# Time Series Extrinsic Regression of Plasma Parameters For a Candidate Fusion Reactor Diagnostic

Ethan Attwood[1] (✉), J.W.S. Cook[2,3], and Peter Hill[1]

[1] York Plasma Institute, University of York, York, UK {ethan.attwood@york.ac.uk}
[2] UK Atomic Energy Authority, Culham Campus, Culham, UK
[3] Centre for Fusion, Space and Astrophysics, University of Warwick, Coventry, UK

**Abstract.** The most promising classes of fusion reactors use magnetic fields to confine plasmas many times hotter than the core of the Sun to achieve net electricity yield. These extreme conditions require constant monitoring due to the potential for catastrophic damage to reactor components. Monitoring systems currently occupy an unfeasible amount of reactor wall space for next-generation devices, which require this area for the power cycle and self-sustaining generation of fuel. Therefore, more information-efficient diagnostics are required, of which one candidate is Ion Cyclotron Emission (ICE) spectra. ICE is a ubiquitous phenomenon in magnetically confined hot plasmas, shown to correlate with fusion reaction rate and the properties of energetic ions, and has been detected using common hardware. While the corpus of experimental data is limited, ICE spectra can be synthesised using Particle-In-Cell (PIC) plasma simulation codes. PIC codes are computationally intensive, so maximising the information extracted from small parameter sets is critical. In this work, scans over experimentally-relevant ranges of magnetic field strengths, plasma densities, and energetic ion velocity characteristics are used to build a versatile dataset of ICE spectra. State-of-the-art time series extrinsic regression (TSER) algorithms from Aeon are then trained to solve the inverse problem of extracting underlying plasma properties from ICE spectra. We demonstrate the capabilities of TSER models to learn information never previously recovered from ICE spectra, providing a route to a novel diagnostic and a proof-of-concept for applying time-series approaches to the enormous wealth of sequential data in physics.

**Keywords:** plasma physics · plasma diagnostics · simulation · particle-in-cell · alpha particles · fast ions · ion cyclotron emission · nuclear fusion · fusion energy · machine learning · time series extrinsic regression

## Introduction

Magnetically confined fusion devices require a network of diagnostics to determine plasma properties critical to safe and effective operation. Next-generation devices known as tokamaks are expected to use a large majority of reactor wall space for a lithium blanket used to breed tritium, a heavy isotope of hydrogen projected to form up to 50% of fusion fuel [1]. However, nearly 20% of the outboard vessel area of the international fusion engineering project ITER is currently used for diagnostics, damaging the tritium-breeding ratio required for the transition from research-scale to grid-scale fusion devices [45]. Diagnostic solutions demonstrating a high technical readiness and providing data on a wide variety of plasma characteristics using as few ports as possible are more important than ever.

Ion Cyclotron Emission (ICE) is a widespread plasma phenomenon, occurring in virtually all large hot magnetized plasmas [13,19,31,36] and experimentally observed as electromagnetic emission in the radio frequency range with spectral peaks at harmonics of the ion cyclotron frequency (proportional to the ionic species and strength of the magnetic field). Typically, ICE requires a minority fast ion population in cyclotron resonance with an underlying motion of the majority, slower ion population: the fast Alfvén wave. When this resonance propagates approximately perpendicular to the background magnetic field, it gives rise to the magnetoacoustic cyclotron instability (MCI), the leading mechanism underpinning ICE [18,7,9,40]. An abundant source of fast ions in a tokamak is the fusion reaction itself, thereby correlating ICE observations with fusion reaction rate as well as with the magnetic field strength and the spatial and velocity distribution of energetic ion populations via the finer details of the kinetic instability [13,32]. Two methods of plasma heating, ion cyclotron resonance [31] and neutral beam injection [38], are also known to excite ICE.

The Particle-In-Cell (PIC) method [26,5,2] is a common technique in computational plasma physics for self-consistently solving Maxwell's equations, the fundamental relations of classical electromagnetism. PIC evaluates electric and magnetic fields on a grid simultaneously with currents and charge densities provided by plasma macroparticles, which are then subject to the resultant forces governing their motion. PIC is ideal for computationally efficient models of nonlinear kinetic phenomena evolving at very fast timescales such as ion gyroperiods, and has been widely applied to ICE investigations [12,7,38,8]. Resolution constraints placed on the number and size of grid cells, particle numbers and simulation durations make large parameter scans with PIC computationally expensive albeit tractable. Reasonable simplifications can be made such as lowering dimensionality, which we apply in this work by rendering our problem in so-called 1D3V: one spatial dimension of variation is allowed while vector fields, including velocities, retain three components. In this paper, we exploit a fully kinetic approach, i.e. all species are represented with discrete macroparticles. This maximises fidelity of the physical interactions of particles and waves crucial to accurately modelling the MCI [11]. In order to include some two dimensional physics, we run four 1D3V simulations for each set of parameters: one perpendicular to the magnetic field, 90°, and three at oblique angles, 92°, 94° and 96°. This retains one dimensional nonlinear physics and introduces linear two dimensional physics at a substantially lower computational cost than 2D3V simulations. Experimental ICE measurements will tend to reflect the nonlinear regime of the MCI due to the development of complex wave-wave interactions past an initial single-digit number of ion gyroperiods [7] as well as the intricacies of the machine geometry, wave propagation, and sink and source terms such as further fusion reactions, neutralisation and exiting of the bulk plasma. Fortunately, analytic approaches and periodic 1D3V PIC code simulations that mimic an infinite homogeneous plasma have been able to explain many aspects of ICE, including how the spectral features that arise during the linear phase of the instability leave their imprint on the nonlinear spectrum [12,7,38,8]. Nonlinearity, enabled by PIC, is required in order to account for the magnitude of energy transfer from fast ions to the background and for wave-wave coupling of power from higher to lower harmonics[10] and hence is essential for comparing simulation results to experimental observations.

While simulations can supplement or replace the body of data available from plasma experiments, there are no comprehensive relationships from theory linking spectral features to useful plasma quantities, so interpretation requires machine learning to solve the inverse problem of recovering the parameters used to generate a spectrum. This can be investigated through traditional regression models such as decision tree ensembles, or through more specialised time-series approaches such as Time-Series Extrinsic Regression (TSER): the prediction of a continuous scalar value from an ordered sequence [41]. TSER exploits knowledge of the relationship between ordered datapoints such as their relative amplitudes, occurrence of repeated patterns, rolling forecasts and seasonality. It is differentiated from the related field of forecasting by predicting an external variable underlying the time-series properties rather than predicting its continuation [41,24].

## Methods

For this study we used EPOCH, a well-established PIC code in plasma literature [2]. We performed parameter scans to form ICE-relevant datasets, on which machine learning models were trained to infer meaningful physical properties for fusion scientists and engineers. Parameter sets can be designed specifically for individual magnetised fusion devices, a natural first step in the development of synthetic diagnostics. We chose parameters proximal to those used in the Joint European Torus (JET) due its provision of ICE observations from tritium experiments with a significant proportion of fusion alpha-particles ($He^{2+}$) which, as opposed to lower energy externally-heated ions, feature increased ICE intensity for this first of a kind study.

The dataset was exploited by regressing spectra against the input simulation parameters. ICE signals derived from simulations are power spectra calculated by Fourier transforming temporal field data, which we call synthetic spectra as we use them as analogues for experimentally observed spectra. We trained models capable of determining the following four quantities from synthetic ICE spectra: the equilibrium magnetic field strength ($B_0$); electron density ($n_e$); concentration of fusion-born alpha-particles ($n_\alpha/n_e$); and alpha-particle velocity distribution pitch $\lambda = u_{\|}/\sqrt{(u_{\|}^2 + u_{\perp}^2)}$, defined as the ratio of parallel bulk speed to the initial speed, which determines their distribution function in Equation 1 (parallel and perpendicular in this case indicating orientation relative to the background magnetic field). The latter quantity is significant as it is predominantly the particle's perpendicular velocity component that contributes to excitation of the MCI.

In our simulations, three species were modelled: electrons and background fuel deuterons ($^2$H), both initially at 10 kiloelectronvolts (keV; approximately $1\times10^8$ degrees K), and a minority population of energetic alpha-particles defined by the distribution function,

$$f(v_{\|}, v_{\perp}) \propto \exp\left(-\frac{(v_{\|} - u_{\|})^2}{v_{th,\|}^2} - \frac{(v_{\perp} - u_{\perp})^2}{v_{th,\perp}^2}\right), \tag{1}$$

which initialises velocities according to bulk speeds parallel and perpendicular to the magnetic field direction, $u_{\|}$ and $u_{\perp}$ respectively, and thermal speeds $v_{th,\|}$ and $v_{th,\perp}$ following the well-known Maxwell-Boltzmann distribution for an ideal gas. This distribution of particle velocities, known as a drifting "ring-beam", is frequently used in ICE studies to replicate experimental observations to a high degree of agreement [8]. A homogeneous magnetic field was initialised predominantly in $z$, perpendicular to the spatial domain $x$, such that the MCI is most strongly excited [22].

To generate the datasets, the above four quantities were sampled using a Latin hypercube strategy, minimum and maximum values for which are shown in Table 1. Electron density and alpha-particle fraction were sampled uniformly in log space. Static parameters used in all simulations are listed in Table 2. Initially 100 parameter sets were simulated; later 10 more were added using another Latin hypercube near to the known parameter values of JET pulse #26148 ($B_0 = 2-3$T, $\lambda = 0.3-0.7$, $n_e = (1-2)\times10^{19}$m$^{-3}$, $n_\alpha/n_e = 10^{-4}-10^{-3}$) as documented in Ref. [13]. Reference values for other major fusion experiments are shown in the rightmost column for $B_0$ and $n_e$ in Table 1. As ICE originates from a variety of fast ion sources in both the tokamak core[36] and edge plasma regions [8], there are many valid operational combinations of parameters such as density and species temperature. This dataset trades specificity to one particular regime for a versatile training set representative of a broad sample of ICE origins.

All simulations were run on ARCHER2 [4], the UK Tier 1 HPC resource, with 250 macroparticles per species per cell for 16 deuteron gyroperiods, with a diagnostic output rate allowing resolution of data in frequency space up to 80 deuteron gyrofrequencies, exceeding the frequency range of common experimental ICE spectra[13,38]. EPOCH default settings were used except for the enabling of $\delta f$, a technique used to suppress numerical noise [2,28], and using higher-order particle shape functions which more realistically model the diffusion of charge through space [2]. A comparison of a subset of particle and field energy diagnostics with and without using $\delta f$ for the background deuterons and electrons showed no appreciable impact on the rate or magnitude of MCI evolution. A full-$f$ treatment was preserved for the fast alpha population. Running

**Table 1.** Simulation parameter scan variable ranges, means and standard deviations (SD). Electron density ($n_e$) and alpha-particle fraction ($n_\alpha/n_e$) were sampled uniformly in $\log_{10}$ space. The rightmost columns show typical values of $B$ (toroidal, on-axis) and $n_e$ in previous D-T experiments on JET [23], or projected for ITER [39] and STEP [43]. Densities are averages across major radius and operational configurations. Minimum and maximum values are shown at the precision to which they were specified; mean and SD's were calculated post-sampling and shown to 3 significant figures.

| Parameter | Minimum | Maximum | Mean | Standard Deviation | JET | ITER | STEP |
|---|---|---|---|---|---|---|---|
| $B_0$ [T] | 1.0 | 5.0 | 2.95 | 1.12 | 3.4 | 5.3 | 3.2 |
| $\lambda$ | 0.01 | 0.99 | 0.499 | 0.273 | – | – | – |
| $n_e$ [$\log_{10}(\mathrm{m}^{-3})$] | 19 | 20 | 19.6 | 19.4 | 19.7 | 19.5 | 20.0 |
| $n_\alpha/n_e$ [$\log_{10}$] | $-4$ | $-2$ | $-2.70$ | $-2.62$ | – | – | – |

**Table 2.** Static parameters used in all simulations. Times and frequencies were normalised to alpha-particle gyroperiod ($\tau_{c,\alpha}$) and gyrofrequency ($\Omega_{c,\alpha}$) respectively.

| Symbol | Description | Value |
|---|---|---|
| $T_{D,e}$ | background temperature [eV] | 10000 |
| $E_\alpha$ | alpha-particle energy [eV] | 3500000 |
| $p_\alpha$ | alpha-particle momentum [kgms$^{-1}$] | $\sqrt{2m_\alpha q_e E_\alpha}$ |
| $p_{\mathrm{beam}}$ | parallel momentum ($u_\parallel m_\alpha$) [kgms$^{-1}$] | $p_\alpha\lambda$ |
| $p_{\mathrm{ring}}$ | perpendicular momentum ($u_\perp m_\alpha$) [kgms$^{-1}$] | $p_\alpha\sqrt{1-\lambda^2}$ |
| $p_{\mathrm{spread}}$ | momentum spread ($v_{th} m_\alpha$) [kgms$^{-1}$] | $0.06p_{\mathrm{ring}}$ |
| $t_{\mathrm{sim}}$ | simulation time [$\tau_{c,\alpha}$] | 16 |
| $N_\Omega$ | Nyquist frequency [$\Omega_{c,\alpha}$] | 80 |
| $n_{\mathrm{samples}}$ | number of time samples | $2t_{\mathrm{sim}}N_\Omega$ |
| $n_{\mathrm{ppc}}$ | number of particles per species per cell | 250 |

all simulations combined took approximately 1.55M CPU hours. Analysis of the outputs common to all simulations took approximately 600 CPU hours split across ARCHER2 and Viking, the University of York's compute cluster.

The experimentally observed power spectra considered here and our synthetic spectra are a function of frequency only; all wavenumber (the reciprocal of wavelength) information is lost. The actions of plasma instabilities are visualised in electromagnetic field perturbations; in this case the parallel magnetic field $\delta B_z$, electrostatic component $\delta E_x$ and Alfvénic electric component $\delta E_y$. These power spectra are integrated over one-dimensional wavenumbers and summed across all four propagation angles of our 110 cases. A representative spectrum of the compressional magnetic field component $\delta B_z$, which provides the principle analogue for the MCI as measured by a magnetic field probe, is shown in Figure 1 and illustrates that the dominant spectral features can arise from different propagation angles. Trends in peak grouping, splitting, and amplitude do not vary straightforwardly with the parameters or with variation in propagation angle. The wave dispersion caused by the MCI is controlled largely by density and

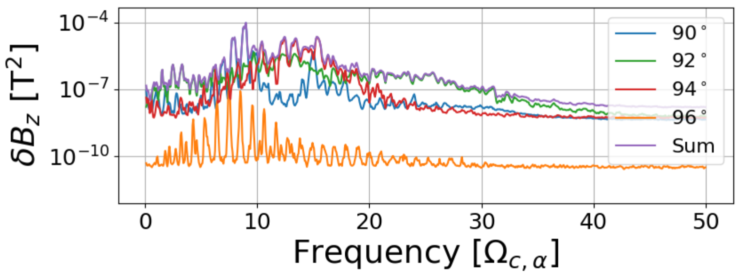


**Fig. 1.** Frequency power spectra in $\delta B_z$ for four propagation angles of parameter set 75. Successive peaks are seen at integer harmonics of the deuteron/alpha cyclotron frequency, with peaks at $8-16\Omega_{c,\alpha}$ (frequency normalised to the alpha-particle cyclotron frequency), indicative of MCI activity. All four angles share a region of activity below $12\Omega_{c,\alpha}$. The perpendicular case, $90^\circ$ shows two distinct regions of excitation around $8\Omega_{c,\alpha}$ and $15\Omega_{c,\alpha}$ and the $96^\circ$ spectral amplitudes are less powerful by a factor of 1000. The parameters are $B_0 = 2.0$T , $\lambda = 0.49$, $n_e = 1.4 \times 10^{19}$ and $n_\alpha = 2.8 \times 10^{-3} n_e$.

magnetic field strength. However, the complicated and sometimes disconnected regions of instability are determined further by factors arising from the fast ion velocity distribution and the energy exchange between electrons, fuel ions and waves [11]. The mapping from spectra to these physics parameters includes details of these complex linear and nonlinear physical processes, which necessarily must be encapsulated by well-performing regression models.

To mitigate against overfitting and give a representative assessment of regression performance, a leave-one-out cross-validation (LOOCV) strategy was employed, where only one test case is removed from data preparation and training. This maximises utility of the relatively small set of computationally expensive nonlinear simulations. Error metrics are then evaluated over all LOOCV folds to generate holistic measures of model performance. To provide a baseline regression expectation, the problem was applied to four traditional regression algorithms available in scikit-learn [37]: two linear models, Ridge (RR) [27] and Lasso (LR) [44], and two decision tree ensembles, the Random Forest (RFR) [6] and Gradient Boosting (GBR) [21] Regressors. As these algorithms do not natively handle multivariate vector input, the three EM spectra and their frequency ranges were concatenated and passed as flattened spectra of shape (`n_cases`, `n_timepoints`×`n_channels`), handling each timepoint as a feature. The ranges of regression results by coefficient of determination ($R^2$) and root mean square error (RMSE) are shown in Table 3.

While $B_0$ and alpha concentration are well recovered using traditional approaches, pitch and density are very poorly predicted. Intuitively, models may benefit from awareness of spectral characteristics such as the shape and relative amplitude of peaks and self-interactions between integer frequency harmonics (appearing as, effectively, seasonality), and hence we have investigated Time-Series Extrinsic Regression (TSER). Regression of sequential data against continuous external values is a young field, having only been formalised by Tan et

**Table 3.** Range of $R^2$ and RMSE scores obtained by conventional (non-time-series) regression algorithms across a leave-one-out cross-validation for each target parameter. There was no significant difference in regression of $B_0$, while the RandomForestRegressor (RFR) was the strongest predictor for pitch and density, and the GradientBoostingRegressor (GBR) performed best for predicting alpha concentration.

| Parameter | Worst $R^2$ | Worst RMSE | Worst model | Best $R^2$ | Best RMSE | Best model |
|---|---|---|---|---|---|---|
| $B_0$ | 0.99 | $2.8 \times 10^{-2}$ | RFR | 1.0 | $3.0 \times 10^{-7}$ | LR |
| $\lambda$ | $-1.2$ | 1.5 | RR | 0.1 | 0.96 | RFR |
| $n_e$ | $-0.014$ | 1.0 | LR | 0.42 | 0.77 | RFR |
| $n_\alpha/n_e$ | 0.22 | 0.90 | LR | 0.83 | 0.42 | GBR |

al. in 2021 [41]. This is a novel approach in plasma physics, having to our knowledge only once previously been applied to the reconstruction of missing electron temperature data [46]. We used the TSER [3,24,33] capabilities of the Aeon time series analysis toolkit [34] to apply these algorithms to the novel area of ICE parameter inference. Each TSER model was trained on four input channels ($\log_{10}$ of the three EM spectra on a frequency range normalized to the per-case gyrofrequency; the fourth channel mapped gyrofrequencies to Hz) and regressed against a single output channel at a time (each of the plasma parameters in Table 3). In total, the dataset consisted of 110 cases, each with 3 EM spectra plus one frequency axis of 1280 points, meaning the total frequency space per case contained 5120 points.

Our pipeline was run on 14 extrinsic regression algorithms available in Aeon: Catch22 [30], Hydra [16], KNeighborsTimeSeries (KNN) [34], Random Convolutional Kernel Transform (ROCKET) [14], MiniROCKET [15], MultiROCKET [42], MultiROCKET-Hydra [16,42], QUANT [17], Random Interval [34], Random Interval Spectral Ensemble [29], Random Dilated Shapelet Transform (RDST) [25], Summary [34], Time-Series Forest (TSF) [20], and the Dummy regressor, which predicts the mean of the data. A minority of regressors in Aeon were not selected due to the tractability of training 440 models for the LOOCV. Figure 2 shows the coefficient of determination ($R^2$) across predictions for every fold. Hydra was the best performing algorithm by $R^2$ for two out of four simulation parameters and was within one standard error by RMSE of the top performance on all parameters. All predictions made by Hydra are shown in Figure 3. On a 13th generation 4.6GHz Intel i5-1335U, training of 110 folds of Hydra took approximately 2.5 CPU hours, while prediction of a single value took 100ms. Code, data and all regression results can be found in a Zenodo repository at `https://doi.org/10.5281/zenodo.21837478`.

## Discussion

From experiments using the synthetic $\delta B_z$, $\delta E_x$ and $\delta E_y$ spectra across four propagation angles, Aeon's implementations of Hydra, MultiROCKET-Hydra

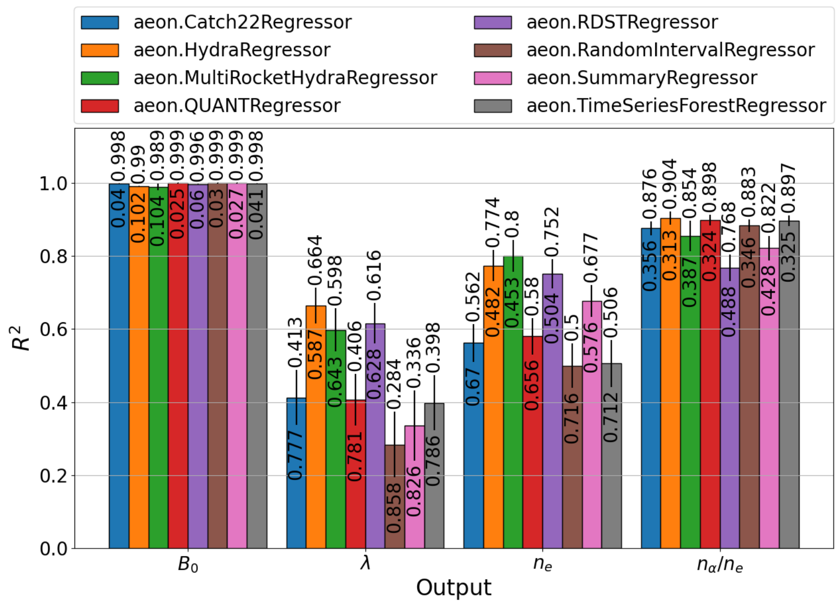


**Fig. 2.** TSER results ($R^2$ and RMSE) of a leave-one-out cross-validation of eight algorithms against four target parameters. Highest accuracy algorithms in LOOCV and experimental comparison were selected based on the coefficient of determination (higher is better) across all predictions. Error bars show the mean standard error of prediction across all folds of the LOOCV. $R^2$ values are shown above each error bar, with RMSE (lower is better) below.

(MRH) and RDST (excepting alpha concentration) perform the best on prediction of the four simulation inputs to within statistical significance, as defined by one standard error from the mean prediction. RMSE has the desirable property of units equal to the original data, so prediction errors can be interpreted in terms of physical quantities. A summary of these is shown in Table 4, where $B_0$, pitch, density and fast alpha concentration are predicted to within $0.6\sigma$ of their true values.

To provide a comparison to experiment, the ICE intensity spectrum from JET pulse No. 26148 was recovered from Figure 2 of Ref. [13]. This JET shot used deuterium and 11% tritium fuel and was taken close to the time of peak fusion activity, supporting its use as an experimental benchmark for ICE stimulated predominantly by fusion alphas in a primarily deuterium plasma. This spectrum resolved frequencies up to ∼187MHz and was obtained using a monopole antenna detecting the near-perpendicular component of $\vec{B}$. A comparable training set was built by truncating simulated $\delta B_z$ power spectra to this frequency range (resolving less than the first $4\Omega_{c,\alpha}$ in the cases with highest magnetic field). The unit of ICE intensity in the experimental data is dB without a stated physical baseline, so all simulated spectra were converted to a dB scale relative to

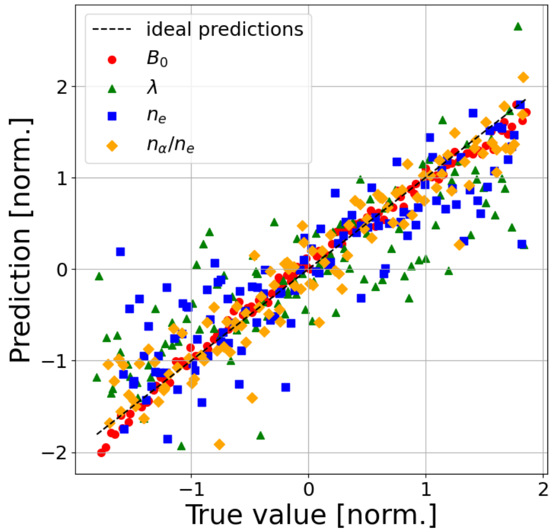


**Fig. 3.** Hydra-generated normalised predictions of each parameter. The datapoints are the single test points in each fold of a leave-one-out cross-validation.

**Table 4.** Best RMSE values and errors (ranges for log-sampled variables) achieved during extrinsic regression of power spectra against simulation input parameters, including denormalised RMSE values in original units. KNN is Aeon's implementation of a standard 1-nearest-neighbour algorithm. Note that for $B_0$, 11/13 algorithms produced statistically identical $R^2$ values of greater than 98%.

| Plasma Parameter | Best Model | RMSE (normalised) | RMSE (denormalised) |
|---|---|---|---|
| $B_0$ | KNN | 0.025 | 0.028T |
| $\lambda$ | Hydra | 0.589 | 0.160 |
| $n_e$ | MRH | 0.453 | $[2.17, 4.00] \times 10^{19}\text{m}^{-3}$ |
| $n_\alpha/n_e$ | Hydra | 0.313 | $[6.20, 14.0] \times 10^{-4}$ |

their global power minimum ($2.23 \times 10^{-12}\text{T}^2$). The $17 \pm 0.5$MHz spacing of the spectral peaks was assumed (by the authors of Ref. [13]) to represent the ground-truth alpha-particle gyrofrequency and converted to a magnetic field strength, $B = 2.21$T. Alpha-particle velocity pitch of 0.4 was estimated from Figure 15 of Ref. [13]. The density at the point of emission is inferred in Ref. [13] to be $1.7 \times 10^{19}\text{m}^{-3}$. The alpha-particle fraction was evaluated by Cottrell et al. using the plasma transport simulation code TRANSP and stated in Figure 18 of Ref. [13]. Training was run over all algorithms used during the previous LOOCV. Results from the nine best performing algorithms are shown in Figure 4. The best results by mean absolute error per parameter are shown in Table 5.

Frequency bandwidth presents a significant limitation when testing against experimental data. The temporal length and output frequency of the simulations

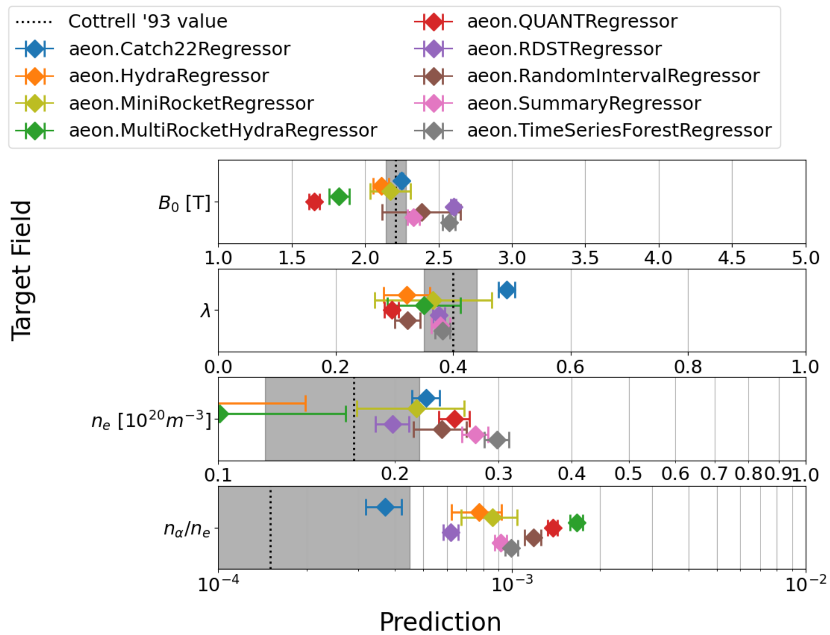


**Fig. 4.** Predictions made by TSER models trained on $\delta B_z$ spectra from all simulations truncated to equal frequency bandwidths ($\sim 0-187$MHz) against JET pulse No. 26148 ICE data extracted from Ref. [13]. Experimental values (excluding alpha-particle velocity pitch) are shown with dotted lines, with uncertainties denoted by shaded regions. The uncertainty in $B_0$ was estimated from measurements of the spectral peak spacing. The uncertainty in pitch was estimated from inspection of the shaded regions in Figures 15 and 18 of Ref. [13]. The uncertainty in density was found from correlation with the range of possible major radius values in Figure 7 [13]. The uncertainty in alpha-particle concentration was based on rationalising the radial density profiles in Figure 7 [13] with estimates from TRANSP, and the large variation between core and edge values. The x-axis spans the training data simulation range.

**Table 5.** Values shown are the true experimental value, the best prediction by mean absolute error of the best performing model in LOOCV (Hydra), the best prediction obtained, and the model making the best prediction. Uncertainty values for predictions are calculated as one standard deviation across 20 repeats. Uncertainty in the experimental data is transcribed or inferred from Ref. [13]; see Figure 4.

| Plasma Parameter | True Value | Hydra Prediction | Best Prediction | Best Pred. Model |
|---|---|---|---|---|
| $B_0$ [T] | $2.21 \pm 0.07$ | $2.11 \pm 0.05$ | $2.25 \pm 0.01$ | Catch22 |
| $\lambda$ | $0.40 \pm 0.05$ | $0.32 \pm 0.04$ | $0.38 \pm 0.01$ | TSF |
| $n_e$ [$10^{19}m^{-3}$] | $1.7 \pm 0.5$ | $0.9 \pm 0.5$ | $2.0 \pm 0.1$ | RDST |
| $n_\alpha/n_e$ [$10^{-4}$] | $1.5 \pm 3.0$ | $7.7 \pm 1.5$ | $3.7 \pm 0.5$ | Catch22 |

were set dynamically to ensure sufficient resolution up to $80\Omega_{c,\alpha}$. In performing the LOOCV analysis, the TSER algorithms are passed spectra discretized onto an axis normalized to the per-simulation cyclotron frequency, and hence this normalization cannot provide discriminative frequency information. This information is restored by including a fourth channel mapping cyclotron harmonics to Hz. When training the models for use with an experimental JET spectrum, however, the synthetic spectra are truncated and interpolated onto the frequency axis of the experimental observations. Hence higher frequency information is discarded; in the highest-$B_0$ cases, the bandwidth is less than $5\Omega_{c,\alpha}$.

Figure 4 shows an overlap between one standard deviation of MiniROCKET's prediction range and the experimental value's uncertainty for all parameters aside from alpha concentration, in which only Catch22 predicted to within the experimental error. Hydra also predicted these parameters well, with MultiROCKET-Hydra working well for pitch and density to within error. In the later case, these were the only two algorithms to predict below the experimental value. Alpha concentration was generally predicted least successfully, though the experimental value is near the lower limit of the training data. MultiROCKET-Hydra was one of the most accurate classifiers in a landmark comparison in Ref. [35] of 40 time-series algorithms, and was almost 200 times faster to train than its comparably-accurate competition. The speed of the ROCKET family of algorithms enhances their appeal for potential online training applications; for example, during experimental campaigns on novel fusion devices with sparse information on the mapping between their operational regimes and observed spectra. Overall, this alignment in results supports the hypothesis that training on power spectra from particle-in-cell codes is a suitable application for the time-series approach and bears similarity to standard time series classification benchmarks. Many of the regression algorithms in Aeon were adapted from classification variants [34].

While the consistently highest performing models are from the convolution family of TSER algorithms (Hydra, ROCKET and their derivatives), it is notable that RDST, a shapelet-based approach, outperforms them on three out of four parameters in Figure 4. The authors of Ref. [35] note that despite some methodological similarity between shapelet and convolutional algorithms in their approach to feature extraction, a key difference is that shapelets are direct subseries of the training data whereas convolutions are computed from the entire dataseries range. This in turn may shed some light on the physics involved; the significant features determined by pitch, density and alpha-particle concentration manifest more in narrow frequency band phrases such as the qualities of peak splitting and the character of their prominence than in a more holistic spectral form, such as how power is distributed throughout the frequency bandwidth. Variation in these quantities controls the quantity of fast ions with velocity and orientation sufficient to excite the MCI, and is likely to correlate with absolute amplitude, a difficult quantity to rationalise between simulation and experiment with uncertain dB scales. Spectra at lower powers appear noisier as they approach the resolution limit, which may be recognised by shapelet-based approaches. Given the data range, collection and power scale limitations making the problem more difficult, this is a successful outcome. LOOCV performance indicates the capacity of a model to learn the training data, and the experimental comparison indicates the training data's proximity to a real spectrum and thus the predictive power of models trained only on synthetic spectra.

Several caveats apply to these results. The periodic simulation domain implies a homogeneous, infinite 1D plasma without source or sink effects, complex geometry, spatially-varying background electromagnetic fields or the potential for waves to convert into other modes. 2D effects were incorporated by combining spectra from four angles near perpendicular, but given the qualitative differences observed even in spectra separated by only 2°, there is likely to be additional phenomena observed in emission from intermediate ranges and their subsequent propagation, attenuation and detection. Experimental spectra are likely to contain peaks affected by Doppler shifts [22], doublet-splitting [11] and nonlinearities. These physical effects increase the complexity and fine structure of the spectra that arise from small differences in parameters, which must be disambiguated by TSER algorithms and may either reduce or strengthen their predictive power.

Regardless, Aeon's algorithms' success at predicting parameters based on an experimental spectrum, with all the real world physics it contains, having been trained only on simulated analogues is a credit to the robustness of the approach. A total data corpus of 110 simulations is small in machine learning terms for scanning four parameters simultaneously. The non-trivial inversion of this problem is made challenging by the covariance of ICE spectral characteristics with the various input parameters, which could be mitigated by increasing the dataset size and restricting the parameters to known experimental combinations. Improvements to a diagnostically-relevant ICE model could be achieved by: running fully kinetic 2D simulations in a narrower parameter space tailored to a particular device; increasing the training size as much as computational resources allow while scanning additional variables known to impact ICE such as temperature and the ratio of fast ion $v_\perp$ to significant physical thresholds such as the Alfvén wave group velocity [38]; adding tritium to the background species; varying the background electron and fuel ion temperatures, potentially independently; and using a dynamic temporal simulation length to ensure the MCI reaches saturation in all cases. It should also be noted that the particular plasma parameters considered herein were chosen based on both their operational usefulness and qualitative estimations of their likelihood to influence the MCI. In particular, if other information can disambiguate the up-down symmetry in a tokamak, $B_0$ and density information can spatially localise the point of emission. There may be many other quantities encoded in ICE spectra recoverable through this workflow.

## Conclusions

By combining state-of-the-art temporal machine learning and first-principles nonlinear kinetic plasma physics, this article describes progress towards exploiting ICE spectra as a tool for inferring bulk and fast ion plasma parameters for the first time. We ran 110 plasma parameter sets of nonlinear 1D3V fully-kinetic PIC simulations at four propagations angles near perpendicular to $B_0$ to efficiently build composite power spectra with 2D spatial information. These simulations comprise fusion alpha-particles in the drifting “ring-beam” class of distribution function with thermal fuel deuterons and electrons, a scenario known to produce spectra characteristic of ICE. This took 1.5 million CPU hours, enabling most labs to generate comparable training sets and build similar inference models for their devices. TSER models were trained in a novel plasma physics context to predict $B_0$ strength (peak LOOCV $R^2 = 0.99$), electron density ($R^2 = 0.80$) and alpha-particle velocity pitch ($R^2 = 0.66$) and concentration ($R^2 = 0.90$) from synthetic $\delta B_z$, $\delta E_x$ and $\delta E_y$ frequency spectra derived from simulations. Comparison to a gold-standard JET ICE spectrum from Cottrell et al. [13] required only one log-scale $\delta B_z$ spectrum truncated in frequency to 187 MHz, well below today’s diagnostic bandwidth capability. Despite these obstacles, the best performing algorithms were able to, based on the experimental spectrum only, predict the emitting plasma’s $B_0$ to within 2%, alpha velocity pitch to within 5% and density to within 18%. The inference time for the best performing models (Hydra, MultiROCKET-Hydra, RDST) is around 100ms, making this approach suitable for analysis of ICE spectra, given a pre-trained prediction model, between plasma pulses of ongoing experimental campaigns. With conceivable performance and resolution optimisations, it could be extended to real-time control applications for the next generation of fusion power plants.

**Acknowledgments.** This work was funded jointly by the UK Atomic Energy Authority and as part of the EPSRC Centre for Doctoral Training in Fusion Energy Science and Technology (Fusion Power CDT) with project reference EP/S022430/1. EPOCH was in part funded by the EPSRC grants EP/G054950/1, EP/G056803/1, EP/G055165/1, EP/ M022463/1 and EP/P02212X/1. This work used the ARCHER2 UK National Supercomputing Service (`https://www.archer2.ac.uk`) through EPSRC grant EP/X035336/1. The Viking cluster was also used during this project, which is a high performance compute facility provided by the University of York. We are grateful for computational support from the University of York, IT Services and the Research IT team.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

**Data Availability.** This work used EPOCH1D branch 5.0-devel (commit `f42ebadc`) available at `https://github.com/epochpic/epoch/tree/5.0-devel` and aeon version 1.3.0 available at `https://github.com/aeon-toolkit/aeon/tree/v1.3.0`. Relevant analysis code, data and results can be accessed at `https://doi.org/10.5281/zenodo.21837478`.

## References

1. Abdou, M., Riva, M., Ying, A., Day, C., Loarte, A., Baylor, L., Humrickhouse, P., Fuerst, T.F., Cho, S.: Physics and technology considerations for the deuterium–

tritium fuel cycle and conditions for tritium fuel self sufficiency. Nuclear fusion **61**(1), 013001 (2021)
2. Arber, T.D., Bennett, K., Brady, C.S., Lawrence-Douglas, A., Ramsay, M.G., Sircombe, N.J., Gillies, P., Evans, R.G., Schmitz, H., Bell, A.R., Ridgers, C.P.: Contemporary particle-in-cell approach to laser-plasma modelling. Plasma Physics and Controlled Fusion **57**(11), 1–26 (Nov 2015)
3. Bagnall, A., Middlehurst, M., Forestier, G., Ismail-Fawaz, A., Guillaume, A., Guijo-Rubio, D., Tan, C.W., Dempster, A., Webb, G.I.: A hands-on introduction to time series classification and regression. In: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. pp. 6410–6411 (2024)
4. Beckett, G., Beech-Brandt, J., Leach, K., Payne, Z., Simpson, A., Smith, L., Turner, A., Whiting, A.: Archer2 service description (Dec 2024). `https://doi.org/10.5281/zenodo.14507040`
5. Birdsall, C.K., Langdon, A.B.: Plasma physics via computer simulation. CRC Press (1991)
6. Breiman, L.: Random forests. Machine learning **45**(1), 5–32 (2001)
7. Carbajal, L., Dendy, R.O., Chapman, S.C., Cook, J.W.S.: Linear and nonlinear physics of the magnetoacoustic cyclotron instability of fusion-born ions in relation to ion cyclotron emission. Physics of Plasmas **21**(1) (2014)
8. Chapman, B., Dendy, R.O., Chapman, S.C., Holland, L.A., Irvine, S.W.A., Reman, B.C.G.: Comparing theory and simulation of ion cyclotron emission from energetic ion populations with spherical shell and ring-beam distributions in velocity-space. Plasma Physics and Controlled Fusion **62**(5), 055003 (2020)
9. Chapman, B., Dendy, R.O., Chapman, S.C., McClements, K.G., Ochoukov, R.: Origin of ion cyclotron emission at the proton cyclotron frequency from the core of deuterium plasmas in the asdex-upgrade tokamak. Plasma Physics and Controlled Fusion **62**(9), 095022 (2020)
10. Chapman, B., Dendy, R.O., Chapman, S.C., McClements, K.G., Yun, G.S., Thatipamula, S.G., Kim, M.: Nonlinear wave interactions generate high-harmonic cyclotron emission from fusion-born protons during a kstar elm crash. Nuclear Fusion **58**(9), 096027 (2018)
11. Cook, J.W.S.: Doublet splitting of fusion alpha particle driven ion cyclotron emission. Plasma Physics and Controlled Fusion **64**(11), 115002 (2022)
12. Cook, J.W.S., Dendy, R.O., Chapman, S.C.: Particle-in-cell simulations of the magnetoacoustic cyclotron instability of fusion-born alpha-particles in tokamak plasmas. Plasma Physics and Controlled Fusion **55**(6), 065003 (2013)
13. Cottrell, G.A., Bhatnagar, V.P., Da Costa, O., Dendy, R.O., Jacquinot, J., McClements, K.G., McCune, D.C., Nave, M.F.F., Smeulders, P., Start, D.F.H.: Ion cyclotron emission measurements during jet deuterium-tritium experiments. Nuclear Fusion **33**(9), 1365 (1993)
14. Dempster, A., François, P., Webb, G.I.: Rocket: exceptionally fast and accurate time series classification using random convolutional kernels. Data Mining and Knowledge Discovery **34**(5), 1454–1495 (2020)
15. Dempster, A., Schmidt, D.F., Webb, G.I.: Minirocket: A very fast (almost) deterministic transform for time series classification. In: Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining. pp. 248–257 (2021)
16. Dempster, A., Schmidt, D.F., Webb, G.I.: Hydra: competing convolutional kernels for fast and accurate time series classification: A. dempster et al. Data Mining and Knowledge Discovery **37**(5), 1779–1805 (2023)

17. Dempster, A., Schmidt, D.F., Webb, G.I.: Quant: A minimalist interval method for time series classification. Data Mining and Knowledge Discovery **38**(4), 2377–2402 (2024)
18. Dendy, R.O., McClements, K.G., Lashmore-Davies, C., Cottrell, G., Majeski, R., Cauffman, S.: Ion cyclotron emission due to collective instability of fusion products and beam ions in tftr and jet. Nuclear fusion **35**(12), 1733 (1995)
19. Dendy, R., McClements, K.: Ion cyclotron emission from fusion-born ions in large tokamak plasmas: a brief review from jet and tftr to iter. Plasma Physics and controlled fusion **57**(4), 044002 (2015)
20. Deng, H., Runger, G., Tuv, E., Vladimir, M.: A time series forest for classification and feature extraction. Information Sciences **239**, 142–153 (2013)
21. Friedman, J.H.: Greedy function approximation: a gradient boosting machine. Annals of statistics pp. 1189–1232 (2001)
22. Gorelenkov, N., Cheng, C.: Alfvén cyclotron instability and ion cyclotron emission. Nuclear fusion **35**(12), 1743–1752 (1995)
23. Gormezano, C., Challis, C., Joffrin, E., Litaudon, X., Sips, A.: Chapter 4: Advanced tokamak scenario development at jet. Fusion science and technology **53**(4), 958–988 (2008)
24. Guijo-Rubio, D., Middlehurst, M., Arcencio, G., Silva, D.F., Bagnall, A.: Unsupervised feature based algorithms for time series extrinsic regression. Data Mining and Knowledge Discovery **38**(4), 2141–2185 (2024)
25. Guillaume, A., Vrain, C., Elloumi, W.: Random dilated shapelet transform: A new approach for time series shapelets. In: International Conference on Pattern Recognition and Artificial Intelligence. pp. 653–664. Springer (2022)
26. Hockney, R.W., Eastwood, J.W.: Computer simulation using particles. CRC Press (1988)
27. Hoerl, A.E., Kennard, R.W.: Ridge regression: Biased estimation for nonorthogonal problems. Technometrics **12**(1), 55–67 (1970)
28. Lanti, E., Ohana, N., Tronko, N., Hayward-Schneider, T., Bottino, A., McMillan, B.F., Mishchenko, A., Scheinberg, A., Biancalani, A., Angelino, P., et al.: Orb5: a global electromagnetic gyrokinetic code using the pic approach in toroidal geometry. Computer Physics Communications **251**, 107072 (2020)
29. Lines, J., Taylor, S., Bagnall, A.: Time series classification with hive-cote: The hierarchical vote collective of transformation-based ensembles. ACM Transactions on Knowledge Discovery from Data (TKDD) **12**(5), 1–35 (2018)
30. Lubba, C.H., Sethi, S.S., Knaute, P., Schultz, S.R., Fulcher, B.D., Jones, N.S.: Catch22: Canonical time-series characteristics: Selected through highly comparative time-series analysis. Data mining and knowledge discovery **33**(6), 1821–1852 (2019)
31. McClements, K.G., Brisset, A., Chapman, B., Chapman, S.C., Dendy, R.O., Jacquet, P., Kiptily, V.G., Mantsinen, M., Reman, B.C.G., Contributors, J.: Observations and modelling of ion cyclotron emission observed in jet plasmas using a sub-harmonic arc detection system during ion cyclotron resonance heating. Nuclear Fusion **58**(9), 096020 (2018)
32. McClements, K.G., Dendy, R.O., Lashmore-Davies, C.N., Cottrell, G.A., Cauffman, S., Majeski, R.: Interpretation of ion cyclotron emission from sub-alfvénic fusion products in the tokamak fusion test reactor. Physics of Plasmas **3**(2), 543–553 (1996)
33. Middlehurst, M., Bagnall, A.: Extracting features from random subseries: A hybrid pipeline for time series classification and extrinsic regression. In: International

Workshop on Advanced Analytics and Learning on Temporal Data. pp. 113–126. Springer (2023)
34. Middlehurst, M., Ismail-Fawaz, A., Guillaume, A., Holder, C., Guijo-Rubio, D., Bulatova, G., Tsaprounis, L., Mentel, L., Walter, M., Schäfer, P., Bagnall, A.: Aeon: a python toolkit for learning from time series. Journal of Machine Learning Research **25**(289), 1–10 (2024), `http://jmlr.org/papers/v25/23-1444.html`
35. Middlehurst, M., Schäfer, P., Bagnall, A.: Bake off redux: a review and experimental evaluation of recent time series classification algorithms. Data Mining and Knowledge Discovery **38**(4), 1958–2031 (2024)
36. Ochoukov, R., Bobkov, V., Chapman, B., Dendy, R., Dunne, M., Faugel, H., García-Muñoz, M., Geiger, B., Hennequin, P., McClements, K., et al.: Observations of core ion cyclotron emission on asdex upgrade tokamak. Review of Scientific Instruments **89**(10) (2018)
37. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
38. Reman, B.C.G., Dendy, R.O., Akiyama, T., Chapman, S.C., Cook, J.W.S., Igami, H., Inagaki, S., Saito, K., Yun, G.: Interpreting observations of ion cyclotron emission from large helical device plasmas with beam-injected ion populations. Nuclear Fusion **59**(9), 096013 (2019)
39. Sips, A.C.C., for the Steady State Operation, the Transport Physics topical groups of the International Tokamak Physics Activity: Advanced scenarios for iter operation. Plasma Physics and Controlled Fusion **47**(5A), A19 (apr 2005). `https://doi.org/10.1088/0741-3335/47/5A/003`, `https://doi.org/10.1088/0741-3335/47/5A/003`
40. Slade-Harajda, T.W., Cook, J.W.S., Dendy, R.O., Chapman, S.C.: Effects of deuterium-tritium mix on linear growth rates of the magnetoacoustic cyclotron instability in fusion plasmas. Physics of Plasmas **32**(8) (2025)
41. Tan, C.W., Bergmeir, C., Petitjean, F., Webb, G.I.: Time series extrinsic regression: Predicting numeric values from time series data. Data Mining and Knowledge Discovery **35**(3), 1032–1060 (2021)
42. Tan, C.W., Dempster, A., Bergmeir, C., Webb, G.I.: Multirocket: multiple pooling operators and transformations for fast and effective time series classification: Cw tan. Data Mining and Knowledge Discovery **36**(5), 1623–1646 (2022)
43. Tholerus, E., Casson, F., Marsden, S., Wilson, T., Brunetti, D., Fox, P., Freethy, S., Hender, T., Henderson, S., Hudoba, A., Kirov, K., Koechl, F., Meyer, H., Muldrew, S., Olde, C., Patel, B., Roach, C., Saarelma, S., Xia, G., the STEP team: Flattop plasma operational space of the step power plant. Nuclear Fusion **64**(10), 106030 (aug 2024). `https://doi.org/10.1088/1741-4326/ad6ea2`, `https://doi.org/10.1088/1741-4326/ad6ea2`
44. Tibshirani, R.: Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society Series B: Statistical Methodology **58**(1), 267–288 (1996)
45. Todd, T.: Diagnostic systems in demo: engineering design issues. In: AIP Conference Proceedings. vol. 1612, pp. 9–16. American Institute of Physics (2014)
46. Wang, M., Wan, C., Lu, J., Yu, Z., Xiao, B., Li, Y., He, X., Luo, Z., Yuan, Q., Hu, Y., et al.: Time series extrinsic regression for reconstructing missing electron temperature in tokamak. Nuclear Fusion **65**(7), 076008 (2025)