 Density pedestal prediction model for tokamak plasmas. *Nucl. Fusion*, **64**, 076025.
- SAARELMA, S., CONNOR, J.W., BILKOVA, P., BOHM, P., FIELD, A.R., FRASSINETTI, L., FRIDSTROM, R., KIRK, A. & JET Contributors 2023 Testing a prediction model for the H-mode density pedestal against JET-ILW pedestals. *Nucl. Fusion* **63**, 052002.
- SAARELMA, S., FRASSINETTI, L., BILKOVA, P., CHALLIS, C.D., CHANKIN, A., FRIDSTRÖM, R., GARZOTTI, L., HORVATH, L., MAGGI, C.F. & CONTRIBUTORS, J.E.T. 2019 Self-consistent pedestal prediction for JET-ILW in preparation of the DT campaign. *Phys. Plasmas* **26**, 072501.
- SIMPSON, J., MOULTON, D., GIROUD, C., GROTH, M. & CORRIGAN, G. 2019 Using EDGE2D-EIRENE to simulate the effect of impurity seeding and fueling on the upstream electron separatrix temperature. *Nucl. Mater. Energy* **20**, 100599.
- SNYDER, P.B. 1999 Gyrofluid theory and simulation of electromagnetic turbulence and transport in tokamak plasmas. PhD thesis, 1999.9944640.
- SNYDER, P.B., GROEBNER, R.J., LEONARD, A.W., OSBORNE, T.H. & WILSON, H.R. 2009 Development and validation of a predictive model for the pedestal height. *Phys. Plasmas* **16**, 056118.
- SNYDER, P.B. & HAMMETT, G.W. 2001 A Landau fluid model for electromagnetic plasma microturbulence. *Phys. Plasmas* **8**, 3199.
- SNYDER, P.B., WILSON, H.R., FERRON, J.R., LAO, L.L., LEONARD, A.W., MOSSESIAN, D., MURAKAMI, M., OSBORNE, T.H., TURNBULL, A.D. & XU, X.Q. 2004 ELMs and constraints on the H-mode pedestal: peeling–ballooning stability calculation and comparison with experiment. *Nucl. Fusion* **44**, 320.

- SOLANO, E.R., *et al.* Recent progress in L–H transition studies at JET: tritium, helium, hydrogen and deuterium. *Nucl. Fusion* **62**, 076026.
- STANGEBY, P.C. 2000 *The Plasma Boundary of Magnetic Fusion Devices*, vol. 10, p. 1201. CRC Press.
- STEFANIKOVA, E., *et al.* 2018 Effect of the relative shift between the electron density and temperature pedestal position on the pedestal stability in JET-ILW and comparison with JET-C. *Nucl. Fusion* **58**, 056010.
- TANG, W.M., CONNOR, J.W. & HASTIE, R.J. 1980 Kinetic-ballooning-mode theory in general geometry. *Nucl. Fusion* **20**, 1439.
- TERRY, P.W. 2000 Suppression of turbulence and transport by sheared flow. *Rev. Mod. Phys.* **72**, 109.
- TOLD, D., JENKO, F., XANTHOPOULOS, P., HORTON, L.D. & WOLFRUM, E. 2008 Gyrokinetic microinstabilities in ASDEX Upgrade edge plasmas. *Phys. Plasmas* **15**, 102306.
- TRINCZEK, S., PARRA, F.I., CATTO, P.J., CALVO, I. & LANDREMAN, M. 2023 Neoclassical transport in strong gradient regions of large aspect ratio tokamaks. *J. Plasma Phys.* **89**, 905890304.
- WAGNER, F., *et al.* 1982 Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak. *Phys. Rev. Lett.* **49**, 1408.
- WALTZ, R.E. 2010 Nonlinear subcritical magnetohydrodynamic beta limit. *Phys. Plasmas* **17**, 072501.
- WANG, H.Q. *et al.* 2018 Effects of divertor geometry on H-mode pedestal structure in attached and detached plasmas in the DIII-D tokamak. *Nucl. Fusion* **58**, 096014.
- WILSON, H.R., CONNOR, J.W., FIELD, A.R., FIELDING, S.J., MILLER, R.L., LAO, L.L., FERRON, J.R. & TURNBULL, A.D. 1999 Ideal magnetohydrodynamic stability of the tokamak high-confinement-mode edge region. *Phys. Plasmas* **6**, 1925.
- ZHANG, Y.Z. & MAHAJAN, S.M. 1992 Edge turbulence scaling with shear flow. *Phys. Fluids B* **4**, 1385.
- ZOCCO, A., LOUREIRO, N.F., DICKINSON, D., NUMATA, R. & ROACH, C.M. 2015 Kinetic microtearing modes and reconnecting modes in strongly magnetised slab plasmas. *Plasma Phys. Control. Fusion* **57**, 065008.