

Array magnetics modal analysis for the DIII-D tokamak based on localized time-series modelling

K E J Olofsson¹, J M Hanson¹, D Shiraki¹, F A Volpe¹, D A Humphreys²,
R J La Haye², M J Lanctot², E J Strait², A S Welander², E Kolemen³
and M Okabayashi³

¹ Applied Physics and Applied Mathematics, Columbia University, New York City, NY 10027-6900, USA

² General Atomics, PO Box 85608, San Diego, CA 922186-5608, USA

³ Princeton Plasma Physics Laboratory, PO Box 451, Princeton, NJ 08543-0451, USA

E-mail: ko2319@columbia.edu

Received 14 February 2014, revised 29 May 2014

Accepted for publication 16 June 2014

Published 14 July 2014

Abstract

Time-series analysis of magnetics data in tokamaks is typically done using block-based fast Fourier transform methods. This work presents the development and deployment of a new set of algorithms for magnetic probe array analysis. The method is based on an estimation technique known as stochastic subspace identification (SSI). Compared with the standard coherence approach or the direct singular value decomposition approach, the new technique exhibits several beneficial properties. For example, the SSI method does not require that frequencies are orthogonal with respect to the timeframe used in the analysis. Frequencies are obtained directly as parameters of localized time-series models. The parameters are extracted by solving small-scale eigenvalue problems. Applications include maximum-likelihood regularized eigenmode pattern estimation, detection of neoclassical tearing modes, including locked mode precursors, and automatic clustering of modes, and magnetics-pattern characterization of sawtooth pre- and postcursors, edge harmonic oscillations and fishbones.

Keywords: inference methods, magnetohydrodynamics, microinstabilities

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

1. Introduction

Magnetic probes arranged in linear Mirnov arrays in the toroidal and poloidal directions constitute a basic diagnostic system in tokamaks [1]. These arrays typically pick up non-axisymmetric deformations of the nominally axisymmetric toroidal plasma as they rotate past the arrays. Many magnetohydrodynamic (MHD) instabilities in tokamak plasmas exhibit a global near-rigid spatial structure. This means their spatial pattern can be inferred from the array signal time series. In general, the problem to be solved is a simultaneous multivariate probe array decomposition in both frequency and space (modal shape). The most reliable classifier of MHD modes is the toroidal period

number n . Less reliable is the poloidal period number m . This is because eigenmodes may be composed of multiple m -number components with an equilibrium-dependent and radially varying relative mixture. For a certain class of resistive MHD modes, tearing modes (TMs) [2], the m -number may be a reliable classifier. TMs are radially localized modes and can be associated with a certain resonant $q = m/n$ flux surface in the plasma. See figure 1 for an illustration of tokamak geometry and a non-axisymmetric localized TM perturbation. Other MHD modes, such as ideal kinks, are of more radially global character and may not be well described by a single m [3]; instead it may be the case that their poloidal spectrum depends strongly on the radial coordinate within the plasma. Figures 1(a) and (b) combined provide definitions of




Figure 1. Depiction of the DIII-D tokamak geometry. (a) shows a cut-away drawing of the DIII-D tokamak vessel with an $m/n = 3/2$ TM structure. The TM should be thought of as embedded within a nominally axisymmetric plasma. Non-axisymmetric perturbations such as the $3/2$ can be detected by the poloidal magnetic field fluctuations picked up as the modes rotate toroidally. The toroidal direction is denoted by ϕ . The pick-up direction b_θ for the magnetic probes is tangential to the wall and indicated by the arrows in the bottom of (b). Magnetic fluctuations are recorded as db_θ/dt . The geometrical poloidal angle θ is here defined from the machine toroidal axis centre. The coordinate R goes along the horizontal plane. $R = 0$ is on the torus major axis. The coordinate Z measures vertical elevation, with $Z = 0$ ($\theta = 0, \pi$) coinciding with the torus midplane.

the poloidal and toroidal angular coordinates θ and ϕ . The radial coordinate within the plasma can be thought of as the magnitude of the θ -angled vector in figure 1(b).

It is important to develop improved techniques to extract information from the magnetics array fluctuations. Both (i) plasma physics offline analyses and (ii) real-time (online) control decision can benefit from more accurate and robust signal processing. The ability to disentangle the magnetics array data into discrete ‘source terms’ of distinct classes, say m/n , has direct applications in real-time tokamak control. As an example, the triggering, aiming, and modulation of millimetre-wave power sources (e.g. gyrotrons) in neoclassical tearing mode (NTM) control systems can conceptually be based on the signature of a particular mode’s existence, $m/n = 3/2$ or $2/1$ predominantly [4–6]. This type of control needs to have a low false-positive trigger rate for it to be more efficient than a pre-emptive strategy with respect to the average required gyrotron power.

Another motivation for improved magnetics array analysis toolkits is to develop early warning indicators for mode-locking events. Mode-locking usually involves a saturated $2/1$ mode that slows down, locks and brings the plasma rotation to a halt [1]. The braking of the mode rotation is due to induced electromagnetic drag from the external wall and/or the interaction with residual error fields. This often has a disruptive effect on the tokamak plasma. Sometimes this can also happen without a rotating precursor phase. The mode is then said to be ‘born’ locked. An accurate early warning may allow a disruption-mitigation system to take more effective action; either a ‘soft-landing’/disruption-mitigation strategy [7, 8], or even some type of disruption avoidance. An example of the latter may be to apply magnetic perturbations to align the mode for mitigation using millimetre-wave driven currents [9]. Reliable signal processing methods that can extract and classify discrete modes with associated frequencies, amplitudes and phases could serve as a building block in such a system.

Fast Fourier transform (FFT)-based cross-coherence of a pair of toroidally separated magnetic probes is typically used to assess fast MHD activity in tokamaks [10, 11]. The DIII-D tokamak magnetic diagnostics is detailed in [11] including an outline of a typical FFT-based analysis. Array magnetics analyses based on singular value decompositions (SVDs) [12] are found in [13–17]. A sparse-approximation approach to MHD classification is found in [18, 19] (in particular focused to Alfvén eigenmodes). Multi-band Kalman-filter-based approaches are developed in [20, 21]. None of these previous studies attempts time-series modelling in the sense that is attempted here. The approach described in this work is based on subspace techniques which are known to be effective for line-spectrum estimation [22].

This work presents a new array analysis technique for poloidal field magnetic probes. The idea is to model a ‘low-rank’ signal that explains the array correlations. The method employs a combination of signal-processing techniques: (i) stochastic subspace identification [23–25] (SSI) with (ii) operational modal analysis (OMA) [26, 27] and (iii) least-squares estimation. The method can be regarded as comprised of two stages: (a) feature extraction and (b) classification. Stage (a) consists of (i) through (iii) introduced above. Stage (b) includes an assortment of data analysis tools that aids the interpretation and visualization of the raw output of stage (a).

This paper is organized as follows. Section 2 details the algorithms that constitute the new method. Two numerical examples in section 2.4 illustrate the benefit of the new method compared with two established techniques. Section 3 outlines the use of the developed code on DIII-D magnetic probe array data. Section 4 concludes the paper.

2. Array analysis method

The set of algorithms detailed in this section will be referred to as *eigspec* for brevity. This also happens to be the name

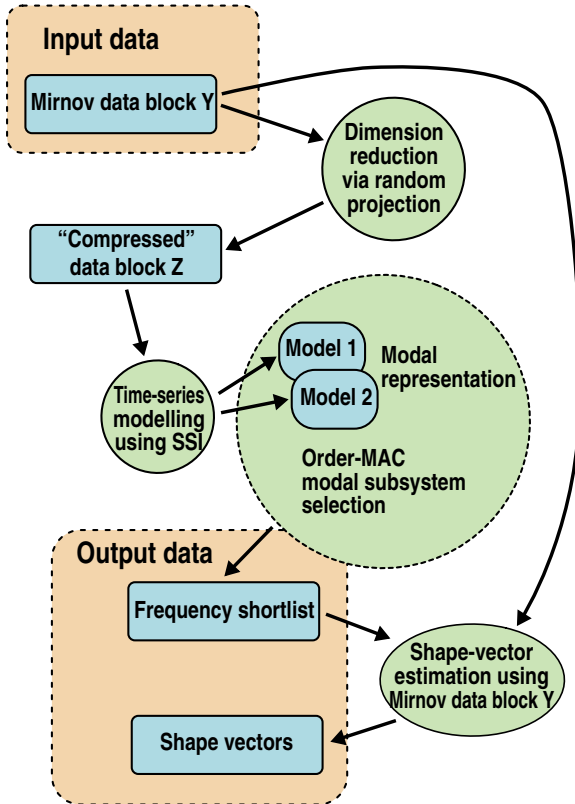


Figure 2. Flowchart of the feature extraction part of the `eigspec` analysis code. The method consists of a data channel-dimensionality reduction step followed by time-series modelling, frequency selection, and modal shape-vector estimation.

of the code deployed for data analysis usage at the DIII-D tokamak. In this section, the symbol m will exclusively denote the number of magnetic probes being analysed. It does not refer to the poloidal period number as it does elsewhere. The sub-matrix selection colon notation is as follows. For a matrix M , $M(a_1 : a_2, b_1 : b_2)$ means the sub-matrix spanning rows a_1 through a_2 and columns b_1 through b_2 . A colon by itself means that all available indices are used at that position. The imaginary unit is denoted by ι , $\iota^2 = -1$. A hatted quantity \widehat{X} explicitly denotes that X is an estimate.

2.1. Feature extraction

The feature-extraction process can be subdivided into three main steps: (i) time-series modelling, (ii) combined frequency decomposition and mode-selection and (iii) shape estimation. Step (iii) is only necessary if step (i) is computed on a reduced-dimension version of the full set of magnetic probes. If the full set of magnetic probes is used in step (i), then step (ii) would also achieve (iii) as a direct by-product. The reason for using a dimension reduction of the signal is computational complexity as will become clear in section 2.1.3 below. The process is illustrated in the flowchart in figure 2. In the next three sections steps (i) through (iii) will be detailed.

2.1.1. Time-series modelling. Let $\mathcal{Y}_N = \{\mathbf{y}(k)\}_{k=1}^N$ be a block of magnetic probe data where $\mathbf{y}(k) \in \mathbb{R}^{m \times 1}$ is an

m -dimensional vector sample indexed by discrete time k . The components of $\mathbf{y}(k)$ are distinct magnetic probe samples at various toroidal and poloidal locations, at time k . For this block of time-series data a discrete-time state-space model of the following form

$$\mathbf{x}(k+1) = A_r \mathbf{x}(k) + K_r e(k) \quad (1a)$$

$$\mathbf{y}(k) = C_r \mathbf{x}(k) + e(k) \quad (1b)$$

is to be estimated. Signal model (1a) and (1b) has a state $\mathbf{x}(k) \in \mathbb{R}^{r \times 1}$ and the estimation procedure assumes that $e(k)$ is a zero-mean sequence of white-noise random variables with some unknown covariance. Note that the time-series state-space dimension is r . In what follows, the method parameters denoted by f , p are respectively called the future and past horizons. The estimation starts by forming a data matrix D_{fp} with a certain shift structure from the block \mathcal{Y}_N :

$$D_{fp} = \begin{bmatrix} \mathbf{y}(1) & \mathbf{y}(2) & \dots & \mathbf{y}(N-p-f+1) \\ \mathbf{y}(2) & \mathbf{y}(3) & \dots & \mathbf{y}(N-p-f+2) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}(p+f) & \mathbf{y}(p+f+1) & \dots & \mathbf{y}(N) \end{bmatrix}. \quad (2)$$

Now the first mp rows of D_{fp} are denoted by D_p and the last mf rows are denoted by D_f . D_{fp} is thus partitioned as

$$D_{fp} = \begin{bmatrix} D_p \\ D_f \end{bmatrix} \in \mathbb{R}^{(mp+mf) \times (N-p-f+1)} \quad (3)$$

and the idea is to find a linear predictor of D_f given D_p . A basic SSI method goes as follows [24].

$$R_{fp} = D_f D_p^T \in \mathbb{R}^{mf \times mp} \quad (4a)$$

$$R_{fp} = U \Sigma V^T \quad (4b)$$

$$X_r = \Sigma_r^{1/2} V_r^T D_p. \quad (4c)$$

Here Σ is a diagonal matrix, $\Sigma_r = \Sigma(1:r, 1:r)$, and $V_r = V(:, 1:r)$. At this point the columns of X_r could be used as the state sequence $\mathbf{x}(k)$ in equations (1a) and (1b). This means linear regression can be used to obtain the matrices (A_r, K_r, C_r) . But if only the modes and array shapes of equations (1a) and (1b) are of interest, only (A_r, C_r) are required. An alternative method is therefore to use the same SVD as in equation (4b) but instead calculating

$$O_r = U_r \Sigma_r^{1/2} \quad (5)$$

where $U_r = U(:, 1:r)$. Here O_r can be interpreted as the ‘extended observability matrix’ [24, 25] which has a particular shift structure that allows direct extraction of (A_r, C_r) :

$$D_1 A_r = D_2 \quad (6a)$$

$$\widehat{C}_r = O_r(1:m, :) \quad (6b)$$

$$\widehat{A}_r = (D_1^T D_1)^{-1} D_1^T D_2 \quad (6c)$$

where

$$D_1 = O_r(1:(mf-m), :) \quad (7a)$$

$$D_2 = O_r((m+1):mf, :). \quad (7b)$$

In `eigspec`, method (5) through (7b) is implemented as the default calculation.

2.1.2. Modal analysis and selection. Time-series model (1a) and (1b) can be represented in modal coordinates $\mathbf{q} = W^{-1}\mathbf{x}$ as

$$\mathbf{q}(k+1) = \tilde{A}_r \mathbf{q}(k) + \tilde{K}_r \mathbf{e}(k) \quad (8a)$$

$$\mathbf{y}(k) = \tilde{C}_r \mathbf{q}(k) + \mathbf{e}(k) \quad (8b)$$

where \tilde{A}_r is diagonalized ($A_r = W\tilde{A}_r W^{-1}$), $\tilde{K}_r = W^{-1}K_r$ and $\tilde{C}_r = C_r W$. The modal subsystems are collected from (i) real-valued eigenvalues and (ii) complex-conjugate pairs of eigenvalues. The eigenvalues λ_i are found on the diagonal of \tilde{A}_r .

In `eigspec`, the i th mode shape vector is, nominally, the complex vector found as a column of \tilde{C}_r , associated with a complex-conjugate mode i . It is only the actual shape vector if no random projection technique was used. See next section 2.1.3 for the details on this. Regardless of whether random projection was used or not, the next important problem to solve is to determine which linear system modes to keep and which to discard. To do so, a coherence-like quantity is introduced that evaluates the similarity of two complex-valued vectors \mathbf{v} , \mathbf{w} . This quantity is denoted as $\text{MAC}(\mathbf{v}, \mathbf{w})$ (acronym for modal assurance criteria) in order to connect with the modal analysis literature where its original usage is to compare finite-element computational modal shapes with experimentally obtained shapes [28, 29]. In `eigspec`, the MAC is used both for obtaining non-spurious features from time-series data, and for clustering and comparison of different non-spurious shape vectors as will be seen in section 2.2.

The MAC is defined as

$$\text{MAC}(\mathbf{v}, \mathbf{w}) = \frac{(\mathbf{v}^\dagger \mathbf{w})(\mathbf{w}^\dagger \mathbf{v})}{(\mathbf{v}^\dagger \mathbf{v})(\mathbf{w}^\dagger \mathbf{w})} = \frac{|\mathbf{v}^\dagger \mathbf{w}|^2}{\|\mathbf{v}\|_2^2 \|\mathbf{w}\|_2^2} \quad (9)$$

where $(\cdot)^\dagger$ denotes complex-conjugate transpose and $\|\mathbf{v}\|_2^2 = \mathbf{v}^\dagger \mathbf{v}$. $\text{MAC}(\mathbf{v}, \mathbf{w})$ is real-valued in $[0, 1]$. It holds $\text{MAC}(r_v \exp(i\theta_v) \mathbf{v}, r_w \exp(i\theta_w) \mathbf{w}) = \text{MAC}(\mathbf{v}, \mathbf{w})$ (the MAC ‘metric’ is independent of both phase and amplitude). Modal shape similarity should imply that the MAC is close to unity, and $\text{MAC}(\mathbf{v}, \mathbf{v}) = 1$ for $\mathbf{v} \neq \mathbf{0}$. `eigspec` uses a specific procedure denoted Order-MAC to remove spurious modes in the frequency extraction process. It works as follows.

Apply SSI described in section 2.1.1 to obtain the pair time-series model matrices (A_{r_1}, C_{r_1}) and (A_{r_2}, C_{r_2}) with orders $r_1 < r_2$ (i.e. two state-space models are estimated for the same block data). Extract the modal shape vectors from both models (r_1 and r_2) as outlined above in this section 2.1.2. For each mode in the smaller model with dimension r_1 , find its best match (using the MAC value [9]) shape vector in the larger model with dimension r_2 . If the MAC value of the best match is higher than a threshold value close to 1, say 0.998, then add the mode of the smaller model to a shortlist of modes. If the MAC value is smaller than the threshold, ignore the mode.

The above description mentions a single threshold value to ‘filter out’ spurious modes. In addition, `eigspec` can use a second threshold value to check that the eigenvalue λ_i of the mode is nearly invariant with respect to r_1 and r_2 . Specifically, the value $\max(0, 1 - |\lambda_i^{(r_1)} - \lambda_i^{(r_2)}|/|\lambda_i^{(r_1)}|)$ has to exceed the second threshold for a mode to be shortlisted. Order-MAC is based on observations in OMA [27, 30] that spurious modes

can be strongly dependent on the estimation parameters (here the subspace dimension) but that physical modes are often insensitive to the same parameter.

2.1.3. Random projection and shape-vector estimation. `eigspec` uses the SSI and MAC algorithms outlined above in a particular way to significantly speed up the calculations. It is observed that multiplying the channel data ‘spatially’ by a constant matrix does not alter the frequency content of a block of data. This can be exploited to reduce the channel dimension of the analysis, which is important since the computational complexity of SSI grows rapidly with the channel dimension m and the f , p horizons. In fact, the time required to numerically calculate the SVD of the $mf \times mp$ matrix R_{fp} in the basic SSI procedure (4b) above is proportional to $m^3 p^2 (f + p)$ using a typical algorithm [12]. The cubic dependence on the array dimension m implies that the SVD calculation would be 8 times quicker if the array sample was ‘compressed’ to half its original size, say $l = m/2$. It turns out to be quite effective to use a random projection matrix which is constructed independently of the actual data in the block Y . The idea of using a random projection comes from recent developments in compressive sampling [31]. The Mirnov array in DIII-D has about $m \approx 40$ channels (depending on how many of them are working properly and whether all of them are used in the analysis) but it has been found that using about $l \approx 10$ channels (from random projection) essentially gives the same results (but is nominally about 64 times faster).

Specifically, if $\mathbf{y}(k) \in \mathbb{R}^{m \times 1}$ is a magnetic probe array sample at time k , and

$$Y = [\mathbf{y}(1) \quad \dots \quad \mathbf{y}(N)] \quad (10)$$

is the block data then introduce a random projection matrix $\Psi \in \mathbb{R}^{l \times m}$ with $l \leq m$ and form the smaller block data

$$\Psi Y = Z = [z(1) \quad \dots \quad z(N)]. \quad (11)$$

The random projection matrix Ψ is constructed such that all its rows are orthogonal. Specifically, $\Psi \Psi^T = I_l$ where I_l is the l -dimensional identity matrix. In the analysis code `eigspec`, the lower dimensional signal $z(k)$ is analysed instead of the actual full-array sample $\mathbf{y}(k)$. This however obfuscates the shape vectors. To solve this problem, SSI with Order-MAC is used only to obtain a shortlist of frequencies $\Omega = \{\omega_j\}_{j=1}^{d \lesssim r_1/2}$ for each block of projected data Z . The frequency shortlist Ω is then used to approximately factorize the full data Y as a low-rank matrix as follows:

$$Y = \sum_{i=1}^d \mathbf{d}_i \mathbf{s}_i^T + E = D S^T + E \quad (12)$$

where \mathbf{d}_i is the i th column of D , \mathbf{s}_i is the i th column of S , and E is a residual matrix. The columns of S are prescribed from Ω : each ω_j contributes two columns to S : $s_j(k) = \cos(\omega_j k)$, $s_{j+1}(k) = \sin(\omega_j k)$. Finally, the columns of D are found by least-squares estimation

$$\hat{D} = (Y S)(S^T S)^{-1} \quad (13)$$

and the columns $\widehat{\mathbf{d}}_i, \widehat{\mathbf{d}}_{i+1}$ are used as the shape vector for the corresponding frequency ω_j which was associated with the creation of the columns \mathbf{s}_i and \mathbf{s}_{i+1} . The use of expression (13) minimizes the Frobenius norm of the residual matrix $E = Y - \widehat{D}S^T$ in equation (12).

The final stage of a typical modal extraction is to calculate a root-mean-square (RMS) amplitude for the mode with frequency ω_j . This is done as

$$\text{RMS}_j = \sqrt{(\widehat{\mathbf{d}}_i^T \widehat{\mathbf{d}}_i + \widehat{\mathbf{d}}_{i+1}^T \widehat{\mathbf{d}}_{i+1})/m}. \quad (14)$$

When the raw Mirnov data is being analysed (time derivative of the poloidal magnetic field component) the amplitude (14) need to be divided by ω_j in order to obtain the overall RMS mode amplitude for mode j .

2.2. Post-processing techniques

The output from the feature extraction phase is a small set of modes for each block Y of time-series data. A full-array analysis would repeat this block analysis over the entire tokamak shot dataset by repeatedly shifting the position of the analysis window by a user-determined increment and aggregating the results. The basic visualization of the results would be to scatter the extracted frequencies versus the block-centre time and similarly for RMS amplitudes in other plots. With each extracted frequency ω_j there is also an associated shape vector defined as

$$\mathbf{v}_j = \widehat{\mathbf{d}}_i + i\widehat{\mathbf{d}}_{i+1} \quad (15)$$

using the notation introduced for equation (13). This opens up several other options for visualization, exploration and classification. Section 2.2.1 outlines a clustering technique that can group the entire set of modes found in the shot into classes based on the shape vectors (15) only. The grouping algorithm does not need to know the actual coordinates of the magnetics array probes, but its performance depends implicitly on them (since they determine how separable the modes are). Section 2.2.2 explains a non-parametric regression technique that exploits the Mirnov array probe coordinates in a particular way so that a periodic modal shape can be extrapolated without any reference to poloidal or toroidal period numbers.

2.2.1. Spectral clustering. The complex-valued shape vectors (15) can be used in combination with the MAC similarity measure to construct a sparse k -nearest-neighbour (k -NN) similarity (graph) matrix which can be processed using spectral clustering algorithms [32, 33]. The point in doing so is to automate the classification of the various modes found in a shot to efficiently visualize and recognize structure and pattern. Let $\mathcal{V}_B = \{\mathbf{v}_b\}_{b=1}^B$ be a set of B shape vectors extracted from a (time slice of a) shot using the SSI/Order-MAC method explained above. The k -NN MAC-based similarity matrix S is then constructed via the following two steps.

- 1 For each shape vector \mathbf{v}_b , $b = 1 \dots B$, compute $\text{MAC}(\mathbf{v}_b, \mathbf{v}_i)$, for all i excluding b and store its k most similar shape-vector indices (those with the k highest MAC scores).

- 2 Assemble a sparse S so that $s_{b_1, b_2} = \text{MAC}(\mathbf{v}_{b_1}, \mathbf{v}_{b_2})$ if \mathbf{v}_{b_1} is in the k -NN list for \mathbf{v}_{b_2} or \mathbf{v}_{b_2} is in the k -NN list for \mathbf{v}_{b_1} . Otherwise $s_{b_1, b_2} = 0$.

The similarity matrix can then be processed using a sparse-matrix version of the procedures detailed in [32]. This involves a first step when the matrix S is converted to a graph Laplacian, a second stage where a sparse eigenvalue problem is solved and a final stage where the eigenvector coordinates found in stage two are clustered using a standard K -means algorithm [34]. But this is feasible even for B (the number of shape vectors) up to several tens of thousands on typical workstations. A crucial clustering parameter is the number of clusters to be formed.

2.2.2. Gaussian process regression. Shape vector (15) is regarded as a spatially discretized sample of a continuous periodic plasma eigenmode shape at a particular time. Gaussian process regression (GPR) [34, 35] turns out to be useful for estimating this continuous mode shape. GPR is a technique where the regularizing properties are removed from a direct parameterization and placed in a kernel instead. A kernel can be thought of as a description of the covariance structure of a random field. The kernel

$$k(\mathbf{x}, \mathbf{x}') = k(\theta, \theta', \phi, \phi') = k_p(\theta, \theta')k_t(\phi, \phi') \quad (16)$$

where $\mathbf{x} = (\theta, \phi)$ and $\mathbf{x}' = (\theta', \phi')$ and

$$k_p(\theta, \theta') = \exp(-2 \sin^2((\theta - \theta')/2)/\sigma_\theta^2) \quad (17a)$$

$$k_t(\phi, \phi') = \exp(-2 \sin^2((\phi - \phi')/2)/\sigma_\phi^2) \quad (17b)$$

can be used to model $2\pi \times 2\pi$ -periodic patterns on (θ, ϕ) with two length scales $\sigma_\theta, \sigma_\phi$ [34].

Let $\mathbf{x}_i = (\theta_i, \phi_i)$ be the i th ‘training point’ with the known target value $f_i \in \mathbb{R}$. Let \mathbf{x} be an arbitrary coordinate where an unknown value of a continuous field is to be evaluated, given the discrete training points (i.e. the real and imaginary components of the shape vector (15)). Given the ‘hyper-parameters’ $\sigma_\theta, \sigma_\phi$ and the noise level σ_1 the kernel (16) is then used as follows.

Construct

$$\mathbf{K}_{11} = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & k(\mathbf{x}_1, \mathbf{x}_2) & \dots & k(\mathbf{x}_1, \mathbf{x}_m) \\ k(\mathbf{x}_2, \mathbf{x}_1) & k(\mathbf{x}_2, \mathbf{x}_2) & \dots & k(\mathbf{x}_2, \mathbf{x}_m) \\ \vdots & \vdots & \ddots & \vdots \\ k(\mathbf{x}_m, \mathbf{x}_1) & k(\mathbf{x}_m, \mathbf{x}_2) & \dots & k(\mathbf{x}_m, \mathbf{x}_m) \end{bmatrix} + \sigma_1^2 \mathbf{I}_m \quad (18)$$

and

$$\mathbf{k}_{12}(\mathbf{x}) = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}) \\ k(\mathbf{x}_2, \mathbf{x}) \\ \vdots \\ k(\mathbf{x}_m, \mathbf{x}) \end{bmatrix} \quad (19)$$

for the m training points (components of the shape vector). Set $k_{22} = k(\mathbf{x}, \mathbf{x})$. The predicted noise-free value at the arbitrary coordinate \mathbf{x} is then

$$\widehat{f}(\mathbf{x}) = \mathbf{k}_{12}(\mathbf{x})^T \mathbf{K}_{11}^{-1} \mathbf{f} \quad (20)$$

where \mathbf{f} is the vector of target values f_i . The variance of the predicted value at the arbitrary coordinate \mathbf{x} is

$$\widehat{\sigma}^2(\mathbf{x}) = k_{22} - \mathbf{k}_{12}(\mathbf{x})^T K_{11}^{-1} \mathbf{k}_{12}(\mathbf{x}). \quad (21)$$

The pattern and shape estimation routines in `eigspec` use the expressions (20) and (21) to draw smooth estimates of the toroidal and poloidal shapes of an eigenmode. In this case the \mathbf{x}_i 's are the magnetic probe locations on the tokamak internal surface and the target values \mathbf{f} are either the real or the imaginary parts of the shape vector. The smooth toroidal and poloidal shapes are then obtained using (20), with $2\widehat{\sigma}$ error bars from (21), by evaluation at \mathbf{x} 's along the toroidal and poloidal arrays respectively.

Importantly, there is a consistent method of selecting the set of smoothing hyper-parameters. The logarithm of the likelihood function (log-likelihood) of σ_θ , σ_ϕ , σ_1 , given the data \mathbf{x}_i , f_i ($i = 1 \dots m$) is

$$l(\sigma_\theta, \sigma_\phi, \sigma_1) = \frac{1}{2} \log \det K_{11} + \frac{1}{2} \mathbf{f}^T K_{11}^{-1} \mathbf{f} + \frac{m}{2} \log(2\pi) \quad (22)$$

and it should be as large as possible. So direct optimization of (22) can be used to regularize the continuous field estimation for the mode shape (20).

After estimating periodic array-aligned poloidal and toroidal shape interpolation on fine grids in the sense above, a 2D pattern is extrapolated as a rank-2 matrix as follows. Let \mathbf{p}_{re} , \mathbf{p}_{im} be the poloidal-array real and imaginary shape estimates. Let \mathbf{t}_{re} , \mathbf{t}_{im} be the toroidal-array real and imaginary shape estimates. The mode pattern is constructed as

$$\begin{aligned} P &= \text{Re} \{ (\mathbf{p}_{\text{re}}^T + \mathbf{p}_{\text{im}}^T) \times (\mathbf{t}_{\text{re}} + \mathbf{t}_{\text{im}}) \} \\ &= \mathbf{p}_{\text{re}}^T \mathbf{t}_{\text{re}} - \mathbf{p}_{\text{im}}^T \mathbf{t}_{\text{im}} \end{aligned} \quad (23)$$

and the matrix P is drawn as a contour plot. Pattern estimation (23) is referred to as an extrapolation for the following reason. GPR could be used to evaluate the field over the poloidal-toroidal plane directly. However, since the density of Mirnov db_θ/dt -array probes outside the toroidal and poloidal sub-arrays is basically zero at DIII-D (figure 5), this will not work so well.

2.3. Additional comments

The modal analysis performed by `eigspec` provides a complex-valued eigenvalue λ_i for an array mode (8a) and (8b). In principle there is information on both growth/damping rates and the frequency of a particular mode. It turns out that in typical DIII-D Mirnov array analysis, this extra information (compared to e.g. FFT-based frequency analysis) does not seem to be so easy to exploit. It is nearly always the case that a shortlisted mode has an eigenvalue magnitude $|\lambda_i|$ very close to unity (this means it is nearly a stationary-amplitude oscillation). Fine details on amplitude dynamics appear to be easier to extract by shifting the block-analysis window by small increments in time instead.

It has also been observed from investigation using basic robust statistics techniques that short outlier-like bursts of Mirnov activity from edge localized modes (ELMs) do not

seem to have a significant effect on the performance of the `eigspec` code. A sub-division of a Mirnov data block into shorter segments and a segment exclusion method based on the minimum covariance determinant (MCD) criterion [36], did not yield a significant change of the modal analysis outputs. A tentative explanation for this is that the ELM signature is comparatively broadband (both in time- and space frequencies). This makes its effect on a line-spectrum code relatively small.

2.4. Numerical examples

2.4.1. Example comparison with cross-spectrum approach.

Approaches based on short-time discrete Fourier transforms (DFTs, FFTs) are limited in frequency resolution by the length of the signal block that is analysed. But the Mirnov array data is rarely a stationary signal, so it becomes important to have the capability to resolve frequencies from quite short segments of data in order to follow transient and time-varying MHD mode activity. An example illustrates the capability of the present method to go below the typical non-dimensional DFT resolution $\Delta\omega = 2\pi/N$, where N is the number of samples. Two rotating modes are recorded with a signal-to-noise ratio (SNR) of 20. The recording time is $N = 200$ samples which gives a DFT $\Delta\omega = 0.0314$ rad/sample. The mode frequencies are $\omega_1 = 0.41$ and $\omega_2 = \omega_1 + \Delta\omega/2 \approx 0.4257$ (the mode numbers are $n_1 = 2$ and $n_2 = 1$ in this case, and the phases are randomized). Two Mirnov probes, selected from a synthetic toroidal array and displaced by $\Delta\phi = 33$ degrees, are used. Figure 3(a) shows the time traces of these two channels y_1, y_2 . As anticipated due to the DFT resolution, the two frequencies are merged into one peak in the cross-power spectrogram of the two probes. This is seen in figure 3(b). Figure 3(c) shows the cross-coherence function for this case. Also displayed in figure 3(b) are two vertical lines. These two lines are the estimated frequencies obtained by the subspace-based method for this segment of data. A zoomed-in version of the peak cross-power spectrum and the two estimated lines are shown in figure 3(d). This example shows that in situations where the FFT-based cross-spectrum approach would have difficulties, the new method can still perform well. Note that toroidal mode-number estimates are also obtained by utilizing knowledge of the geometrical displacement of the probes. The DIII-D tokamak normally records Mirnov channels at a sample rate $f_s = 200$ kHz. This means that $N = 200$ samples equals 1 ms block length. And the resolved frequency difference $\Delta\omega/2$ corresponds to 500 Hz. In this particular numerical example the estimated line spectrum frequencies were $\widehat{\omega}_1 = 0.4094$, $\widehat{\omega}_2 = 0.4283$ rad/sample (which are close to the true source frequencies $\omega_1 = 0.41$, $\omega_2 \approx 0.4257$ as defined above).

2.4.2. Example comparison with the SVD approach.

Another popular approach to general magnetic probe array analysis is the SVD of the block data matrix (sometimes known as biorthogonal decomposition). The SVD may however not be so well suited for general modal analysis since it imposes constraints on the orthogonality of frequencies with respect




Figure 3. Example showing that the new method can resolve frequencies that are more closely spaced than what is achievable by the standard cross-spectrum-based method. Both methods are applied to two channels (a). The cross-spectrum approach would look for peaks in the cross-power (b) and check that the cross-coherence (c) is above some confidence threshold. However only one peak is seen in (b). The subspace-based method gives a line spectrum with two closely spaced frequencies shown in red; (b) and zoom-in (d).

to the time frame (as does DFT too) and an orthogonality constraint in a spatial sense. This means that the SVD can fail both in near frequency-degenerate cases as well as in spatially degenerate cases. The reason for this is that the SVD is a low-rank matrix approximation technique in the sense of the Frobenius norm (least squares) [12]. This implies that the chief criteria for SVD is to capture as much signal variation as possible with few terms. In general, this is not the same as extracting sinusoidal components from a data matrix. Figure 4 illustrates a case where the SVD fails to accomplish modal analysis (but seems to achieve data compression) whereas the new method, denoted SSI, succeeds in extracting the modal patterns. In this example the block data length is again $N = 200$ samples. The normalized frequency of the first mode is $\omega_1 = 0.41$ rad/sample. The second mode frequency is set to $\omega_2 = \omega_1 + \Delta\omega/8$ (corresponding to a real frequency separation of 125 Hz and a real signal length of 1 ms for the typical sample rate of $f_s = 200$ kHz). The block data matrix is composed of multi-channel time-series data obtained from a toroidal array with 20 angularly equidistant synthetic probes. The SNR of the data is 20. The toroidal mode numbers are as in the first example in section 2.4.1. Note that the order of the extracted modes is arbitrary. In figure 4, it is clearly

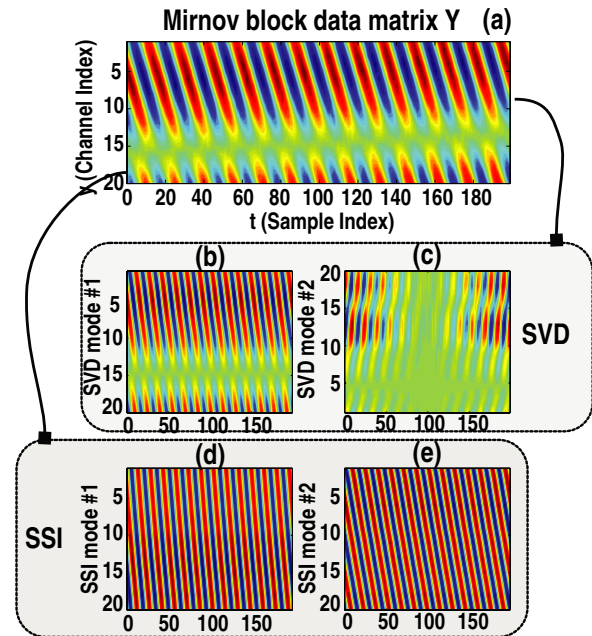


Figure 4. Example illustrating that the new method is able to resolve modes more reliably than a direct SVD of the Mirnov data matrix. The Mirnov data beating-effect pattern (a) is indeed the superposition of the two recovered SSI modes (d) and (e). It is seen that the first SVD mode (b) looks a lot like the full data Y in subfigure (a) already, but this is not the primary objective of a modal decomposition.

seen that the SVD ‘modes’ (b)–(c) fail to resemble the planar-wave-looking patterns that was used to construct the Mirnov data matrix Y (a). The SSI modes (d) and (e) are very close to the ‘true’ modes. As in the previous example, the method was not informed that there actually were two modes to be recovered. This is another problem encountered in the use of direct SVD: how many singular vectors to choose and which to pair?

3. DIII-D applications

Figure 5 summarizes the installation of the poloidal field magnetic probes on the DIII-D tokamak that are used in this work [11, 37]. The angular probe positions (θ, ϕ) are illustrated in figure 1. The blue-marker sensors constitute the toroidal outer midplane array. The green-marker sensors record poloidal spatial variation. The blue dotted and the green dotted lines respectively show the directions along which the GPR regularization technique of section 2.2.2 evaluate the eigenmode spatial structure. These lines are along the high-probe density directions. All sensors in figure 5 are typically used for optimizing the hyper-parameters according to equation (22).

All the modal pattern plots to be shown in this section are normalized in amplitude. Specifically, the GPR smooth shape regularization procedure is applied on the unit vector $\hat{v}_j = v_j / \sqrt{v_j^\dagger v_j}$ where v_j is defined by equation (15). There is no compensation done for the shape estimates with respect to magnetic probe distance to the plasma surface

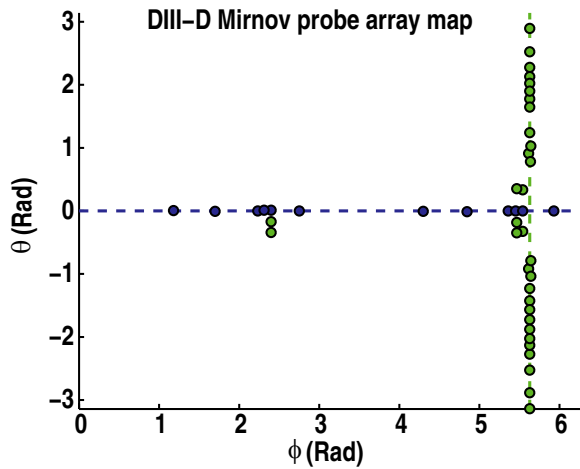


Figure 5. Locations of the available Mirnov probes that directly digitizes the poloidal-component magnetic-field time derivative db_{θ}/dt . The coordinates are the inner vessel surface location expressed in the 2π -periodic (θ, ϕ) -space.

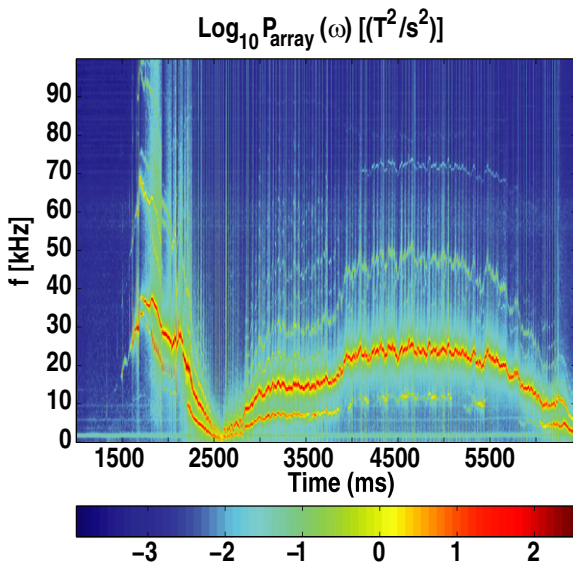


Figure 6. Averaged spectrogram of the Mirnov array for the NTM control experiment discharge #150792. 34 probes were available for analysis of this shot.

or discontinuities of the probe field direction, since that is not essential for the present purposes. This means that the modal patterns shown below are not exactly, but approximately, representative of the plasma surface pattern. Whenever forward modelling with synthetic diagnostics at the actual magnetic probe locations is used for comparing the empirical results with theory (usually in the sense of linear eigenfunctions, but more complete nonlinear simulations are conceivable too) then the inference problem from probe location to plasma surface is not needed.

3.1. Automatic clustering of MHD activity in an NTM control experiment

Figures 6 and 7 overview the results obtained by running the eigspec analysis on the DIII-D NTM control shot #150792.

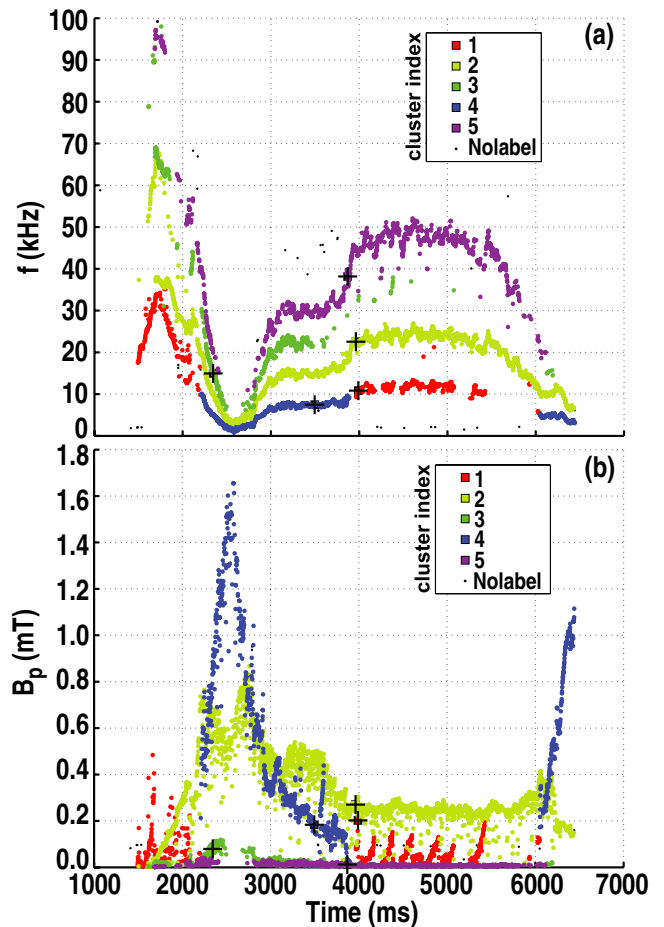


Figure 7. DIII-D shot #150792 analysed with the eigspec code. A 2/1 NTM (cluster 4, in blue) grows rapidly and slows down. Before it reaches the locked state the ECCD-actuated control system manages to shrink the 2/1 amplitude so it speeds up and disappears. The 2/1 NTM reappears and locks soon after ECCD is turned off at $t = 5600$ ms. (a) shows the extracted modes plotted in a frequency plot. The modes have been clustered into 5 different groups based on pairwise shape-vector similarities. (b) shows the RMS amplitudes of the shape vectors clustered exactly as in (a). The cross-hairs mark the cluster medoids (i.e. the most representative, or ‘mean’, shape vector for each cluster).

Figure 6 shows an FFT spectrogram that has been averaged over all available magnetic probes for this shot. Figure 7 shows the eigspec feature extraction results and can be understood to expand the spectrogram picture of figure 6.

The purpose of this experiment is to generate a $m/n = 2/1$ NTM that triggers a ‘catch & subdue’ control algorithm [38]. The control algorithm monitors the $n = 1$ RMS amplitude component (which is derived from the toroidal array only). When this amplitude exceeds a preset value, the gyrotrons are powered-up for electron cyclotron current drive (ECCD) on its best estimate for where the $q = 2$ surface is in the plasma. This type of control requires a trigger on a small NTM signature for it to efficiently remove the NTM. In theory, if the NTM is detected before its amplitude is above a certain critical marginal level, it can be shrunk relatively fast [39, 40]. If the catch is at an amplitude higher than this marginal level, the catch is significantly more power- and time-consuming. It has been observed that prompt distinction of the 2/1 NTM is challenging

due to the presence of sawtooth (ST)- and fishbone-related internal kink signatures with the same toroidal period number $n = 1$. This additional $n = 1$ signature has an amplitude which is of the order of the desired trigger level for the NTM signature and easily produces false positives. A further complication is that the frequency is also comparable with the expected range of the anticipated NTM frequency.

In this shot, ECCD is triggered by the $n = 1$ RMS level at $t = 2230$ ms. As seen in figure 7 this manages to make the 2/1 mode disappear before $t = 4000$ ms (the 2/1 mode is labelled by cluster index 4, explained below). The ECCD is turned off at $t = 5600$ ms and the 2/1 mode reappears around $t = 6000$ ms.

The estimation parameters for the SSI Order-MAC code for this example is as follows. The block length $N = 400$ corresponds to a real-time block duration of 2 ms since the sample frequency is 200 kHz. The block increment is equal to the block length, so there is no overlap (but all data is used). The random projection dimension is $l = 12$ and the future and past horizons are respectively $f = 10$ and $p = 20$. The Order-MAC orders are $r_1 = 12$ and $r_2 = 20$. This means that the maximum number of modes that can be extracted from each block is 6 (at two states per mode). The spurious mode threshold values are both set to 0.998. The post-processing of the extracted modes was done using a 60-NN MAC similarity matrix. The mode extraction took about 30 s and the sparse-method clustering took around 25 s (on a typical laptop computer). For comparison, about 4 s is required to calculate an FFT-based mean spectrogram as shown in figure 6 (on the same hardware, using the same block size and block stride, and with zero-padded FFT of length 2048).

The clustering of the extracted modes from shot #150792 are shown in figures 7(a) and (b). It turns out that eigspec can achieve an automatic separation of the 2/1 NTM (blue symbols, cluster 4) from the ST-related internal kink (red symbols, cluster 1). The dominant classes of modes are represented as patterns in (θ, ϕ) -space in figure 8. The phases of the patterns are arbitrary. Note that figure 8(b), cluster 2 in figure 7, is a 3/2 pattern that corresponds to the illustration in figure 1 (if it was the only mode in the present shot). The modes labelled as 1 and 4 in figure 7 respectively correspond to the extrapolated patterns of figures 8(a) and (b). It can be seen that these two types of modal patterns show high similarity if only the outboard midplane is considered. This is the case when only the toroidal array is employed. If more poloidal shape information is used then the modal patterns appear to be separable. Specifically, the difference between these modal shapes becomes large near the inboard midplane. The internal kink mode in figure 8(c) even seems to change helicity along the poloidal angle θ . This has previously been dubbed ‘phase folding’ [3]. It will be further commented in section 3.3 below.

The application of the eigspec analysis toolkit on this NTM control experiment suggests that it is possible to develop a classifier system (based on the full-array shape vector (15) extracted by the present algorithm) that can reliably distinguish between the $m/n = 2/1$ mode and the $n = 1$ ST-related internal kink mode. It does not seem enough to restrict the analysis to the toroidal array (both $n = 1$, and both can have




Figure 8. Shot #150792 shape-vector pattern plots using the GPR technique. Cluster 4 has a 2/1 modal pattern (a). Cluster 2 has a 3/2 modal pattern (b). Cluster 1 has a ‘3/1’ modal pattern (c). It is typically interpreted as 1/1. It can be seen that the outboard-midplane ($\theta = 0$) modal patterns for the internal kink (c) and the 2/1 NTM island (a) are quite similar $m \sim 2, 3$, but the internal kink exhibits ‘phase folding’ on the inboard side as illustrated by the dashed lines in (c). The distance from a probe to the plasma surface is largest around the up/down locations at $\theta = \pm\pi/2$ (this explains why the contours are fainter around these coordinates in the above patterns).

comparable frequencies). Detailed poloidal shape information seems to be mandatory to clearly separate these two classes of modal pattern.

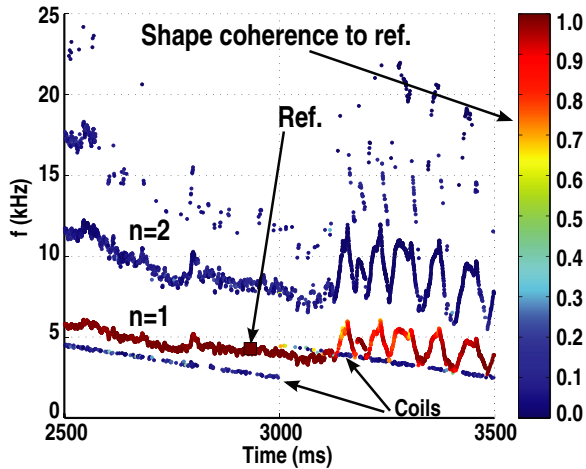


Figure 9. Mirnov analysis overview for DIII-D experiment #153343, timeframe 2500 to 3500 ms. All extracted modes are coloured by their similarity index (the MAC value) with respect to a reference shape vector located at $t \approx 2932$ ms, $f \approx 4.28$ kHz. The time trace with a significant dark red colour has a $n = 1$ toroidal geometry. The dark blue trace at around twice the $n = 1$ frequency is an $n = 2$ trace. The inclined straight lines are internal coil driven excitations.

3.2. Characterization of an $n = 1$ edge harmonic oscillation.

The quiescent H-mode (QH-mode) is characterized by H-mode-like edge plasma conditions but without the typical ELM activity [41, 42]. ELMs are believed to be a serious danger to the plasma wall in large machines such as ITER since they result in intense localized heat loads. The QH-mode may offer a regime that avoids the ELMs but otherwise retains the H-mode confinement advantage. In place of the ELMs, edge transport is facilitated by MHD phenomena known as edge harmonic oscillations (EHOs). The EHO seems to exist in edge plasma conditions close to the onset of ELMs. EHOs have been observed to increase edge particle transport without degrading the edge thermal transport barrier. It is thought of as a saturated kink-peeling mode driven by edge current density and edge rotational shear. EHOs have been observed both as narrow-band oscillations and broadband activity. The experimental characterization of the EHO is usually limited to assigning its toroidal period number n and its oscillation frequency f .

The present analysis toolkit can provide an experimental (external magnetics) characterization of the narrow-band EHO phenomenon. This is exemplified here for DIII-D shot #153343. This experiment attempted to demonstrate interaction with the EHO using internal saddle loop coils (excited at several kHz to match the anticipated EHO frequencies). A shape-similarity map of the extracted modes for the timeframe 2500 to 3500 ms is shown in figure 9. The plot is produced by assigning a colour index to each extracted mode based on its shape vector w and a reference shape vector v . The reference shape vector v is located at $t \approx 2932$ ms and its frequency is $f \approx 4.28$ kHz. The reference shape vector is marked by a red square marker in the figure. The colour index is then simply the MAC value (9) between w and v . The colour index is given the colour according to the colour bar on

the right-hand side in figure 9. The estimation parameters in this case were similar to the NTM clustering example above. The block length was $N = 500$ (2.5 ms at $f_s = 200$ kHz), the block-step size was 50 samples or 0.25 ms. The past and future horizons were $p = 20$ and $f = 10$. The random projection dimension was $l = 12$. The Order-MAC orders were $r_1 = 12$, $r_2 = 20$. The Order-MAC threshold values were set to 0.995.

It can be seen that the similarity measure (9) seems to take a high value for a contiguous trace. This is indeed the evolution of a clear $n = 1$ component in the Mirnov data. The straight lines with some inclination that can be seen in the figure are actually the ‘modes’ extracted for the internal coil coupling to the Mirnov array (frequency sweeps starting at 4.5 kHz at $t = 2500$ ms and coming down to 2.5 kHz at $t = 3000$ ms, and similarly in the timeframe $t = 3000$ ms to $t = 3500$ ms). It turns out here that as this externally applied frequency aligns with the EHO $n = 1$ trace frequency, a strong modulation results, after $t \approx 3100$ ms. In fact both the amplitude (not shown here) and the frequency of the EHO trace are strongly affected. This will be analysed in more detail elsewhere. The point here is that the eigspec toolkit can effectively give an overview of the Mirnov data. Specifically, a modulation effect of the actual EHO shape appears to also exist (i.e. the colour variation of the $n = 1$ time trace in the modulated part of figure 9).

Figure 10 shows the spatial extrapolation results using the GPR technique for the specific reference shape vector used in figure 9. Figures 10(a) and (b) should be understood in relation to figure 5. The dashed lines in figure 5 are the one-dimensional directions along which figures 10(a) and (b) are evaluated. These evaluations (using the GPR technique) give $2\hat{\sigma}$ ‘error bars’ which are indicated in the plots as dashed lines (the area within the dashed curves should really be thought of as a $\sim 95\%$ confidence interval for the Gaussian process, given the hyper-parameters, described in section 2.2.2 and references therein). The markers in the poloidal and toroidal shape plots are the actual shape-vector values. These markers should be thought of as noisy observations. The smooth periodic curves are interpreted as the underlying functions which could have generated the noisy observations. The confidence interval is not for the possible observations but for the underlying smooth periodic function. This implies that the markers can be outside the region between the dashed lines. Figure 10(c) is the combination of the real- and imaginary-part smooth curve pairs in the sense of equation (23).

An MHD stability code that provides EHO eigenmodes could be compared to fully empirical modal plots such as those in figure 10. Perhaps the easiest way to achieve this is to transform the theoretical mode to the actual Mirnov array coordinates as a forward problem (including sensitivity studies).

3.3. ST pre- and postcursor (successor) magnetics analysis

The ST cycle was first observed almost 40 years ago [43]. It starts by a peaking current density due to decreased resistivity with temperature in the plasma core. The peaking is concomitant with a ‘precursor’ $n = 1$ magnetics oscillation.

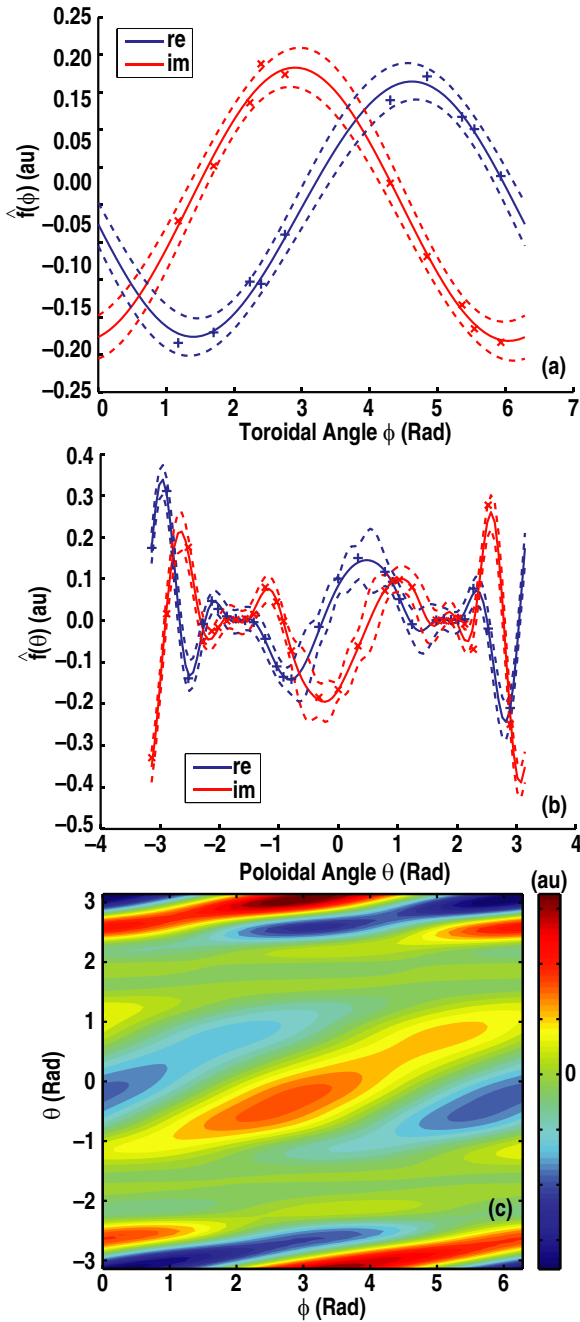


Figure 10. Example EHO spatial shape characterization using the GPR techniques. EHO shot is #153343. The extracted shape vector is the same as the reference vector in figure 9. (a) and (b) illustrate the array-aligned shape estimates obtained with the GPR technique of section 2.2.2. (c) depicts the full modal pattern estimate based on (a) and (b). Reversed I_p compared to the normal DIII-D operation.

The current peaking tends to push the central safety factor q_0 close to 1 and triggers an internal kink or quasi-interchange mode with $m = n = 1$. An abrupt re-arrangement of the plasma core flattens the temperature and current profiles. After this ‘crash’ event, the central q_0 tends to be above 1 and the cycle restarts. The details of the nonlinear ST cycle are complicated and is a field of active research. The internal kink and the ST cycle significantly depend on the shaping parameters of the tokamak plasma, such as ellipticity

and triangularity [44, 45]. Sometimes there is a ‘postcursor’ magnetic oscillation that decays after the crash (these are also known as ‘successor’ oscillations). Sometimes the crash triggers other MHD modes such as NTMs [1].

STs are not believed to pose a direct problem for burning plasma fusion performance. Instead a main concern is that fast ions in hot plasmas appear to stabilize the internal kink mode. This has an effect of producing a slow ST cycle characterized by large ST crashes and a cycle time which is longer than an energy confinement. These ‘monster’ STs can be more prone to trigger NTMs than smaller STs. One recent approach to destabilize STs (to make the ST more frequent, smaller and less prone to trigger NTMs) is to use ECCD [46, 47].

For the present purposes, in connection with the NTM discrimination problem described in section 3.1, it becomes relevant to study the variability of ST precursor and postcursor modal patterns. Another motivation for studying the evolution of ST precursor Mirnov array signal patterns is to provide detailed empirical observations that can be related to ST cycle modelling. In what follows, an ST cycle with both pre- and postcursor Mirnov oscillations is analysed with eigspec. The plasma is an inboard-limited L-mode oval-shaped low- q discharge, #154887, designed to study external kink stability. At $t = 1980$ ms the safety factor key values are $q_{95} \approx 2.1$, $q_{\min} \approx 1.0$. Figure 11 shows an ST cycle as seen on the magnetics from this shot. A rapid growth in amplitude of the $|n| = 1$ component, associated with a nearly constant frequency, is observed for ~ 10 ms. See figures 11(b) and (c) for the evolution of these quantities. Times t_1 and t_2 belong to this precursor growth phase. The ST crash event happens at $t \approx 1999.5$ ms. The postcursor oscillation is observed to decay relatively slowly in amplitude and has a frequency which is higher than the precursor frequency. Times t_3 through t_6 belong to this postcursor decay phase.

An $|n| = 2$ signal component is seen to grow and decay together with the $|n| = 1$ component which carries the main array signal amplitude. This signal component can be interpreted as a spatial harmonic of the ST pre- and postcursor since it has twice the frequency of the $|n| = 1$ signal component (and thus the same toroidal phase velocity). The ST crash also triggers an $|n| = 3$ signal component. This is interpreted as a $4/3$ TM that rapidly decays and disappears. The array analysis code can resolve the evolution of the ST modal patterns present in the overview figure 11. The modal patterns for the $|n| = 1$ array signal component at times t_1 – t_6 are depicted in figure 12. The precursor phase (first row; t_1, t_2) exhibits a strong inboard side pattern that seems to have reversed helicity. This is qualitatively similar to the phase-folded pattern for the ST shape representative in section 3.1. But the present plasma is quite different. The second row (t_3, t_4) exhibits a more kink-like pattern. It appears that the ST crash rapidly changes the inboard-side magnetics array pattern. The third row (t_5, t_6) shows a similar pattern to the second row. But it seems that the outboard side bulge in the modal pattern decays faster than does the inboard side. Plasma #154887 is closer to the inboard side wall than the typical diverted DIII-D plasma. The modal patterns for the $|n| = 2$ array signal component at times t_1 – t_4 are depicted in figure 13. The precursor phase (first row;

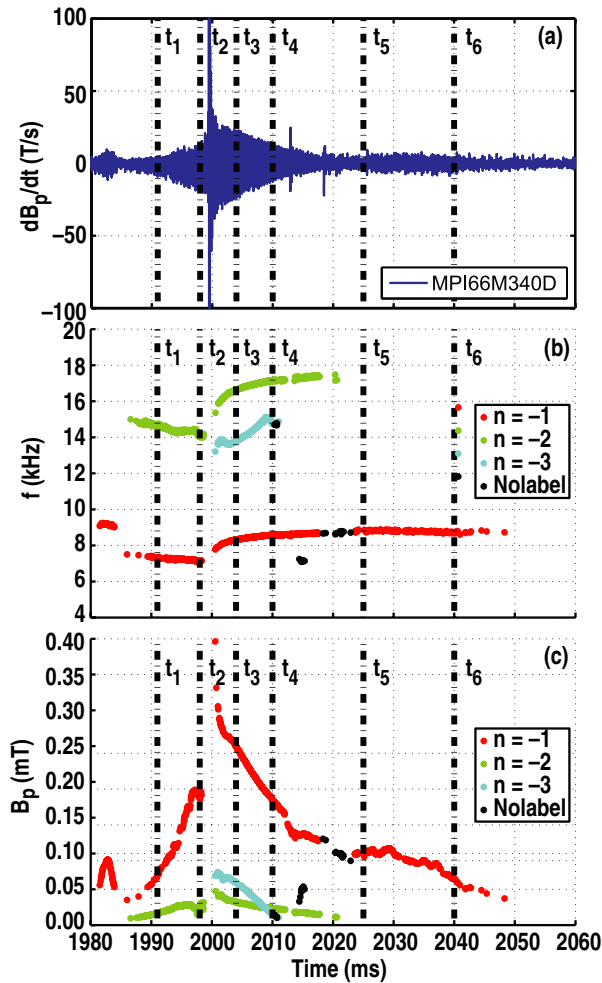


Figure 11. Analysis of a ST cycle as seen on the magnetics array with `eigspec`. The ST crash is at the spike in the time series of figure (a) just after time t_2 . The precursor phase duration is ~ 10 ms, whereas the postcursor oscillation decays over nearly 50 ms. Subplots (b), frequency, and (c), amplitude, are coloured by n -number. The probe in subplot (a) is located at the outboard midplane.

t_1, t_2) exhibits an inboard side pattern with reversed helicity, analogous to the $|n| = 1$ pattern. The second row (t_3, t_4) exhibits a more kink-like pattern, also in analogy with the $|n| = 1$ case. In contrast to the $|n| = 1$ case, the spatial harmonic does not seem to have a clear decay difference for the inboard and outboard sides. The $|n| = 2$ array signal is either not present or too weak to be detected by the method with the current estimation parameters at times t_5, t_6 . Not shown here is the modal pattern for the triggered $|n| = 3$ signal component. It has a dominant $m/n = 4/3$ geometry without any ‘phase folding’.

For all the shots analysed so far with the new code `eigspec`, the precursor phase of an ST cycle is invariably associated with a phase-folded modal pattern. In particular, this is the case for both the H-mode monster STs and ECCD destabilized STs in [47]. A typical phase-folded precursor modal pattern from ECCD-destabilized discharge #145692 is depicted in figure 14. A similar modal pattern is observed for the precursor oscillation in the non-ECCD monster-ST




Figure 12. Evolution of the $n = 1$ array signal component of the ST cycle in shot #154887. The times t_1 – t_6 are defined in figure 11. A phase-folded pattern is observed for the precursor phase t_1 – t_2 . This disappears for the postcursor phase t_3 – t_6 , as indicated by the dashed lines in t_1, t_3 .

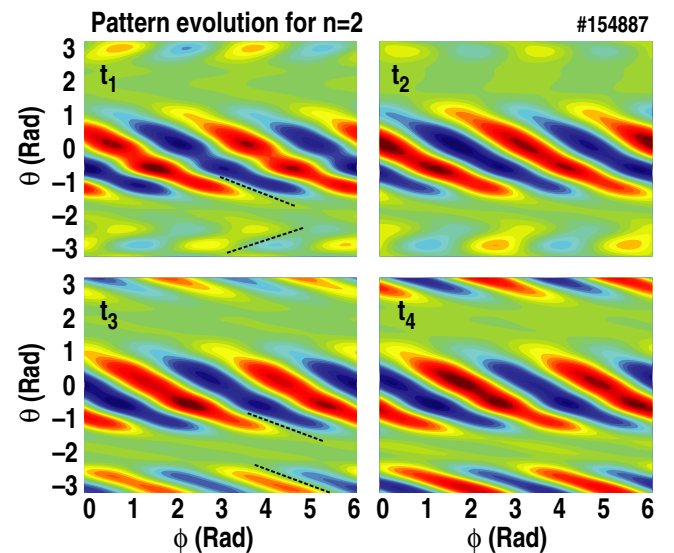


Figure 13. Evolution of the $n = 2$ array signal component of the ST cycle in shot #154887. The times t_1 – t_4 are defined in figure 11. A phase-folded pattern is observed for the precursor phase t_1 – t_2 . This disappears for the postcursor phase t_3 – t_4 , as indicated by the dashed lines in t_1, t_3 .

discharge #145861 from the same study [47]. The details of the modal patterns appears to differ slightly depending on the plasma. This is not surprising since axisymmetric shaping parameters and distances to the Mirnov probes from the plasma

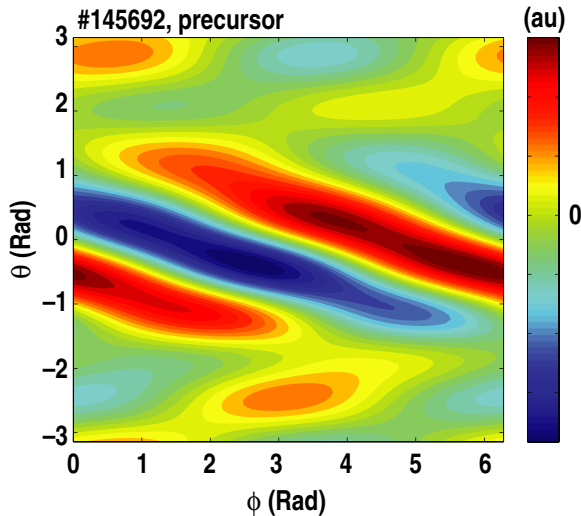


Figure 14. A typical modal pattern for the ST precursor oscillation in discharge #145692, which uses ECCD inside the $q = 1$ surface to destabilize the ST cycle. There is no postcursor oscillation in the ST cycle for this plasma.

surface are expected to influence these details. But the overall characteristics of the precursor modal patterns seem to persist. The simplest interpretation of the phase-folded modal pattern is possibly to think of it as a consequence of the internal ideal kink mode that effectively displaces (tilts) the plasma axis in a $m/n = 1/1$ fashion. The crudest model of this is a single $1/1$ current filament. Such a filament could create a phase-folded pattern. The same inclination of the filament is then seen from opposite sides when viewed from the outboard and inboard respectively.

The next section 3.4 below shows another (related) class of MHD modes that seems to exhibit phase folding on the Mirnov array.

3.4. Fishbone burst magnetics array decomposition

This final DIII-D analysis example illustrates that `eigspec` can be used to decompose dynamic and transient events. Figure 15 shows a single fishbone event [48] from DIII-D shot #155279. The shot is an ITER-like plasma with $q_{95} \approx 5.1$, $q_0 \approx 1.0$ and $\beta_N \approx 2.9$ (at $t = 3205$ ms), designed to explore steady-state scenarios (focus on current-profile optimization). The fishbone instability is believed to be excited by energetic particles when q_{\min} is in the vicinity of 1 [49]. It is considered to be a finite-frequency branch of the ideal internal kink mode [50–52]. It can significantly deteriorate the efficiency of beam heating.

The estimation parameters for the SSI/Order-MAC algorithms in this case were as follows. The block length was $N = 100$ (0.5 ms at $f_s = 200$ kHz), the block-step size was 2 samples (i.e. $10 \mu\text{s}$). The past and future horizons were $p = 25$ and $f = 10$. The random projection dimension was $l = 16$. The Order-MAC orders were $r_1 = 12$, $r_2 = 20$. The Order-MAC threshold values were set to 0.992. Notice that 0.5 ms is about a tenth of the length of the entire fishbone burst as seen in figure 15(a). This means that the `eigspec`

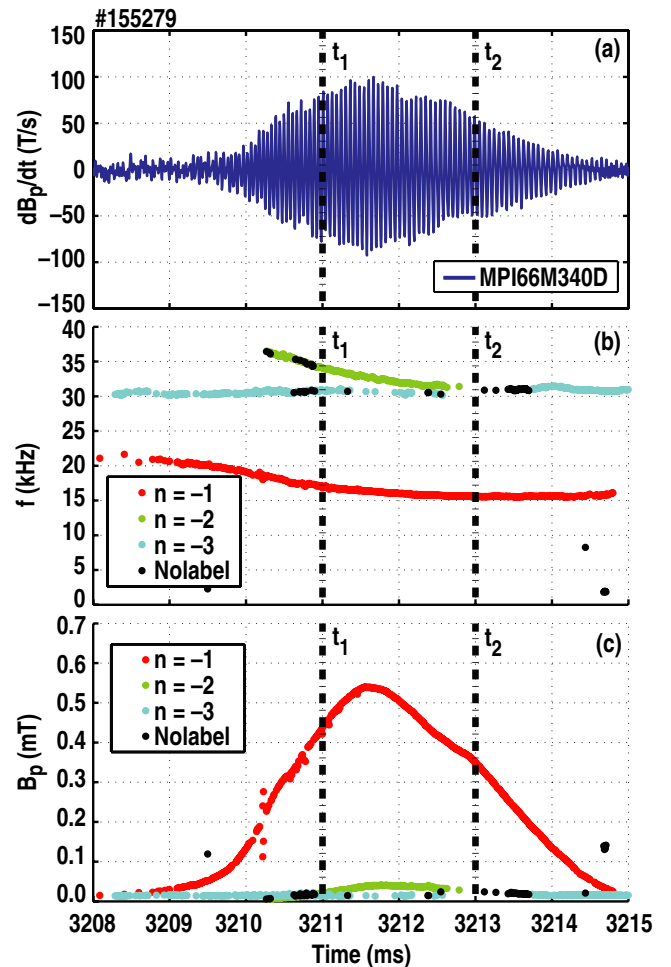


Figure 15. Analysis of a fishbone burst as seen on the magnetics array with `eigspec`. The decomposition of the array signal illustrated by (b) and (c) exemplifies that the code can be used to analyse quite dynamic and short-time transient behaviour. A single outboard-midplane probe time series is shown in subfigure (a). The same modal pattern can be used for both times t_1 and t_2 .

analysis is indeed achieving localized time-series modelling for the fishbone event.

Figures 15(b) and (c) indicate that the fishbone in the present plasma is mainly described by an amplitude envelope and a frequency sweep. The frequency starts at about 20 kHz and sweeps down to about 15 kHz. The amplitude rapidly goes up and then down. The amplitude decay is somewhat slower than the growth. The times marked by t_1 and t_2 in figure 15 are used to explore the modal shape pattern. These times respectively exemplify different fishbone phases with amplitude growth and decay. It turns out that the fishbone event analysed here is well described by a single modal pattern. Thus the amplitude envelope can be simply interpreted as a scaling of the modal pattern. The modal pattern for time t_2 is shown in figure 16. It would look the same if it was plotted for t_1 (modulo a toroidal phase shift). It is observed that the fishbone modal pattern in this analysis exhibits phase folding, in analogy with the ST precursor oscillations above in sections 3.3 and 3.1. The phase folding observed for the fishbone in figure 16 suggests that it has a significant $m/n = 1/1$ internal kink mode character. The frequency of the

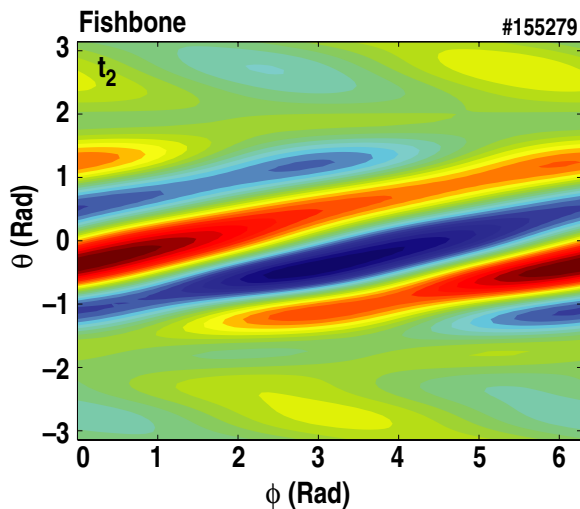


Figure 16. Fishbone modal pattern example. The time t_2 is defined in figure 15. The pattern for the $|n| = 1$ mode at time t_1 looks very similar. Reversed B_T compared to the other shots analysed in this work.

fishbone oscillation is consistent with the plasma core velocity. Compared to e.g. wavelet spectrograms of the fishbone burst as seen in [51] (experimentally) and [52] (theoretically), the signal ‘parametrization’ (figure 15) produced by eigspec seems clearer and may be easier to interpret with respect to models.

4. Conclusions and outlook

A new signal-processing toolkit for magnetic probe array analysis has been developed and deployed for the DIII-D tokamak. The algorithms are based on subspace methods that are known to be effective for narrow-band frequency estimation and multi-channel modal analysis. Importantly, the methods include a modal selection criterion (Order-MAC) that seems to be able to remove most spurious modes in the analysis. This means that the magnetics fluctuation analysis toolkit can output a shortlist of modes that are not obviously realization dependent (i.e. not simply random/spurious modes). This can be exploited in many ways. The development of a neoclassical tearing mode (NTM) detection system is a natural application for these methods. The code also seems to be useful in general magnetic probe array analysis since it is able to provide more accurate results than DFT based methods.

Another possible application is the integration of profile/internal diagnostics. The presented algorithm is based on signal-processing methods that may find good use also in e.g. SXR, electron cyclotron emission (ECE) and beam emission spectroscopy (BES) array analysis; the frequency shortlist derived from the Mirnov array as described in this work could directly be exploited to do conditional narrow-band frequency selective tomographic inversions [53]. Another important extension is to attempt these analysis techniques for slowly rotating MHD activity, which may allow application to additional diagnostics such as charge-exchange recombination spectroscopy (CER), motional Stark effect (MSE), and time-integrated magnetic probe arrays.

Acknowledgments

This work is supported by the US Department of Energy under DE-SC0008520, DE-FG02-04ER54761, DE-FC02-04ER54698 and DE-AC02-09CH11466. The first author thanks B G Penaflor for help with the deployment of parts of the analysis code eigspec at DIII-D, C Paz-Soldan for doing early beta testing, and J King for help with DIII-D magnetics upgrade information. Further thanks goes to Bill Meyer for providing the original DIII-D povray model. DIII-D data shown in this paper can be obtained in digital format by following the links at https://fusion.gat.com/global/D3D_DMP.

References

- [1] Wesson J 2011 *Tokamaks, International Series of Monographs on Physics* 4th edn (Oxford: Oxford University Press)
- [2] La Haye R J 2006 *Phys. Plasmas* **13** 055501
- [3] Kim J S *et al* 2001 *Plasma Phys. Control. Fusion* **43** 1593
- [4] Zohm H *et al* 2007 *Nucl. Fusion* **47** 228
- [5] Prater R *et al* 2007 *Nucl. Fusion* **47** 371
- [6] Kolemen E *et al* 2013 *Fusion Eng. Des.* **88** 2757
- [7] Taylor P L *et al* 1999 *Phys. Plasmas* **6** 1872
- [8] Hollmann E M *et al* 2010 *Phys. Plasmas* **17** 056117
- [9] Volpe F A G *et al* 2009 *Phys. Plasmas* **16** 102502
- [10] Ferron J R and Strait E J 1992 *Rev. Sci. Instrum.* **63** 4799
- [11] Strait E J 2006 *Rev. Sci. Instrum.* **77** 023502
- [12] Trefethen L N and Bau D 1997 *Numerical Linear Algebra* (Philadelphia, PA: SIAM)
- [13] Nardone C 1992 *Plasma Phys. Control. Fusion* **34** 1447
- [14] de Wit T Dudok *et al* 1994 *Phys. Plasmas* **1** 3288
- [15] Kim J S *et al* 1999 *Plasma Phys. Control. Fusion* **41** 1399
- [16] Edgell D H *et al* 2002 *Rev. Sci. Instrum.* **73** 1761
- [17] Hole M J and Appel L C 2007 *Plasma Phys. Control. Fusion* **49** 1971
- [18] Klein A *et al* 2008 *Plasma Phys. Control. Fusion* **50** 125005
- [19] Testa D *et al* 2010 *Europhys. Lett.* **92** 50001
- [20] Alves D and Coelho R 2011 *IEEE Trans. Nucl. Sci.* **58** 1582
- [21] Ives D and Coelho R 2013 *Rev. Sci. Instrum.* **84** 083508
- [22] Stoica P and Moses R L 2005 *Spectral Analysis of Signals* (Englewood Cliffs, NJ: Pearson Prentice-Hall)
- [23] Arun K S and Kung S Y 1990 *SIAM J. Matrix Anal. Appl.* **11** 42
- [24] van Overschee P and De Moor B 1996 *Subspace Identification for Linear Systems* (Dordrecht: Kluwer)
- [25] Qin S J 2006 *Comput. Chem. Eng.* **30** 1502
- [26] Peeters B and De Roeck G 1999 *Mech. Syst. Signal Process.* **13** 855
- [27] Van der Auweraer H and Peeters B 2004 *Proc. 21st IEEE Instrumentation and Measurement Technology Conf. (IMTC Como, Italy)* vol 3 p 2193
- [28] Allemang R J 2003 The modal assurance criterion—twenty years of use and abuse *Sound Vibration Mag.* 14
- [29] Friswell M I, Mottershead J E and Ahmadian H 2001 *Phil. Trans. R. Soc. Lond. A* **359** 169
- [30] Lau J *et al* 2007 *Proc. 25th IMAC (Orlando, FL, Feb. 2007)* p 1029
- [31] Candes E J and Wakin M B 2008 *IEEE Signal Process. Mag.* **25** 21
- [32] Luxburg U 2007 *Stat. Comput.* **17** 395
- [33] Hastie T J, Tibshirani R J and Friedman J J H 2009 *The Elements of Statistical Learning (Springer Series in statistics)* (New York: Springer)
- [34] MacKay D J C 2002 *Information Theory, Inference and Learning Algorithms* (Cambridge: Cambridge University Press)

- [35] Rasmussen C E and Williams C K I 2006 *Gaussian Processes for Machine Learning* (Cambridge, MA: MIT Press)
- [36] Rousseeuw P J and Van Driessen K 1999 *Technometrics* **41** 212
- [37] King J D *et al* 2014 An upgrade of the magnetic diagnostic system of the DIII-D tokamak for non-axisymmetric measurements *Rev. Sci. Instrum.* submitted
- [38] Kolemen E *et al* 2014 *Nucl. Fusion* **54** 073020
- [39] La Haye R J *et al* 2006 *Nucl. Fusion* **46** 451s
- [40] van den Brand H *et al* 2012 *Plasma Phys. Control. Fusion* **54** 094003
- [41] Suttrop W *et al* 2003 *Plasma Phys. Control. Fusion* **45** 1399
- [42] Burrell K H *et al* 2012 *Phys. Plasmas* **19** 056117
- [43] von Goeler S, Stodiek W and Sauthof N 1974 *Phys. Rev. Lett.* **33** 1201
- [44] Eriksson H G and Wahlberg C 2002 *Phys. Plasmas* **9** 1606
- [45] Lazarus E A *et al* 2007 *Phys. Plasmas* **14** 055701
- [46] Lennholm M *et al* 2009 *Phys. Rev. Lett.* **102** 115004
- [47] Chapman I T *et al* 2012 *Nucl. Fusion* **52** 063006
- [48] McGuire K *et al* 1983 *Phys. Rev. Lett.* **50** 891
- [49] Cheng C Z 1992 *Phys. Rep.* **211** 1
- [50] Chen L, White R B and Rosenbluth M N 1984 *Phys. Rev. Lett.* **52** 1122
- [51] Kass T *et al* 1998 *Nucl. Fusion* **38** 807
- [52] Marchenko V S and Reznik S N 2011 *Nucl. Fusion* **51** 122001
- [53] Tobias B *et al* 2013 *Plasma Phys. Control. Fusion* **55** 095006