Fusion Engineering and Design 220 (2025) 115318

Contents lists available at ScienceDirect

Fusion Engineering and Design

journal homepage: www.elsevier.com/locate/fusengdes

On the potential and limitations of Bayesian ensemble algorithms for the decomposition of time series generated by tokamak diagnostics

Michela Gelfusa [a,1,*], Teddy Craciunescu [b,1], Riccardo Rossi [a], Andrea Murari [c,d], on behalf of JET Contributors[2], the EUROfusion Tokamak Exploitation Team[3]

[a] Department of Industrial Engineering, University of Rome "Tor Vergata", via del Politecnico 1, 00133 Roma, Italy
[b] National Institute for Laser, Plasma and Radiation Physics, Bucharest-Magurele 077125, Romania
[c] Consorzio RFX (CNR, ENEA, INFN, Università di Padova, Acciaierie Venete SpA), Corso Stati Uniti 4, 35127 Padova, Italy
[d] Istituto per la Scienza e la Tecnologia dei Plasmi, CNR, Padova, Italy

ARTICLE INFO

Keywords:
Time series decomposition
Bayesian model averaging
Trend analysis
Stationarity
ELMs
Sawteeth
Long pulse operation

ABSTRACT

The vast majority of signals generated by tokamak diagnostics are in the form of time series. Consequently dealing with time-indexed data is a major task, to be tackled daily by both experimentalists and analysts. Decomposing a time series in terms of seasonal components, trends, change-points and noise is therefore a crucial activity, per se and as a preliminary step to further investigations. In the present work, the Bayesian ensemble approach to model decomposition of time series, originally developed for remote sensing of the earth, is applied to various global measurements routinely available in tokamak devices. Among the competitive advantages of the methodology, particularly relevant are its holistic view of the data and the independence from the details of the statistical algorithms and models. The potential of the technique, implemented by the BEAST code, has been assessed with both synthetic signals and experimental data. The approach proves to be very reliable in modelling trends and determining the time locations of abrupt changes even of strongly oscillatory components, such as ELMs and sawteeth. Deployment to assess small drifts confirms the lack of stationarity in tokamak high performance discharges. The difficulties of modelling the details of the sawteeth and irregular ELMs indicate the need to improve the method to deal with seasonal components of complex harmonic content and/or varying frequency. However, the available routines are already very effective in determining the times changes in the ELM regimes.

1. The importance and challenges of time series empirical modelling in tokamaks

The tokamak is an inductive magnetic configuration and consequently its discharges are inherently transient [1]. One of the important challenges in the perspective of practical applications of the concept remains extending the high performance phase of the pulses, to achieve Long Pulse Operation (LPO). Indeed controlling fusion relevant plasmas for long periods, possibly even toward steady-state operation, is an essential element of the international fusion programme, including ITER, on the route to future fusion reactors [2]. To optimise the international effort to achieve this goal, the International Atomic Energy Agency (IAEA) has even established a working group of experts called CICLOP, which stands for Coordination on International Challenges on Long duration Operation [2]. Indeed, the route to viable LPO in tokamaks is full of difficulties, because it will require controlling stable plasma for a duration much longer than the plasma energy confinement time and approaching plasma-wall interaction (PWI) timescales. Operating a reactor grade tokamak device for a few hours, even in a pulsed mode of operation, implies addressing three main orders of problems, often conflicting or at least not easy to harmonise: a) maintaining high performances (sufficiently high neutron yield) b) preserving the integrity of the plasma facing components c) avoiding dangerous instabilities leading to disruptions.

* Corresponding author.
E-mail address: gelfusa@ing.uniroma2.it (M. Gelfusa).

[1] These authors contributed equally to the paper
[2] See the author list of C.F. Maggi et al 2024 Nucl. Fusion 64 112012 https://doi.org/10.1088/1741-4326/ad3e16
[3] See the author list of E. Joffrin et al 2024 Nucl. Fusion 64 112019 https://doi.org/10.1088/1741-4326/ad2be4

https://doi.org/10.1016/j.fusengdes.2025.115318
Received 28 February 2025; Received in revised form 1 July 2025; Accepted 3 July 2025
Available online 21 July 2025

At present, the top performances have been achieved in the so-called high confinement regime (H-mode) [3]. The H-mode is characterised by an edge transport barrier and the build-up of a steep edge pressure gradient (pedestal). The overall plasma performance is determined by the quality of the pedestal, because its height constitutes the boundary condition for the core plasma confinement. However, the pedestal pressure is limited by Edge Localized Modes (ELMs) [4,5], charge-and-fire instabilities that cause the explosive ejection of energy and particles into open field lines, rapidly reaching the plasma facing components. The consequent energy loads and erosion of divertor target plates are deemed excessive even for the most resistant current materials and are therefore unacceptable in future burning plasma devices [6,7]. This situation has led the international fusion programme to dedicate substantial efforts and experimental time in two main directions: a) the development of no-ELM and small-ELM regimes, b) the implementation of active ELM control techniques [8]. However, even robust no-ELM scenarios or active ELM control techniques need to preserve high performance. In the perspective of future devices, this implies for example compatibility with low collisionality, high Greenwald and radiation fractions, dominant electron heating and steady state operation of the divertor.

To preserve the performances of the ITER baseline scenario for long intervals, another instability to carefully control are the sawteeth. Contrary to the ELMs, which are an edge phenomenon, these relaxation instabilities are localised in the centre of the plasma and their signature is a non-linear oscillation of the central plasma parameters [8]. If they are not too large, sawteeth, do not case a significant of degradation of plasma performance. They can even prove to be beneficial in the reactor, because they can help expelling impurities and helium ash from the plasma core. Indeed recently, evidence has emerged showing that sawteeth can play an important role in avoiding the core accumulation of heavy impurities in metallic machines [9,10]. On the contrary, if the sawtooth crashes become too large, they can trigger more deleterious instabilities, such as NTMs, localised at flux surfaces, whose the safety factor q is a rational number (typically $q = 2$ or $q = 3/2$). NTMs typically have a very detrimental effect on plasma confinement and they can even cause disruptions [11]. In the perspective of the reactor, it is therefore essential control the sawteeth quite reliably, and indeed significant efforts have been devoted to this aspect in the last years [12].

Even this brief introduction suggests that carefully monitoring the evolution of the plasma conditions is fundamental and requires analysing all aspects of time indexed data, from drift and slowly varying trends to oscillatory components. The typical form of the information provided by tokomak diagnostics is indeed as time series. Properly understanding and modelling this type of signal is therefore of the highest importance. The analysis of time series has indeed reached high levels of sophistication in the fusion community. All the main techniques, from Fourier analyses to wavelet and empirical mode decomposition have been implemented [13–16]. Even more recent and involved approaches, such as recurrence plots, causality detection and physics informed neural networks have been adapted and deployed [17,18]. Bayesian statistics has also been extensively and very successfully applied to fusion data, including time series, particularly for Integrated Data Analysis (IDA) [19–25].

Notwithstanding the advanced level of many techniques, the first screening of the signals is often performed manually or with quite simplistic approaches. Moreover, all techniques present their limitations and none outperforms all the alternatives in all respects. Indeed, some rely on quite restrictive assumptions, others tend to address preferentially only some aspects. This can lead to limited applicability or scarce statistical reliability. Typical practical examples of the ensuing issues are the difficulties of classifying the ELM type [26], of assessing the effectiveness of sawteeth pacing techniques [27] or of identifying the root causes of disruptions [28,29].

The objective of the present work consists of assessing whether a Bayesian technique, originally developed for remote sensing, can provide a quite complete empirical model of time series generated by the main tokamak diagnostics. Indeed, the method was conceived to decompose a time-indexed signal in terms of seasonal components, trends, change-points and noise. In the following, seasonal components indicate periodic or quasi-periodic elements of the time series, reminiscent of the changes in many environmental phenomena with the seasons. In tokamaks, they can be due to macroscopic charge-and-fire instabilities such as ELMs or sawteeth. Trends indicate slow drifts in the characteristics of a time series and can therefore reveal a progressive deterioration of the discharge properties and performances; impurity accumulation or current diffusion are typical examples of this type of behaviour. These lower frequency variations can be peppered with abrupt change-points, due to sudden instabilities or control problems. Significant variations in the noise level or statistics can reveal either problems in the diagnostic systems or changes in the operational regime. Reliable statistical detection and classification of these elements would therefore be extremely beneficial for any type of subsequent analysis.

In this respects the ensemble Bayesian approach presents various potential advantages: rather than delivering only event times, the method yields full posterior distributions for seasonal amplitudes and phases, provides probabilistic uncertainty bands around repeating patterns, and can jointly model slow secular drifts alongside periodic events in a single coherent framework. In general the technique gives a more holistic view of time series as will be substantiated later in the paper.

In mathematical terms, the empirical interpretation of a time series consists of finding the relationship between a quantity y and time t from the observed data at n points $D = \{t_i, y_i\}$ $i = 1,\ldots, n$ using a statistical or mechanistic (deterministic) decomposition model $y(t) = f(t)$. As mentioned, in general the model comprises four components; a seasonal $S(\bullet)$ component, a trend $T(\bullet)$, abrupt changes and a noise term ε:

$$y(t_i) = f(t_i; \Theta) = S(t_i; \Theta_s) + T(t_i; \Theta_T) + \varepsilon_i, i = 1, \ldots, n \tag{1}$$

In Eq. (1) the parameters Θ_s and Θ_T determine the properties of the seasonal and trend signals and are to be estimated from the data D. Traditionally the parameters Θs and Θ_T implicitly encode also the change-points as abrupt variations of the parameters themselves. In the following, the noise ε is assumed to be Gaussian with a magnitude σ, modelling the information in the data not explained by the seasonal $S(\bullet)$ and trend $T(\bullet)$ terms. The assumption of Gaussian noise is normally well justified in tokamak environments. In any case, even if an exhaustive analysis of the noise statistics and its effects is beyond the scope of the present work, various tests have shown that the results are quite consistent even if the probability distribution of the signal uncertainties departs significantly from Gaussianity.

Most techniques tend to focus preferentially on the detection of only one of the elements in Eq. (1) or rely on strong hypotheses such as strict stationarity. For example, some techniques require specifying the number of change points or the harmonic order of the seasonal components. Recently, new methods conceived to detect only abrupt changes have been implemented to predict the onset of disruptions [30]. Consequently, typical interpretations of time series depend a lot on the choice of the statistical algorithms, assumptions and models [31]. Moreover, in tokamaks the time series can reflect quite strong nonlinearities, which cannot be adequately approximated by linear or piecewise linear models typical of most popular techniques.

Another often underestimated difficulty resides in the fact that the four aforementioned aspects of a system time evolution can influence each other; determining them simultaneously in a coherent way is therefore one of the main advantages of the approach proposed in the present work. Indeed accurate quantification of trends, seasonality, change-points and noise is not a separable problem. Any error in quantifying one term would tend to leak and affect the estimation of the others. For example, it is unlikely that abrupt changes can be properly identified if the trend and/or seasonality are poorly fitted.

There are two additional practical aspects that traditional analysis methods have problems to deal with. Firstly, they have problems quantifying the decomposition uncertainties due to the model misspecification of the four components. Secondly, it should not be neglected that the signals of tokamak diagnostics are typically affected by significant levels of noise. Such high levels of uncertainties can significantly complicate the estimation of the individual elements and their mutual relationship.

The limitations of traditional statistical analysis techniques can be alleviated by deploying Bayesian methods. Indeed, if the four main components —trend, seasonality, change-points and noise— contribute additively to form the experimental signals, Bayesian statistics can be quite effective in inferring the decomposition model of Eq. (1) [32,33]. In the next section, the main characteristics and properties of a powerful Bayesian ensemble algorithm, a Bayesian Estimator of Abrupt change, Seasonal change, and Trend (BEAST) [31], are introduced. The code implementing BEAST is meant to be run autonomously without any need for human intervention. Some alternative frequentist criteria often deployed in time series analysis are also shortly reviewed for completeness sake. The tests of the code with synthetic signals of the type generated by tokamak diagnostics and benchmark experimental data from other systems are described in Section 3. The potential and limitations of the approach to analyse tokamak data are presented in Section 4, in which two particular aspects are discussed: the assessment of stationarity and the changes in seasonal components. These are two characteristics of time series modelling that are often not addressed with sound statistical techniques in the community. The subject of Section 5 is computational demands, before the discussion of the prospects of future applications in the conclusion section.

2. Overview of Bayesian ensemble algorithms for time series decomposition

To render the paper self-contained, this section provides an overview of the Bayesian ensemble approach to model decomposition of time series and in particular of its implementation with the BEAST code. All the details about the method and its interpretations can be found in [31] and references therein.

2.1. The specificities of Bayesian ensemble modelling

In the last decades, a significant body of theoretical and empirical work has shown the effectiveness of the Bayesian paradigm [34–36]. One of its characteristics and potential advantages is that Bayesian inference treats both model parameters and structures as random variables. Describing their parameters probabilistically, Bayesian inference tends not to converge on a single model as the true one [37,38]. On the contrary, each model is assigned a probability of being the true model and this probability can be learned from the data. In the so-called Bayesian model averaging (BMA) methodology, all the candidate models are weighted and contribute to the definition of the final average model. The advantages of the BMA approach range from the capability of approximating complex relationships, not expressible by individual models, to the alleviation of model misspecifications and uncertainties [31]. These positives are not limited to application involving cross sectional data but can be achieved also in the case of time series, as discussed in some detail in the next subsections.

2.2. Parametric decomposition of time series

As mentioned, the analysis presented in the following is based on the assumption that the time series under study are composed of four additive components—seasonality, trend, abrupt changes and noise, as formalised in Eq. (1). In the proposed treatment, the seasonal signal $S(t)$ is modelled as a piecewise harmonic function, defined with respect to p knots in time at ξ_k, $k = 1, \dots p$ that divide the time series into $(p + 1)$ intervals $[\xi_k, \xi_{k+1}]$, $k = 0, \dots, p$ [39].

In each of the $(p + 1)$ segments the model is specified as:

$$S(t) = \sum_{l=1}^{L_k} \left[a_{k,l} \sin\left(\frac{2\pi l t}{P}\right) + b_{k,l} \cos\left(\frac{2\pi l t}{P}\right) \right]; \quad \xi_k \le t < \xi_{k+1}, k = 0, \dots, p \tag{2}$$

where, P is the period of the seasonal signal and L_k is an unknown segment-specific parameter that determines the harmonic order of the kth segment. It should be appreciated that this harmonic model is non-continuous, because the knots ξ_k indicate the change-points at which abrupt changes may take place. Both the total number of change-points p and their locations at times ξ_k are unknown parameters. To summarise this part, the seasonal harmonic curve is completely defined by the following set of parameters:

$$\Theta_s = \{p\} \cup \{\xi_k\}_{k=1,\dots,p} \mathrm{U} \{L_k\}_{k=0,\dots,p} \cup \cup \{a_{k,l}, b_{k,l}\}_{k=0,\dots,p; l=1,\dots,L_k} \tag{3}$$

They include the number and time locations of the seasonal change-points, the harmonic orders of the $(p + 1)$ segments, and the coefficients of the harmonic terms. All these parameters have to be estimated from the data.

The term modelling the trends $T(t)$ is assumed to be a piecewise linear function with respect to m knots at the times τ_j, $j = 1, \dots m$. Therefore, in each interval the trend is a line segment defined by coefficients a_j and b_j:

$$T_j(\mathrm{t}) = a_j + b_j t \text{ for } \tau_j \le \mathrm{t} < \tau_{j+1}, j = 0, \dots, m. \tag{4}$$

Of course, the trend change-points τ_j are not necessarily located at the same times as the seasonal change-points ξ_k. As for the seasonal components, the number of change-points m and their time locations are unknown parameters. Hence, the full set of parameters required to specify the trends is:

$$\Theta_T = \{m\} \cup \{\tau_j\} \cup \{a_j, b_j\} \; j = 0, \dots, m \tag{5}$$

which comprises the number and time locations of trend change-points, and the intercepts and slopes of the individual line segments. As for Θ_S, the parameters of the set Θ_T need to be estimated from the data.

For the sake of intuitive interpretation of the results, it is useful to regroup the parameters of Θ_T and Θ_s into two alternative sets: $\{\Theta_T, \Theta_s\} = \{\boldsymbol{M}, \boldsymbol{\beta_M}\}$. The first set $\mathbf{M}$ includes the parameters required to define the model structure: numbers and time locations of trend and seasonal change-points, and seasonal harmonic orders:

$$\boldsymbol{M} = \{m\} \cup \{\tau_j\}_{j=1,\dots,m} \cup \{p\} U \{\xi_k\}_{k=-1,\dots,p} U \cup \{L_k\} k = 0, -D \tag{6}$$

The second group $\boldsymbol{\beta_M}$ includes the segment-specific coefficients required to determine the exact shape of the trend and seasonal curves, given the model structure. $\boldsymbol{\beta_M}$ is therefore completely determined by the following set of parameters:

$$\beta_{\mathbf{M}} = \{\mathrm{a}_j, \mathrm{b}_j\}_{j = 0\dots,m} \cup \{\mathrm{a}_{k,l}, \mathrm{b}_{k,l}\}_{k = 0,\dots,p;\; l = 1,\dots,\; Lk.} \tag{7}$$

Where the subscript M indicates the dependence of $\boldsymbol{\beta_M}$ on the model structure.

With this re-grouping, the general linear model of Eq. (1) can be rewritten as [31]:

$$\mathbf{y(t_i) = xM(t_i)\beta_M + \sigma_i} \tag{8}$$

where $\boldsymbol{xM(t_i)}$ and $\boldsymbol{\beta_M}$ are dependent variables and associated coefficients.

The form of Eq. (8) reveals that the task of decomposing a time series is now converted into a model selection problem: the determination of an appropriate model structure $\boldsymbol{M}$, including the numbers and time locations of the abrupt changes, and the harmonic orders of the seasonal and trend components.

From a conceptual point of view, once a model structure $\boldsymbol{M}$ is selected, it is relatively straightforward to estimate its coefficients $\boldsymbol{\beta_M}$. However, in practice, the number of possible model structures for Eq. (8) is extremely large. Even for a time series of moderate length (e.g., $n \approx 100$), it becomes computationally impossible to find the best one by optimising traditional model selection criteria such as the AIC and BIC [40]. The BEAST code circumvents this difficulty by resorting to Bayesian inference, as described in the next subsection.

2.3. The Bayesian formulation of time series decomposition by beast

One of the specificities of Bayesian modelling consists of considering all the unknown parameters as random variables. In the present treatment, this includes the model structure $\boldsymbol{M}$, the coefficients $\boldsymbol{\beta_M}$, and the data noise $\boldsymbol{\sigma}$. However the goal is not only to obtain the best values of these parameters but also, and more importantly, their posterior probability distribution: $p(\boldsymbol{\beta_M},\sigma,\boldsymbol{M} \mid D)$. Invoking Bayes' theorem, this posterior can be written as the product of a likelihood and a prior model:

$$p(\beta_M,\ \sigma,\ \mathbf{M} \mid D) \propto \mathbf{p}(D \mid \beta_M,\ \sigma,\ \mathbf{M})\pi\ (\beta_M,\ \sigma,\mathbf{M}). \tag{9}$$

The likelihood $p(D| \boldsymbol{\beta_M},\sigma,\boldsymbol{M})$ is the probability of observing the data given the model parameters $\boldsymbol{\beta_M}$, $\boldsymbol{\sigma}$, and $\boldsymbol{M}$. Since the form of the likelihood is determined by the general linear model of Eq. (8) and owing to the assumed normality of error distribution $\boldsymbol{\sigma}$, this likelihood can be expressed as a Gaussian:

$$p(D|\ \boldsymbol{\beta_M},\ \boldsymbol{\sigma},\ \mathbf{M}) = \prod_{i=1}^{n} N(y_i;\ \boldsymbol{x}(t_i)\ \boldsymbol{\beta_M},\ \boldsymbol{\sigma}) \tag{10}$$

The Bayesian formulation of the problem finally requires specifying the second term of Eq. (9), which is the prior distribution of the model parameters. By definition, one can write:

$$\pi\ (\beta_M,\ \sigma,\mathbf{M}) = \pi\ (\beta_M,\ \sigma \mid \mathbf{M})\pi(\mathbf{M}). \tag{11}$$

allowing to elicit the conditional prior $\pi\ (\boldsymbol{\beta_M},\sigma|\ \boldsymbol{M})$ and the model prior $\pi(\boldsymbol{M})$ separately.

The prior encodes the available information about possible values of the model parameters. In case reliable knowledge is not available, non-informative flat priors are the common choice. This is the alternative adopted in the rest of the paper. In more detail, the normal-inverse Gamma distribution has been implemented for $\pi\ (\boldsymbol{\beta_M},\ \sigma \mid \boldsymbol{M})$, while for the prior model structure $\pi(\boldsymbol{M})$, it has been assumed that the numbers of change-points are any nonnegative integers equally probable.

Given the chosen likelihood and prior models, for the posterior of the model parameters one obtains [31]:

$$p(\boldsymbol{\beta} M,\ \boldsymbol{\sigma},\ \nu,\ \mathbf{M}|\ D) \propto \prod_{i}^{n} N\big(y_i;\ \mathbf{x}_M(t_i)\boldsymbol{\beta}_M,\ \boldsymbol{\sigma}_\beta\big)\pi\ (\boldsymbol{\beta}_M,\ \boldsymbol{\sigma},\ \nu|\boldsymbol{M})\ \pi\ (\boldsymbol{M}) \tag{12}$$

Unfortunately, the posterior distribution $p(\boldsymbol{\beta_M},\ \boldsymbol{\sigma},\ \boldsymbol{M}|D)$ derived in Eq. (12) is analytically intractable and therefore it is necessary to make recourse to the typical Monte Carlo Chain Monte Carlo (MCMC) technique to generate realizations of random samples for posterior inference [41]. The MCMC sampling algorithm implemented to obtain the results reported in the following is a hybrid sampler that embeds a reversejump MCMC sampler (RJ-MCMC) into a Gibbs sampling framework. All the details about RJ-MCMC can be found in [42].

The aforementioned MCMC algorithm can generate a chain of N posterior samples $\{\boldsymbol{M}^{(i)},\boldsymbol{\beta_M}^{(i)},\boldsymbol{\sigma}^{(i)}\}\ i = 1,\ \cdots,\ N$. This output captures all the information essential for the inference of time series, including trends, seasonal variations, abrupt changes and noise.

In mathematical terms, the structure of the sampled model $\boldsymbol{M}^{(i)}$, such as time locations of the abrupt changes and seasonal harmonic orders, can be directly translated into the model's covariates $\boldsymbol{x}_{\boldsymbol{M}}^{(i)}\ (t)$ and their associated coefficients $\boldsymbol{\beta_M}^{(i)}$. Each sampled model $\boldsymbol{M}^{(i)}$ indeed gives one estimate of the time series dynamics $\boldsymbol{x}_{\boldsymbol{M}^{\omega}}^{(i)}\ \boldsymbol{\beta}_{\boldsymbol{M}}^{(i)}\ (t)$. Combining the individual estimates provides not only a final BMA estimate but also uncertainty measures, obtained averaging all the sampled models. In a similar way, it is also possible to provide the uncertainties calculated as a sample-based estimate of the variance of the models [31].

It should also be appreciated that, although each single model $\boldsymbol{M}^{(i)}$ is a piecewise function, the BMA global estimate, being the combination of all the individual models, can approximate nonlinear signals arbitrarily well. Moreover, since the covariates $\boldsymbol{xM}^{(i)}(t)$ and the model coefficients $\boldsymbol{\beta_M}$ are a combination of the individual elements of the trend and seasonal signals, they allow recovering the global trend and seasonal components.

To appreciate the potential of BEAST and its variants, it should not be overlooked that the sampled model structure permits also to provide an empirical distribution of the number of change-points (in addition to their localisation), a quantity the other techniques have to assume a priori. Allowing for different harmonic orders in different time intervals is another additional flexibility of the approach, distinguishing it from most existing algorithms, which typically are based on a pre-set, fixed order for the seasonal signal. Moreover, the chains $\{\tau_{k=1,\ldots,\ m(i)},\ \xi_{k=1,\ldots,\ p(i)}\}$ of the N sampled individual models indicate the exact time localisations, at which abrupt changes interleave the trends or seasonal components. From these chains, it is also possible to provide many statistics such as the probability that a change-point occurs at a certain time or within an interval, or the probability of observing an abrupt variation in a trend if another change-point has already occurred [31].

2.4. Some alternative indicators for time series analysis: AIC and HQC

The Bayesian approach to time series decomposition can be conceptualised as an alternative to the frequentist tradition, in which model selection criteria are quite consolidated. Two representative examples, the Akaike Information Criterion (AIC) and the Hannan-Quinn Criterion (HQC), have been implemented to provide alternatives for comparison. Both criteria are single number scores that can be used to select the model that is most likely to be the best for interpreting a given data set among a list of competing candidates. They are both cost functions to be minimised in the sense that a lower AIC or HQC score is better. They are also relative estimators, meaning that their numerical values are only useful in comparison with scores of other models for the same data set.

The Akaike information criterion, introduced in 1973, is the first model selection criterion that gained widespread acceptance [43]. It can be considered as an extension to the maximum likelihood principle, which is traditionally applied to estimate the parameters of a given model, whose structure and dimension have already been fixed. The qualitative advance of the AIC consists of combining estimation with structural and dimensional determination into a single procedure.

The conceptual basis of the AIC is rooted in information theory. Any statistical model, describing the process that generated the data, will almost never be exact. Consequently, some information is inevitably lost by using the model to represent the process. The AIC, by applying the Kullback–Leibler divergence [44], estimates the relative amount of information lost by a given model. The less information a model loses, the higher is the quality of that model. Therefore, the candidate model that minimises the information loss is to be chosen.

The most widely used version of the AIC criterion is:

$$AIC = n \cdot \ln(MSE) + 2k \tag{13}$$

Where MSE is the mean-squared error of the residuals, the differences between the data and the predictions of the models, n the number of entries in the database and k the number of the model parameters to be estimated. When the indicator is applied to the modelling of time indexed signals, as in the present work, n is the number of samples in the time series. Simple inspection of Eq. (13) reveals why, among a set of candidates, the preferred model is the one with the minimum AIC value. Indeed, AIC rewards goodness of fit as quantified by the MSE. However, it also includes a penalty term that is an increasing function of the number of model parameters. The penalty is of course introduced to discourage overfitting.

The Hannan-Quin criterion [45] is an alternative to (AIC). It can be written as:

$$HQC = n\cdot\ln(MSE) + 2k\ln(\ln(n)) \quad (14)$$

Where the meaning of the symbols is the same as in Eq. (13). The HQC imposes a larger penalty on complex models than the AIC for large samples. HQC is consistent and can avoid the overshooting of the number of parameters that can affect AIC. As the Hanna-Quinn criterion also the Takeuchi Information Criterion (TIC), the Focussed Information Criterion (FIC), the Deviance Information Criterion (DIC) and the Kashyap Information Criterion (KIC) are basically all descendants of AIC [46]. The differences among these criteria are immaterial in the context of the present work, because they would all provide basically the same results for the examples analysed in the following. In any case, a detailed discussion of these criteria properties can be found in [47].

3. Tests with synthetic data and benchmarking with known experimental signals

Even if the accuracy of BEAST has already been verified in various fields and mainly in the analysis of remote sensing of the earth, the specificity of tokamak data requires some dedicated validations with synthetic data. Of the main numerical examples analysed, the ones reported in the following relate to the two main experimental applications discussed in the rest of the paper: stationarity detection and the change in periodic or quasi-periodic components.

The top plot of Fig. 1 shows the potential of the method to decompose a particularly delicate situation: a relatively small change of trend, to which a quite strong oscillatory components is superimposed. This a pattern is quite common in tokamak time series in H mode, when many kinetic quantities, such as the density, are affected by oscillations due to ELMs. From the plot, representative of a wide range of tests, it appears quite clearly that the identification of the trend and its variation are in no way compromised by the presence of even strong seasonal components. Indeed, the Bayesian algorithm manages to identify quite accurately the time location of the change in the trend, with a probability of almost unity, even if it is submerged by a larger oscillation of varying amplitude.

The bottom plot of Fig. 1 is meant to show the capability of the Bayesian ensemble technique to detect abrupt variations in seasonal components. Indeed the characteristics of the oscillatory instabilities such as ELMs and, to a lesser degree sawteeth, can change quite rapidly during a discharge, depending on various conditions. Detecting accurately the time locations of such changes in a statistically sound way is quite important for both the understanding of the physics and the control of the scenarios. Various numerical tests have shown that the locations of the transitions in periodic seasonal components are determined with very good accuracy and again with a probability close to one.

To dissipate any doubt that the routines could have difficulties when dealing with experimental data, Fig. 2 and Fig. 3 report the results for two well-known real life data sets, which are considered a good validation for any decomposition algorithm.

Fig. 2 reports the evolution of the total monthly number of airlines passengers in the US between Jan. 1949 and Dec. 1960 [48]. This is considered a "classic" benchmark for the detection of constant trends when immersed in strong seasonal components. The fit of the trend, reported in green, proves again the robustness of method to identify drifts in presence of disturbances and oscillations.

A completely different complex time series is the case of Fig. 3, which reports the internet traffic from UKERNA (United Kingdom education and research networking association) and represents aggregated traffic in the United Kingdom academic network backbone. It was collected between 19 November 2004, at 09:30 h and 27 January 2005, at 11:11 h. The time series was recorded every 5 min [49]. The critical aspect of

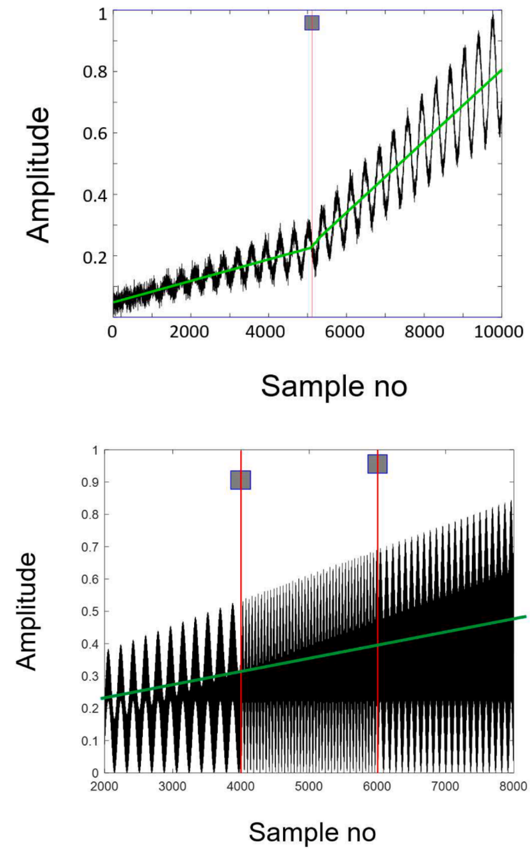

Fig. 1. Top: an oscillatory behaviour reminiscent of ELMs' signature superimposed on a change of trend: the synthetic data is obtained by adding a sinusoid with increasing amplitude and normal noise to two segments with different gradients. The green curve is the trend. The method decomposes the time series perfectly, identifying exactly the location of the gradient change with transition probability equal to 1. Bottom: three oscillatory components simulating different ELM regimes on a constant trend: the highest frequency in the middle is 8 times the one in the first interval and twice that in the third. Again the green curve is the trend. Also the changes of the oscillatory frequencies are detected correctly with high probability. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

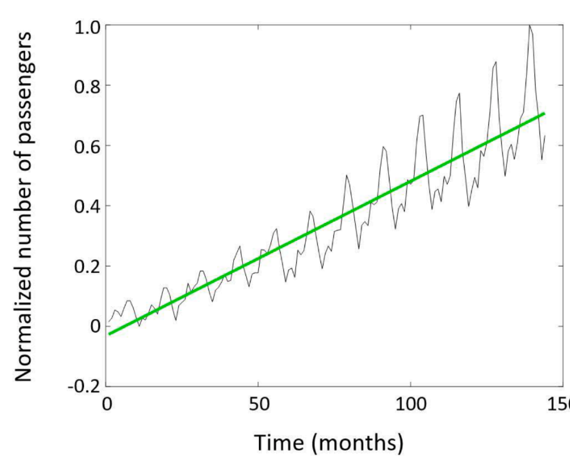

Fig. 2. The evolution of the monthly total of airlines passengers in the US between Jan. 1949 and Dec. 1960 a standard time series benchmark for stationarity tests. The method can correctly detect the trends (the fitted curve is in green). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

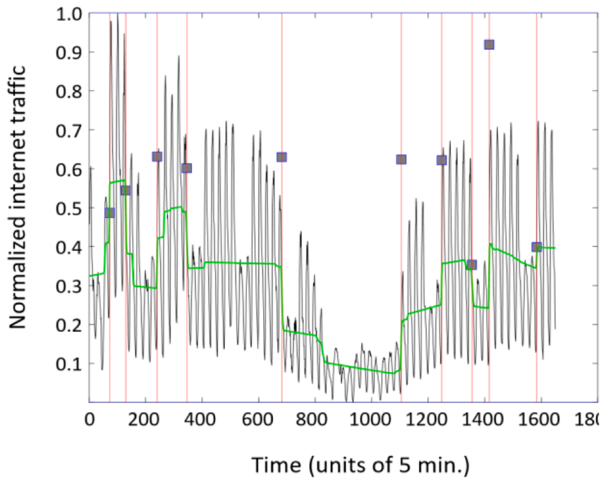

Fig. 3. Internet traffic time series with several daily, weekly and seasonal patterns; e.g. the traffic decrease in late December in the dataset is due to Christmas season (https://doi.org/10.1111/j.1468-0394.2010.00568.x). The method provides a good separation of the different regions. The abrupt changes are marked by the vertical red lines, while the fitted curve in shown. The squares indicate the probabilities of the transitions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

this example is that the data contains several daily, weekly and seasonal patterns. The final fit reported in green shows that the implemented version of BEAST manages to properly identify the slow trends, stationary phases, and the abrupt changes in the oscillatory components. At first sight, this signal seems as challenging as decomposing various quantities in tokamak H mode regime with well-formed Type I ELMs and sawteeth. However, a discussed in Section 4, the properties of sawteeth and ELMs are even more complex to deal with than this type of signal.

4. Results with experimental signals of JET tokamak

In this section, various JET signals are analysed, to show the potential of the Bayesian ensemble method to help addressing three important aspects of time series in tokamaks: the presence of actual stationary phases, the decomposition of sawteeth, which are quasi-periodic oscillations, and the changes of ELM properties. An overview of the discharges discussed in the rest of the section is provided in Appendix A. When appropriate and possible, the focus has been on high power discharges in DD and DT, given their relevance for the next generation of tokamaks. Unless otherwise specified, the abrupt changes reported in the following are those estimated to have a probability equal or higher than 0.4. This choice has been motivated mainly by the objective of obtaining conservative results and avoiding false detections. Of course the choice of this hyper-parameter could be optimised on the basis of the signal nature and the objectives of the analysis. It is also important to note that the code has been run autonomously. With regard to the MCMC settings, the length of individual chains has been set to 20,000 iterations with the first 1000 being discarded as burn-in samples. Chains have been thinned by retaining only every third sample. Moreover, a total of three such chains have been run in parallel, thinned, and lastly merged into a final chain of samples. According to [37], chains mix rapidly and, in most cases, reach stability after only a few thousand iterations. Therefore, this setup is sufficient to ensure convergence. The long length of the chains also removes the effect of initial model choices on the chains. The alpha level of the credible intervals has been set to 95 %. The probability of selecting a re-sampling proposals has been set to 10 %.

4.1. Stationarity assessment and trends

To show the potential of the Bayesian ensemble approach to identify trend and stationary phases, the signal of the plasma current for many JET discharges has been analysed. A quite involved example is reported in Fig. 4. In this discharge, the shape of the plasma current is particularly involved. There is indeed a change of slope during the ramp-up, an overshoot and then two stationary periods before the final ramp-down. The algorithm manages to identify all the trends and stationary phases, locating the transition times between them very accurately. The bottom plot of Fig. 4 shows how this task is much more challenging for the AIC and HQC criteria, which try to identify a single best model and not an average one.

The results reported in Fig. 4 are fully representative of the systematic application of the Bayesian ensemble approach to well controlled quantities, such as the plasma current, the toroidal field and the input power. The algorithms have no major difficulties in identifying slow trends and stationary periods. They manage to locate the transition points accurately and with good confidence. As shown in Section 3, superimposed strong seasonal components or high levels of noise do not compromise the performances of the codes.

The capability of properly identifying slow trends is an extremely relevant competitive advantage when it is important to detect stationarity or the lack of it. Indeed, as already mentioned, the sustainment of long pulses requires the early detection and avoidance of even small drifts, which can compromise the performances if not the stability of the plasma configurations. In this perspective, the algorithms have been applied to various global kinetic quantities. One example is reported in Fig. 5, in which the plasma density of a JET high performance discharge is analysed [50]. This is an important quantity because ITER and DEMO, to achieve their target performances, are expected to operate stably at or even above the Greenwald limit. Such a regime of operation is not easily achieved and maintained. As can be easily derived from the plot of

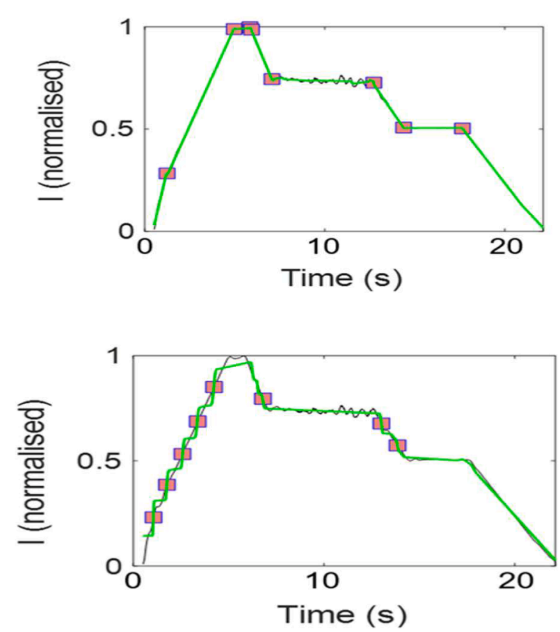

Fig. 4. Top: Bayesian decomposition of the plasma current time series for JET discharge #96947. A quite complicated case, including a current overshoot and a staircase ramp-down has been chosen. The y axis is the normalised amplitude and the x axis represents time. In black the original signal and in green the fit obtained with BEAST. The square indicates the times of the abrupt changes. The quality of the decomposition is apparent. Bottom: the decomposition obtained with the AIC and Hannan-Quinn approximations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

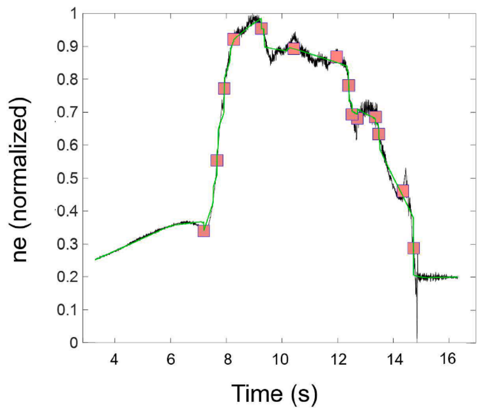

Fig. 5. Bayesian decomposition of the line integrated plasma density time series for JET discharge #99948. The y-axis is the normalised amplitude of the interferometer vertical chord 3 and the x axis represents time. In black the original signal and in green the fit obtained with BEAST. The squares indicate the times of the abrupt changes. The quality of the decomposition is quite good. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

Fig. 5, the plasma density varies and drifts significantly in the high power phase of the experiment even if the plasma is well below the Greenwald density regime envisaged for the proto-reactors.

Another example is the total neutron yield of the DTE2 record discharge #99,950 presented in Fig. 6. Again, in this case various trends are clearly picked up by the algorithm, including a slow but evident reduction of the plasma performances. Really stationary conditions do not last for much more than one second. Since the variations are evident on time scales of a few seconds, such drifts would certainly be inacceptable in the discharges produced in the next generation of devices.

The examples reported in this subsection are in line with the observation reported by the CICLOP working group that high performance phases in tokamaks tend to be quite short lived. Indeed application of the BEAST suite of codes to the JET contribution to the database reported in [2] reveals that it is very difficult to find phases at reasonable performances, whose main global kinetic quantities can be considered stationary for more than a couple of seconds. Even if a systematic analysis has not been possible yet, for lack of validated data, it is quite clear that more subtle quantities, such as the concentration of extrinsic and intrinsic impurities, can be even more delicate to control.

On JET, it has also not been possible to systematically stabilise the current profile in the hybrid scenario for intervals much longer than the current diffusion time. This fact is illustrated in the discharge of Fig. 7, which is simply another typical example of trends that would not be acceptable in the reactor, because it would compromise the configuration on time scales much shorter than the length of the discharge [51, 52]. The main point here is that, since the current diffusion can be quite slow, the Bayesian ensemble approach is particularly useful because, as shown also with the numerical examples, it can detect even weak drifts.

Even if the previous assessment of limited stationarity is factually correct, some cautionary considerations are in pace. Firstly, JET was not designed to perform a systematic a programme for the investigation of long pulse operation. For example various ancillary systems, the additional heating schemes in particular, and the inertially cooled divertor pose significant limitations in this respect. Indeed, even the discharges in the database reported in [2] were not explicitly devoted to this objective. Consequently, a more comprehensive assessment of the issues related to long pulse operation will have to be provided by other devices, such as EAST, KSTAR, WEST and JT-60SA, better equipped to perform this type of studies. The results of the present section have therefore to be interpreted simply as a confirmation of the conclusions reported in [2], but obtained with sound and automated statistical techniques.

4.2. Seasonality: sawteeth

Seasonality components are very important elements of time series. In magnetically confined thermonuclear plasmas, various instabilities, such as sawteeth and ELMs, produce quasi-periodic oscillations of the acquired signals, which certainly belong to the class of seasonal variations. However they both present some peculiarities requiring careful investigation.

The most prominent feature of sawteeth is the clear difference between two time scales within each cycle: an interval of slow rise and a fast crash. The oscillations in the core temperature that they cause therefore resembles a sequence of sawteeth, from which they derive their name. The shape of this seasonal component is very far from the simple sine wave. Even if the period is quite regular, in principle it should be verified to what extent the Bayesian ensemble method can properly identify this type of seasonality.

In the case of sawteeth, the application of the Bayesian averaging version implemented by BEAST produces mixed results, as shown in Fig. 8. Indeed, given the regularity and high amplitude of this type of instability, the algorithm normally manages to identify the changes in

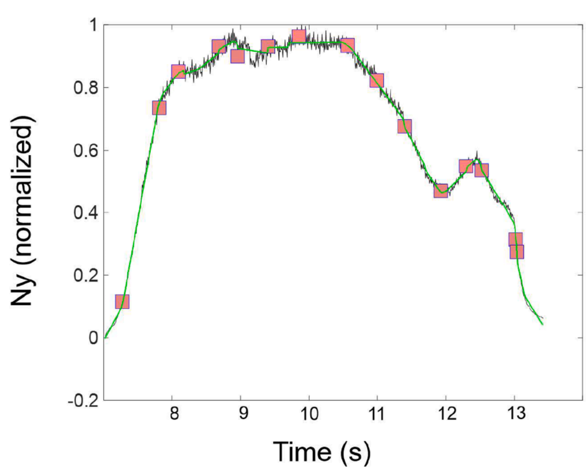

Fig. 6. Bayesian decomposition of the total neutron yield for JET record discharge #99950. The y axis is the normalised amplitude and the *x* axis represents time. In black the original signal and in green the fit obtained with BEAST. The squares indicate the times of the abrupt changes, whose probability is estimated to be higher than 0.3. The quality of the decomposition is apparent. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

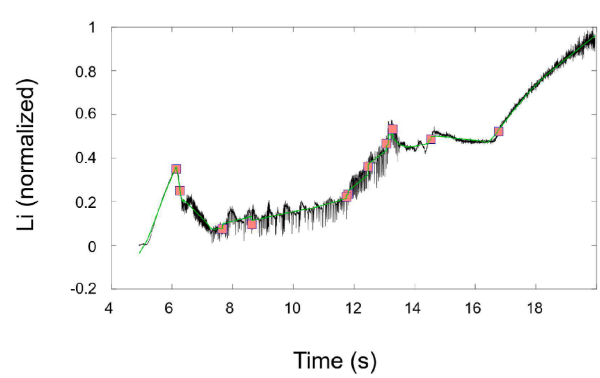

Fig. 7. Bayesian decomposition of the internal inductance for JET discharge #96,368. The y axis is the normalised amplitude and the x axis represents time. In black the original signal and in green the fit obtained with BEAST. The squares indicate the times of the abrupt changes. The quality of the decomposition is evident. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

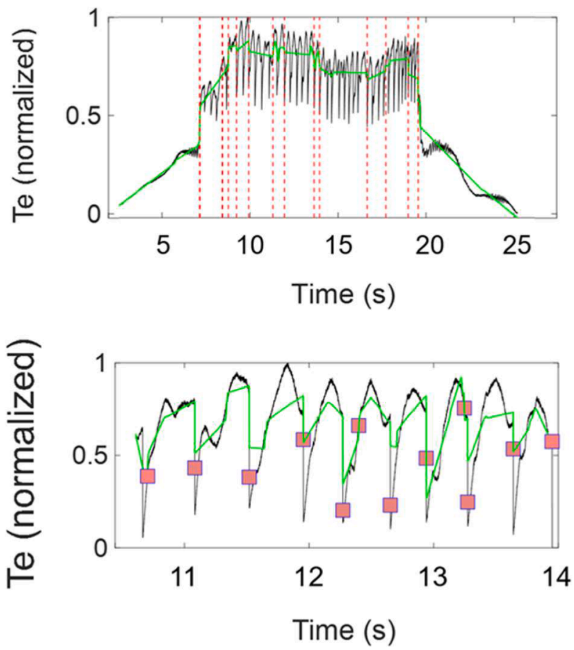

Fig. 8. Bayesian decomposition of the central electron temperature series for JET discharge #89,995. Top: the main variations in the instability behaviour, corresponding to the vertical dashed red lines, are well identified and localised in time. The green line represents the trends. Bottom: the fit of the individual sawteeth cycles, the green line, is not satisfactory, due to the quite involved spectral content of the signal. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

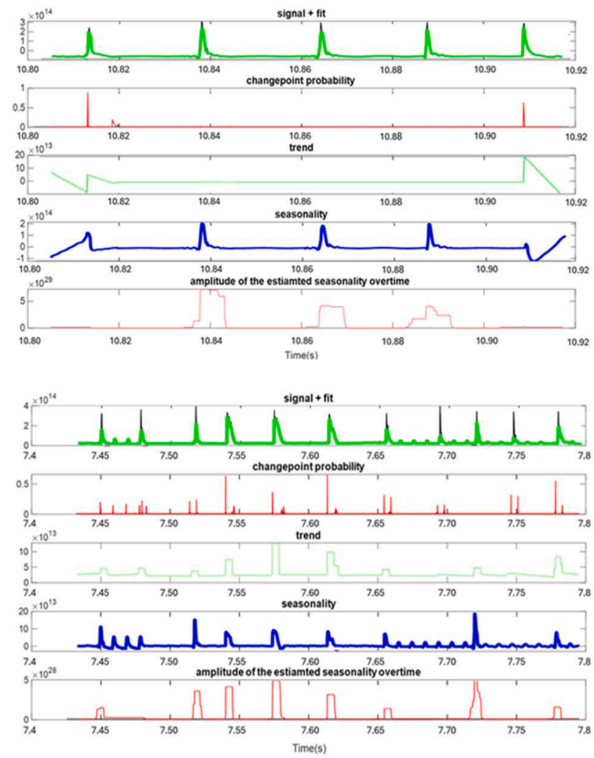

Fig. 9. Bayesian decomposition of the Be signal of JET discharge #94871.

the sawteeth behaviour quite accurately. The time points of the variations in the central electron temperature, measured by the ECE diagnostic, are well localised with a high probability. On the other hand, the individual oscillations are not very well modelled, as can be seen in the bottom plot of Fig. 8. The intervals pertaining to the individual sawteeth are delimited quite accurately but the reconstruction of their overall from is not satisfactory. This is due to the shape of the signal, whose oscillations do not resemble sinusoidal curves and are very asymmetric. This behaviour is very difficult to model with the harmonic functions described in SubSection 2.2.

4.3. Seasonality: ELMs

The decomposition of the ELM reveals a similar limitation of the Bayesian ensemble implementation of the BEAST code.

Indeed, the BEAST codes have been conceived for applications in remote sensing of the atmosphere. In such a context, the seasonal elements present a quite precise periodicity and can be properly modelled with the harmonic series introduced in Section 2.2. Consequently, if the ELMs are almost exactly periodic, the suite of codes manages to model them quite accurately, as shown in the top plots of Fig. 9. The oscillating component is reproduced well and no trend variation is detected. However, the frequency of ELMs can be much less precise and vary significantly during the H-mode phase of the discharges. The difficulty of the Bayesian ensemble implementation to deal with such a situation is shown in the bottom plot of Fig. 9. In such a case, the seasonality elements and their locations in time are picked up reasonably well but there is clear and unphysical spill over on the trend component. This of course is not a realistic description of the real ELM dynamic evolution.

Even if the full Bayesian decomposition is not completely satisfactory for physical studies, the present set of codes allows detecting quite accurately the changes in the seasonal components due to ELMs and their locations in time. A representative example is reported in Fig. 10. The discharge comprises very stationary phases of ELMs oscillations but from one interval to another the frequency and behaviour of the ELMs change significantly. As can be seen from the figure, the Bayesian ensemble approach is very good at detecting the moments of frequency changes and also the general trend of the signal evolution. Moreover, this case confirms that the algorithm does not produce spurious changepoints when the behaviour of an oscillating signal is periodic on top of a stationary background. The localisation of the ELMs and their correspondence to actual variations of the plasma state are confirmed by the plots of Fig. 11, which report the changes in the edge integrated density corresponding to the detected patterns in the BeII signal.

5. Considerations on computational times and resources

As already explained in Section 2, the performed analysis is based on Bayesian regression, which provides a natural framework for variable and model selection. Multiple competing models are synthetized and this enormous model space is explored, by using the Monte Carlo Markov Chain (MCMC) sampling algorithm [41]. A posterior probability of being true can be assigned to each model. MCMC ensures fast computation for a problem that involves even an extremely large number of candidate models. The variability of the number of models depends on the structure of the investigated time series.

For the examples from JET diagnostics analysed in this paper, the computation time varies in the range [0.17s-1.03 s] per 1000 samples of the time series. The fastest analysis is obtained for cases like that one

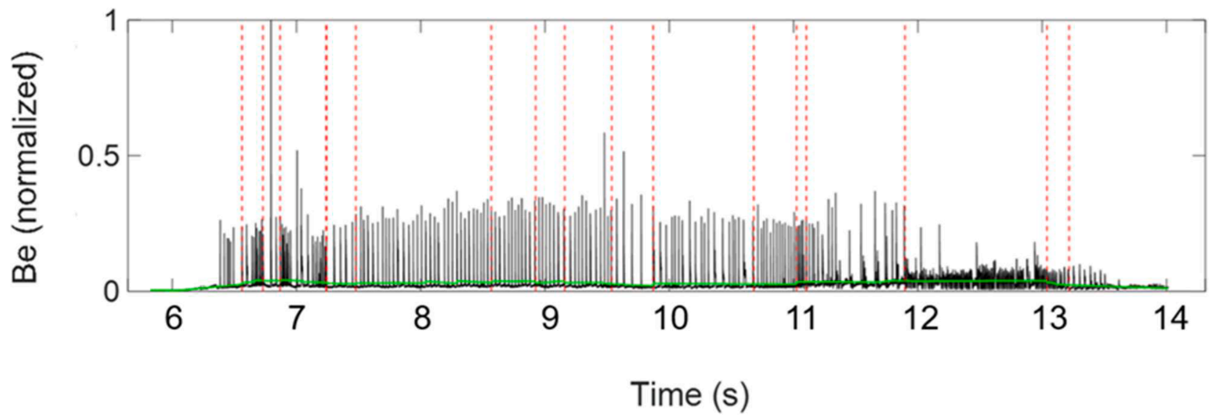

Fig. 10. The seasonality component for JET discharge #94871, in which the oscillations of the ELMs are almost exactly periodic for several quite long time intervals. However the ELM frequency and behaviour vary quite abruptly depending on the phase of the pulse. The BEAST algorithm manages to identify the time location of these abrupt changes quite accurately. A different choice of the threshold would allow detecting also smaller variations in the ELM regimes at about 7 s and 11.3 s.

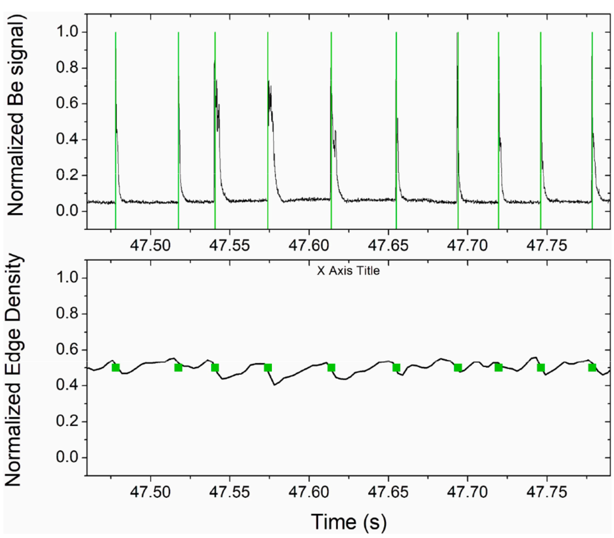

Fig. 11. Correspondence between the ELMs and the variations of the edge plasma density in JET discharge #94871. The green lines and squares indicate the change-points detected by the BEAST code using the BeII signal. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.).

presented in Fig. 4, recorded with a time sampling of ~14 ms, which needs 0.72 s. The most complex time series structure in that one related to the evolution of ELMs in the discharge #94871 (Fig. 9). The sampling time is 0.1 ms and the time series comprises 3×10^6 samples. In this case a total computation time was of 288 s. The reported computation time values have been obtained running the codes in MATLAB 2023on a laptop PC with the following characteristics: processor Intel Xeon w5–2465X, 5.10 GHz with 64 GB of RAM memory. Indeed, the BEAST routines do not require any special allocation of memory. Therefore, from a computational perspective the method is already suitable for the offline investigation of large databases.

It is worth pointing out that JET signals are deemed to be not too different from the ones that will be produced by the next generation of devices. Indeed in larger machines the time scales are expected to be slower and therefore sets of 1000 samples should provide quite good resolution of many phenomena also for control room analysis. On the other hand, the present implementation of the algorithms is not compatible with real-time deployment, with the aim of providing feedback for plasma control systems (see next section).

6. Discussion, conclusions and future developments

In this work, the Bayesian ensemble approach to time series decomposition has been applied to various tokamak signals. Indeed, the accurate and robust determination of seasonal components, trends, abrupt changes and noise in the data of plasma diagnostics is important per se and as a preliminary step to more advanced types of studies. Developing an automatic and generally accepted methodology to undertake this task would be very beneficial for the entire community.

The tests performed with both synthetic signals and experimental data have proven the strengths and limitations of the BEAST suite of codes, implementing a very advanced version of Bayesian ensemble time-series decomposition. The technique is very effective in detecting trends and slow drifts, confirming the difficulties of achieving good stationarity of the kinetic quantities in JET high performance discharges. Less satisfactory has been the application to the seasonal components. Indeed, the method, as presently implemented, can properly model only simple periodic oscillations and has difficulties coping with the details of sawteeth and with ELMs of changing frequencies. However even the available routines can detect quite accurately the time locations of the changes in sawteeth frequency and ELM regimes, even if the decomposition of these seasonal components leaves a lot to be desired from the perspective of physics modelling. Moreover the code can be run autonomously, so avoiding the vagaries of substantial human intervention

The obtained results suggest the following considerations and lines of future activity. Firstly, the suite of codes will have to be rewritten for a full assessment of its potential and to adjust it to the scientific requirements of the fusion community. For example, at the moment the working of the reversible-jump MCMC sampler is quite obscure (it is not possible to determine how the sampler explores different model structures or the number of change-points or the seasonal harmonic order with the iterations). Also BEAST outputs only the best model and not the posterior probabilities of the competing solutions, which would allow assessing which structures the data favour and by what margins.

In terms of physics applications, since ITER and DEMO are expected to produce discharges lasting one hour or longer, it seems imperative to extend the stationary phases of present day experiments [53]. More attention should be paid to preserve high performance conditions and to monitor quantities such as the radiation or the current profile, which cannot be allowed to vary even slowly on the time scales of the next generation of devices [54]. For the offline analysis of these aspects, the present suite of codes implements adequate sufficiently performing solutions (the prospects of real time implementations are discussed later).

On the front of seasonality and oscillations [8], the research to handle ELMs is already very active. In this direction, what seems more pressing is a methodological development of the approach proposed in the present work. The inability of the Bayesian ensemble implementation to properly decompose ELMs of varying frequency should be remedied. Indeed, proper understanding and modelling of the ELMs requires a better discrimination between the seasonal components and the slow trends. A promising approach would consist of replacing the present sine and cosine modelling of seasonality trends with different basis functions, for example splines or wavelets. Indeed, the constraints on the periodicity of the solutions should be relaxed to properly model complex signals such as those generated by ELM instabilities. Such an upgrade would also remedy the difficulties of the present implementation in modelling the details of the sawteeth.

Another important, even if not easy, upgrade of the technique would involve relaxing the assumption that the four components of the time series are purely additive, which is not necessarily true for all the signals generated by complex systems. In this perspective, however, it is not clear to what extent such an upgrade of the methodology would be necessary to model time series and to understand the dynamics of tokamak plasmas. A deeper and even more demanding methodological development relates to the calculation of the probabilities. In the present version of the algorithm they have to be considered relative, in the sense that they quantify the comparative probability of a certain feature (change points for example). Providing calibrated probability, with an accurate and absolute meaning, would be extremely useful. Of course such an improvement is quite difficult, because the Bayesian ensemble method is based on averaging over the candidate models, whose individual prior probabilities are not easy to determine absolutely.

On the front of significant practical applications and as already alluded to, a certainly very important additional line of investigation relates to the development of tools deployable in feedback. Indeed, in discharges of reactor relevant machines, any small drift away from the chosen operational point should be detected as soon as possible to undertake the necessary remedial actions and recover the required performances. In addition, variations in the behaviour of the ELMs and sawteeth (or any other macroscopic instability) should also be notified to the control system very timely, to avoid the onset of dangerous situations, potentially leading to disruptions. Unfortunately, the present version of the algorithms requires computational times incompatible with feedback applications. However for phenomena that can be analysed with of the order of about 1000 samples, the prospects are quite good. Two possible lines of attack seem possible. On the one hand, one could investigate alternative implementations of the codes, with acceleration measures to reduce substantially the computational times. This should be achievable because the method is inherently parallelisable. Consequently on more preforming and possibly dedicated hardware, it is expected that computational times of the order of 1 ms for 1000 samples should be feasible. For the determination of slow trends and the detection of abrupt changes in sawteeth or ELM behaviour, this should be achievable relatively easily. The relevance of such an upgrade for control and handling of off-normal events is obvious [55,56]. However, to properly decompose the seasonal component of course an effective algorithm should be identified first. If its parallelisation proved to be difficult, an alternative direction could be to use the offline suite of codes to train advanced machine learning tools for feedback application. Bayesian physics informed neural networks have been developed [57] but they have not been conceived for real time implementations as a first priority and therefore significant developments in this direction would also be required.

An additional very important potential application involves the monitoring of diagnostics (and ancillary systems in general). Indeed, the detected signals can present drifts or abrupt changes also due to unwanted variations in the hardware of the corresponding measurement equipment. Of course it is very difficult if not impossible to disentangle the plasma from the diagnostic effects on the basis of purely plasma measurements. On the other hand, the proposed technique could be applied to specific signals suitable to monitor the status of the diagnostics. Given their quality, the developed Bayesian ensemble algorithms could help detect drifts and degradations in various instruments and reduce the downtime of many diagnostics, a not negligible problem in most tokamaks. The proposed tools could therefore be deployed in ambitious programmes of predictive maintenance [58,59]. Of course also this type of application would require appropriate adjustments of the technique to the specificities of fusion systems. In particular, it is to be expected that particularly challenging problems, such as the discrimination of the slow trends in the magnetic signals due to simultaneous plasma small changes and neutron induced voltage drifts in the cables, will require further assessment and the combination of more diagnostics.

CRediT authorship contribution statement

Michela Gelfusa: Resources, Methodology, Investigation, Funding acquisition. **Teddy Craciunescu:** Validation, Software, Methodology, Formal analysis, Conceptualization. **Riccardo Rossi:** Writing – review & editing, Validation, Software. **Andrea Murari:** Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

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 Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them. One of the authors (TC) is grateful for financial support received from the Romanian National Nucleu Program LAPLAS VII – contract no 30N/2023.

Appendix A. main global plasma quantities of the discharges discussed in the main text

Pulse 89995.

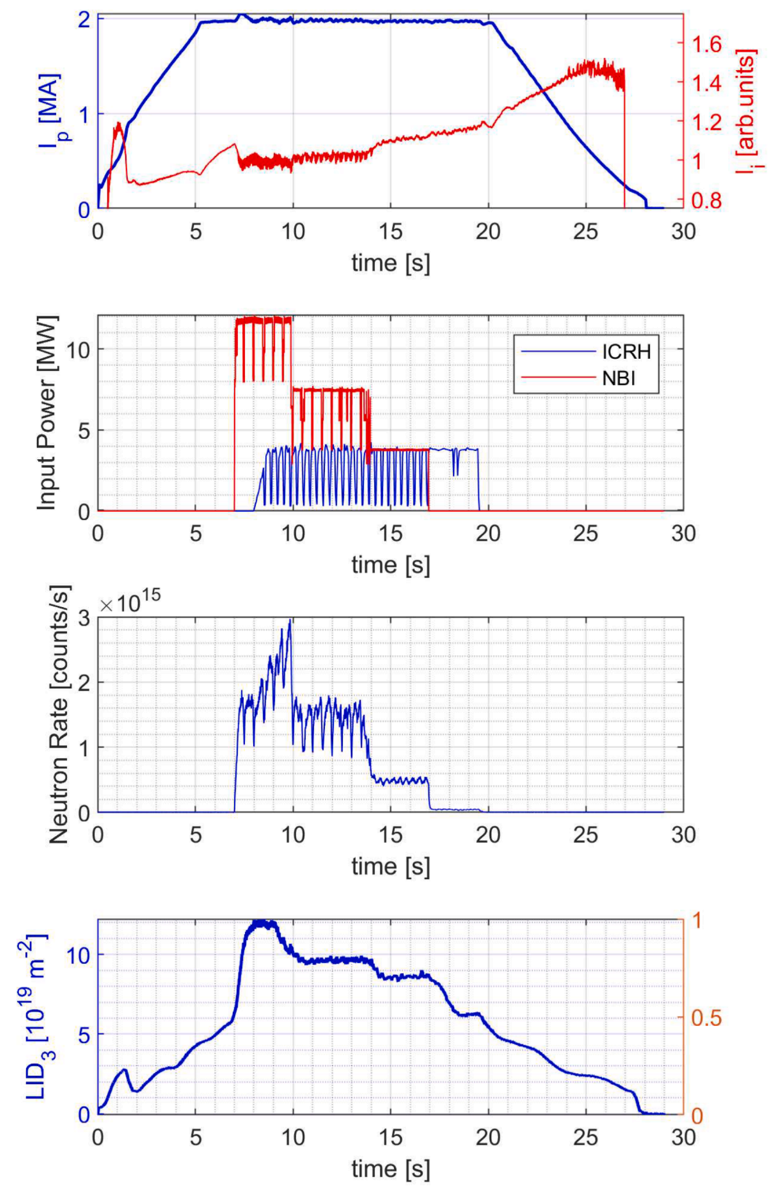

Pulse 94871.

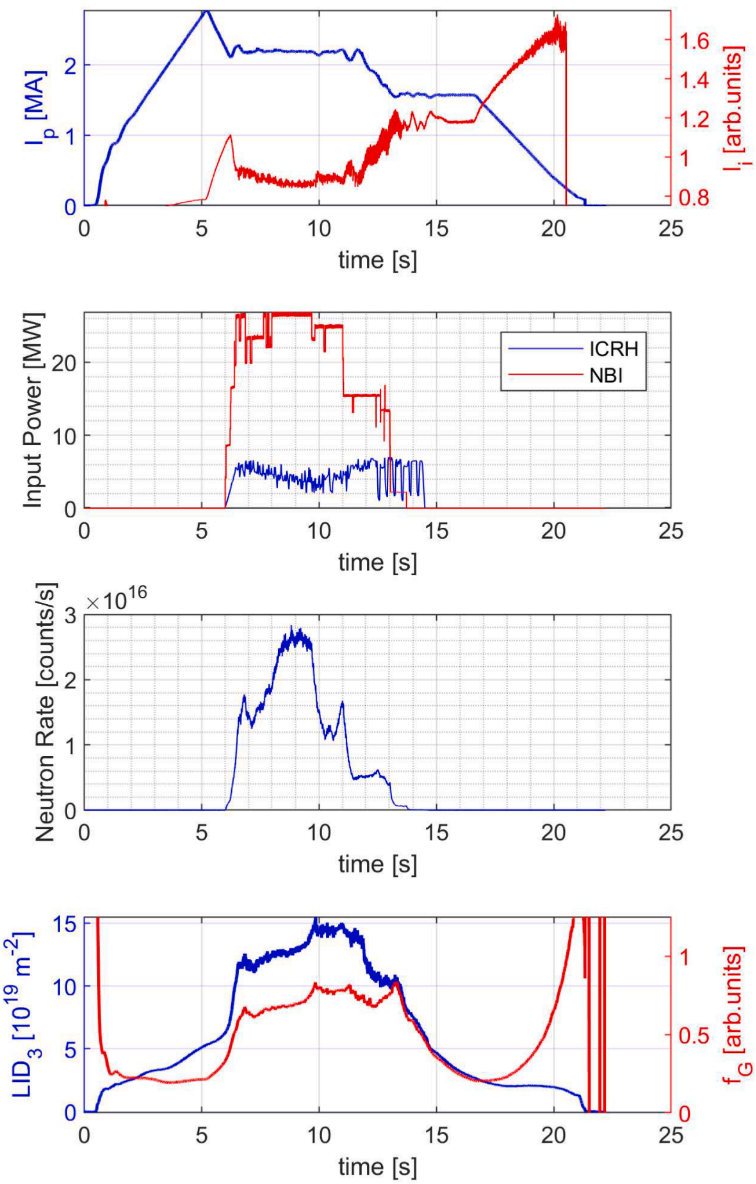
I_p [MA]
l_i [arb.units]
time [s]
Input Power [MW]
ICRH
NBI
time [s]
×10^16
Neutron Rate [counts/s]
time [s]
LID_3 [10^19 m^-2]
f_G [arb.units]
time [s]

Pulse 96368.

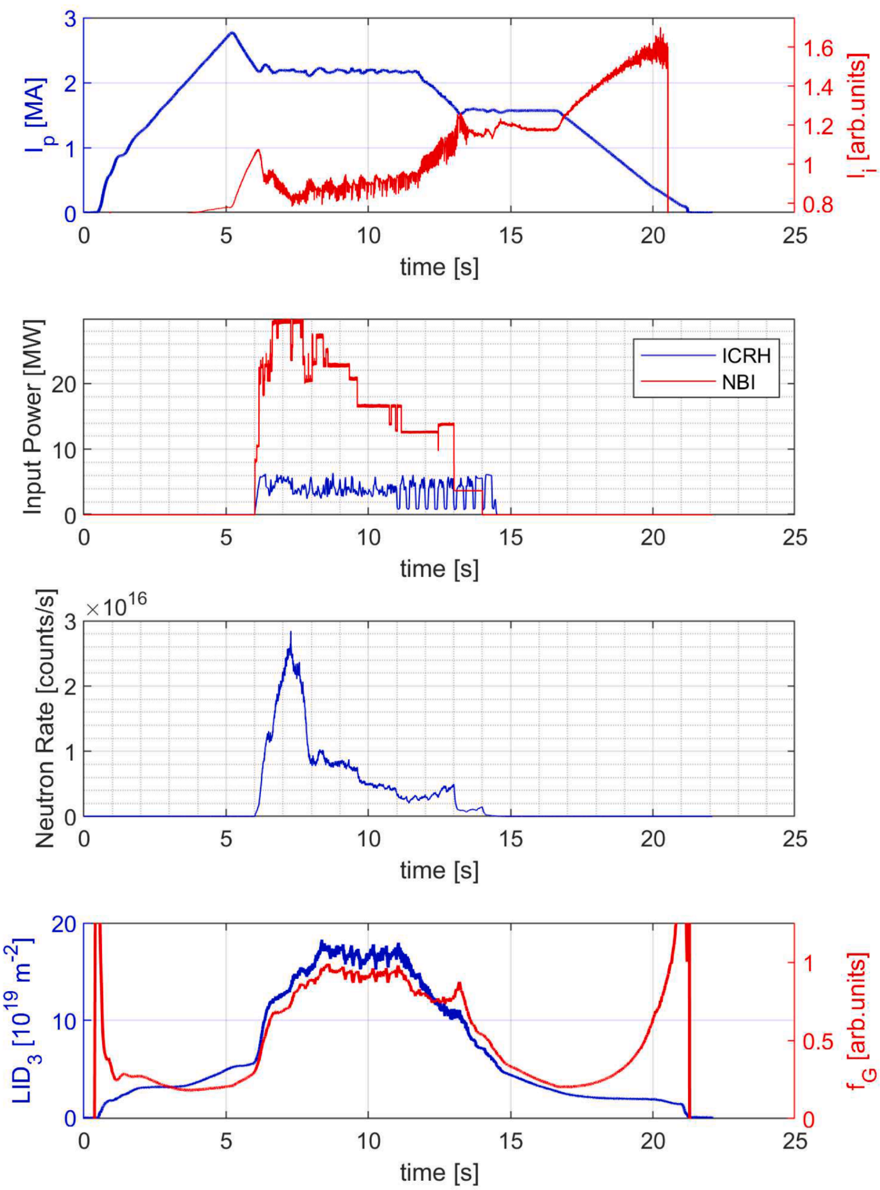
I_p [MA]
l_i [arb.units]
time [s]
Input Power [MW]
ICRH
NBI
time [s]
×10^16
Neutron Rate [counts/s]
time [s]
LID_3 [10^19 m^-2]
f_G [arb.units]
time [s]

Fusion Engineering and Design 220 (2025) 115318

Pulse 96947.

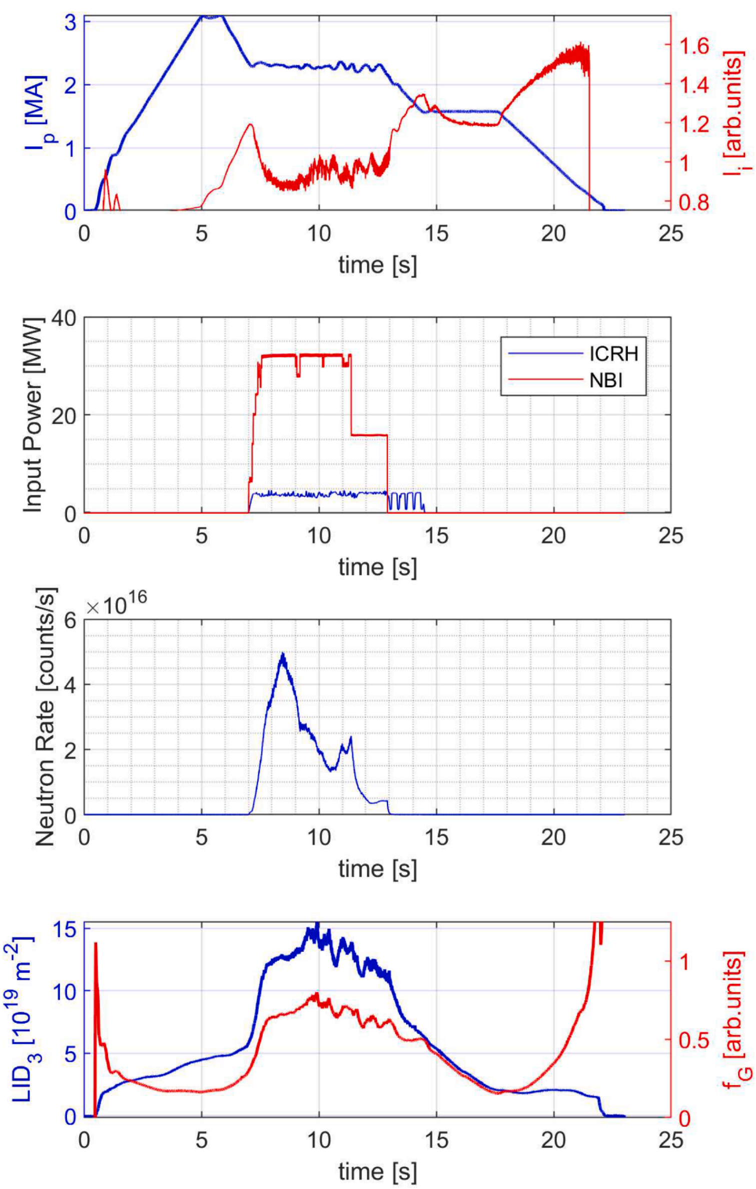

Pulse 99948.

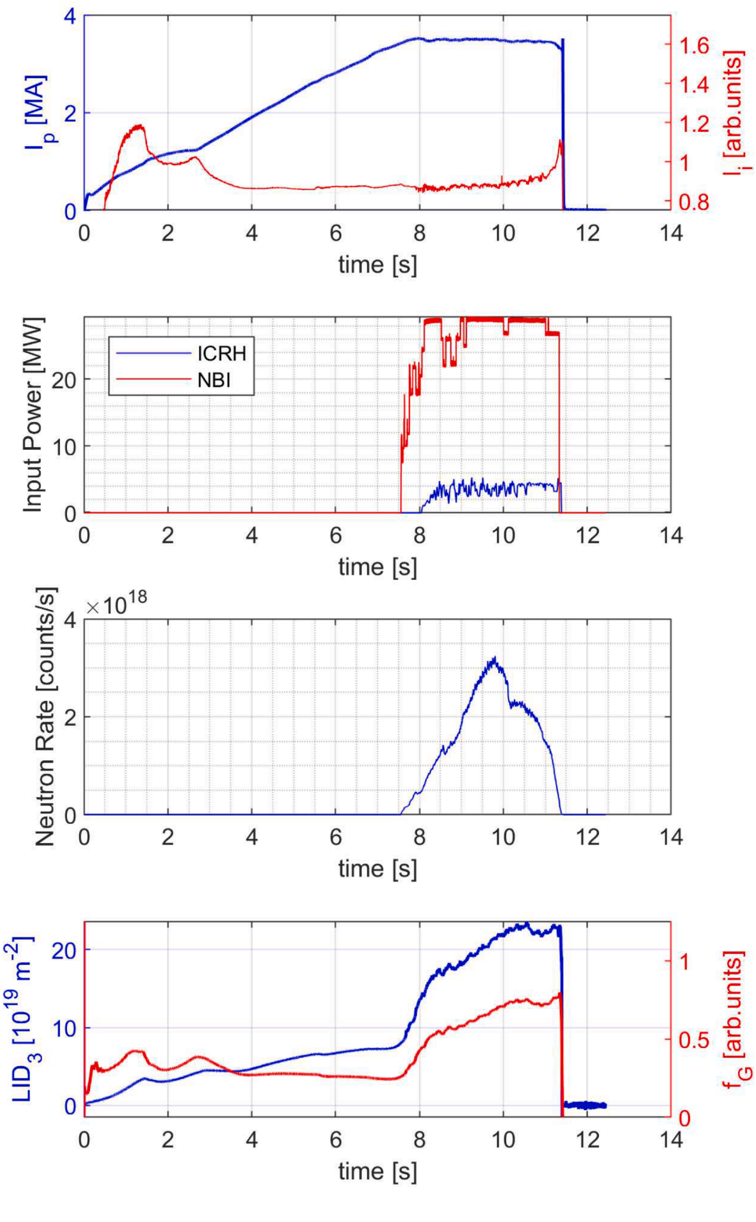

I_p [MA]
l_i [arb.units]
time [s]
Input Power [MW]
ICRH
NBI
time [s]
×10^18
Neutron Rate [counts/s]
time [s]
LID_3 [10^19 m^-2]
f_G [arb.units]
time [s]

Pulse 99950.

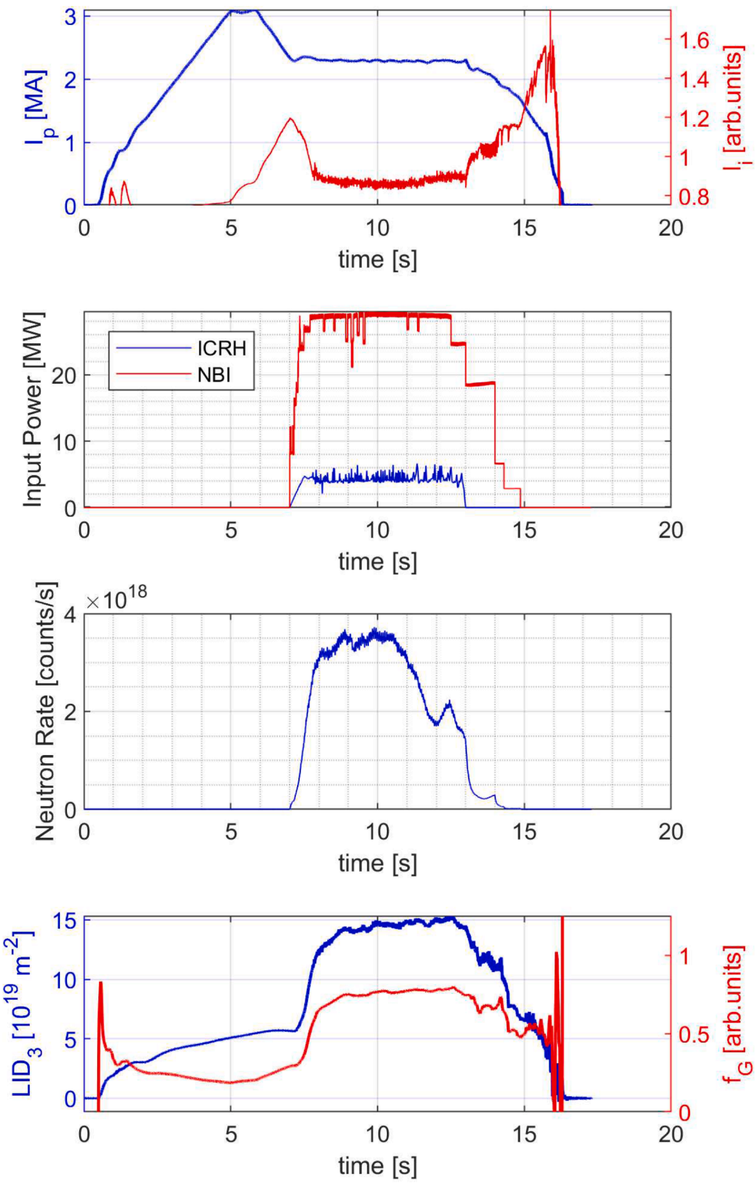

Data availability

Data will be made available on request.

References

[1] J. Wesson, D.J. Campbell, Tokamaks, 4th Edition, Oxford University Press, 2011. ISBN: 0198509227 9780198509226.

[2] X. Litaudon, et al., Nucl. Fusion 64 (2024) 015001, https://doi.org/10.1088/1741-4326/ad0606.

[3] F. Wagner, et al., Regime of mproved confinement and high beta in neutralbeam-heated divertor discharges of the ASDEX tokamak, Phys. Rev. Lett. 49 (1982) 1408.

[4] H. Zohm, Edge localized modes (ELMs), Plasma Phys. Control Fusion. 38 (1996) 105.

[5] A.W. Leonard, Edge-localized-modes in tokamaks, Phys. Plasmas. 21 (2014) 090501.

[6] G. Federici, et al., Overview of the DEMO staged design approach in Europe, Nucl. Fusion 59 (2019) 066013.

[7] M. Siccinio, et al., DEMO physics challenges beyond ITER, Fusion Eng. Des. 156 (2020) 111603.

[8] E. Viezzer, M.E. Austin, M. Bernert, K.H. Burrell, P. Cano-Megias, X. Chen, D. J. Cruz-Zabala, S. Coda, M. Faitsch, O. Février, L. Gil, C. Giroud, T. Happel, G. F. Harrer, A.E. Hubbard, J.W. Hughes, A. Kallenbach, B. Labit, A. Merle, H. Meyer, C. Paz-Soldan, P. Oyola, O. Sauter, M. Siccinio, D. Silvagni, E.R. Solano, Prospects of core–edge integrated no-ELM and small-ELM scenarios for future fusion devices, Nucl. Mater. Energy 34 (2023) 101308. doi.org/10.1016/j.nme.2022.101308.

[9] E. Lerche, et al., Nucl. Fusion 57 (2017) 036027.

[10] M. Lennholm, et al., Demonstration of effective control of fast-ion-stabilized sawteeth by electron-cyclotron current drive, Phys. Rev. Lett. 102 (2009) 115004.

[11] J. Graves, et al., Plasma Phys. Control Fusion. 57 (2015) 014033.

[12] M. Lennholm, et al., Closed loop sawtooth period control using variable ECCD injection angles on Tore Supra, Fusion Sci. Technol. 55 (2009) 45–55.

[13] R. Kube, et al., Near real-time streaming analysis of big fusion data, Plasma Phys. Control Fusion. 64 (2022) 035015, https://doi.org/10.1088/1361-6587/ac3f42.

[14] R.M. Churchill, et al., Deep convolutional neural networks for multi-scale time-series classification and application to tokamak disruption prediction using raw, high temporal resolution diagnostic data, Phys. Plasmas. 27 (2020) 062510.

[15] X. Yang, X. Chen, K. Sun, C. Xiong, D. Song, Y. Lu, L. Huang, S. He, X. Zhang, A wavelet transform-based real-time filtering algorithm for fusion magnet power signals and its implementation, Energies 16 (2023) 4091, https://doi.org/10.3390/en16104091.

[16] R. Coelho, D. Alves, C. Silva, Magnetohydrodynamic and turbulence activity analysis in the ISTTOK tokamak using empirical mode decomposition, Rev. Sci. Instrum. 77 (2006) 10F512, https://doi.org/10.1063/1.2220474.
[17] M. Gelfusa, R. Rossi, A. Murari, Causality detection and quantification by ensembles of time delay neural networks for application to nuclear fusion reactors, J. Fusion Energy 43 (2024) 7, https://doi.org/10.1007/s10894-024-00398-8.
[18] A. Murari, et al., Application of transfer entropy to causality detection and synchronization experiments in tokamaks, Nucl. Fusion 56 (2016) 026006, https://doi.org/10.1088/0029-5515/56/2/026006.
[19] J. Svensson, et al., Current tomography for axisymmetric plasmas, Plasma Phys. Control Fusion. 50 (2008) 085002.
[20] R. Fischer, A. Dinklage, Integrated data analysis of fusion diagnostics by means of the Bayesian probability theory, Rev. Sci. Instrum. 75 (2004) 4237.
[21] A. Dinklage, R. Fischer, J. Svensson, Topics and methods for data validation by means of Bayesian probability theory, Fusion Sci. Technol. 46 (2004) 355.
[22] J. Svensson, A. Werner, Large scale Bayesian data analysis for nuclear fusion experiments, in: IEEE International Symposium on Intelligent Signal Processing, 2007, p. 1.
[23] R. Fischer, C.J. Fuchs, B. Kurzan, et al., Integrated data analysis of profile diagnostics at ASDEX upgrade, Fusion Sci. Technol. 58 (2010) 675.
[24] B.Ph. van Milligen, T. Estrada, E. Ascasíbar, et al., Integrated data analysis at TJ-II: the density profile, Rev. Sci. Instrum. 82 (2011) 073503.
[25] A. Pavone et al. *Machine learning and Bayesian inference in nuclear fusion research: an overview* 2023 Plasma Phys. Control Fusion 65 053001 **DOI:**10.1088/1361-6587/acc60f.
[26] A. Murari, Fabio Pisano, Jesús Vega, Barbara Cannas, A. Fanni, Socorro González, Michela Gelfusa, Massimiliano Grosso, J.E.T Contributors, Extensive statistical analysis of ELMs on JET with a carbon wall, Plasma Phys. Control Fusion 56 (2014) 114007.
[27] A. Murari, et al., On efficiency and interpretation of sawteeth pacing with on-axis ICRH modulation in JET, Nucl. Fusion 57 (2017) 126057, https://doi.org/10.1088/1741-4326/aa87e7.
[28] J. Vega, et al., Disruption prediction with artificial intelligence techniques in tokamak plasmas, Nat. Phys. 18 (7) (2022) 741–750, https://doi.org/10.1038/s41567-022-01602-2.
[29] A. Murari, et al., A control oriented strategy of disruption prediction to avoid the configuration collapse of tokamak reactors, Nat. Commun. 15 (1) (Mar. 2024) 2424, https://doi.org/10.1038/s41467-024-46242-7.
[30] Teddy & Craciunescu, A. Murari, Detection of changes in the dynamics of thermonuclear plasmas to improve the prediction of disruptions, Nonlinear. Dyn. 111 (2022) 1–15, https://doi.org/10.1007/s11071-022-08009-x.
[31] Kaiguang Zhao, Michael A. Wulder, Tongxi Hu, Ryan Bright, Qiusheng Wu, Haiming Qin, Yang Li, Elizabeth Toman, Bani Mallick, Xuesong Zhang, Molly Brown, Detecting change-point, trend, and seasonality in satellite time series data to track abrupt changes and nonlinear dynamics: a Bayesian ensemble algorithm, Remote Sens. Environ. 232 (2019) 111181, https://doi.org/10.1016/j.rse.2019.04.034.
[32] D.G. Denison, Bayesian Methods for Nonlinear Classification and Regression, John Wiley & Sons, 2002.
[33] B.M. Rankin, J. Meola, M.T. Eismann, Spectral radiance modeling and Bayesian model averaging for longwave infrared hyperspectral imagery and subpixel target identification, IEEE Trans. Geosci. Remote Sens. 55 (2017) 6726–6735.
[34] D.G. Denison, C.C. Holmes, B.K. Mallick, A.F. Smith, Bayesian Methods For Nonlinear Classification and Regression, John Wiley & Sons, 2002 (Vol. 386).
[35] B.M. Rankin, J. Meola, M.T. Eismann, Spectral radiance modeling and Bayesian model averaging for longwave infrared hyperspectral imagery and subpixel target identification, IEEE Trans. Geosci. Remote Sens. 55 (12) (2017) 6726–6735.
[36] J. Reiche, A. De Bruin, D. Hoeckman, J. Verbeselt, M. Herold, A Bayesian approach to combine landsat and alos palsar time series for near real-time deforestation detection, Remote Sens. 7 (5) (2015) 4973–4996.
[37] K. Zhao, D. Valle, S. Popescu, X. Zhang, B. Mallick, Hyperspectral remote sensing of plant biochemistry using Bayesian model averaging with variable and band selection, Remote Sens. Environ. 132 (2013) 102–119.
[38] T. Zhou, S.C. Popescu, A.M. Lawing, M. Eriksson, B.M. Strimbu, P.C. Bürkner, Bayesian and classical machine learning methods: a comparison for tree species classification with LiDAR waveform signatures, Remote Sens. 10 (1) (2017) 39.
[39] M.C. Kennedy, A. O'Hagan, Bayesian calibration of computer models, J.R. Statist. Soc. B 63 (2001) 425–464.
[40] G. Claeskens, N.L. Hjort, Model Selection and Model Averaging, Cambridge University Press, 2008.
[41] J.S. Speagle. (2019) *A conceptual introduction to Markov Chain Monte Carlo methods*, arXiv:1909.12313 https://doi.org/10.48550/arXiv.1909.12313J.
[42] P.J. Green, Reversible jump Markov chain monte carlo computation and Bayesian model determination, Biometrika 82 (4) (1995) 711–732.
[43] H. Akaike, Information theory and an extension of the maximum likelihood principle, in: B.N. Petrov, F. Csáki (Eds.), 2nd International Symposium on Information Theory, Tsahkadsor, Armenia, USSR, September 2-8, 1971, Akadémiai Kiadó, Budapest, 1973, pp. 267–281. Republished in Kotz, S.; Johnson, N. L., eds. (1992), *Breakthroughs in Statistics*, vol. I, Springer-Verlag, pp. 610–624.
[44] S. Kullback, R.A. Leibler, On information and sufficiency, Ann. Math. Stat. 22 (1951) 79–86.
[45] E.J. Hannan, B.G. Quinn, The Determination of the order of an autoregression, J. R. Stat. Soc. Series B (41) (1979) 190–195.
[46] K.P. Burnham, D.R. Anderson, Model Selection and Multi-Model Inference: A Practical Information-Theoretic Approach, 2nd edition, Springer, New York, 2002.
[47] J. Ding, et al., Model selection techniques – an overview. IEEE Signal Processing Magazine 1810.09583, 2018 arXiv.org.
[48] https://www.kaggle.com/datasets/rakannimer/air-passengers).
[49] P. Cortez, M. Rio, M. Rocha, P. Sousa, Multi-scale Internet traffic forecasting using neural networks and time series methods, Expert Syst. 29 (2) (2012) 143–155, https://doi.org/10.1111/j.1468-0394.2010.00568.x.
[50] L. Garzotti, et al., Development of high current baseline scenario for high deuterium-tritium fusion performance at JET, in: Preprint: 2023 IAEA Fusion Energy Conf. (London, 16–21 October 2023), 2023. EX/7-3(available at, https://conferences.iaea.org/event/316/papers/27856/files/10983-Paper-IAEA-2023-V5.pdf).
[51] J. Hobirk, et al., The JET hybrid scenario in Deuterium, Tritium and Deuterium-Tritium, Nucl. Fusion 63 (2023) 112001, https://doi.org/10.1088/1741-4326/acde8d.
[52] M. Fitzgerald, et al., Stability analysis of alpha driven toroidal Alfvén eigenmodes observed in JET deuterium-tritium internal transport barrier plasmas, Nucl. Fusion 63 (2023) 112006, https://doi.org/10.1088/1741-4326/acee14.
[53] A. Bécoulet, G.T. Hoang, The technology and science of steady-state operation in magnetically confined plasmas, Plasma Phys. Control Fusion 50 (2008) 124055, https://doi.org/10.1088/0741-3335/50/12/124055.
[54] S.H. Kim, et al., in: Exploration of Long-Pulse and Steady-State Operations in ITER IAEA Technical Meeting Devoted to Long-Pulse Operation of Fusion Devices (IAEA Headquarters, Vienna, Austria, 14–16 November 2022), 2022 (available at, https://conferences.iaea.org/event/258/.).
[55] J. Vega, et al., Disruption prediction with artificial intelligence techniques in tokamak plasmas, Nat. Phys. 18 (7) (2022) 741–750, https://doi.org/10.1038/s41567-022-01602-2.
[56] A. Murari, et al., A control oriented strategy of disruption prediction to avoid the configuration collapse of tokamak reactors, Nat. Commun. 15 (1) (2024) 2424, https://doi.org/10.1038/s41467-024-46242-7.
[57] R.M. Neal, Bayesian Learning for Neural Networks. Lecture Notes in Statistics, Springer, 1996. ISBN-13978-0387947242.
[58] F. Raza, AI for Predictive Maintenance in Industrial Systems, 2023, https://doi.org/10.13140/RG.2.2.27313.35688.
[59] A. Theissler, et al., Predictive maintenance enabled by machine learning: use cases and challenges in the automotive industry, Reliab. Eng. Syst. Saf. (2021), https://doi.org/10.1016/j.ress.2021.107864.