

UC Irvine

UC Irvine Previously Published Works

Title

Inversion methods for fast-ion velocity-space tomography in fusion plasmas

Permalink

<https://escholarship.org/uc/item/0rk8g5q3>

Journal

Plasma Physics and Controlled Fusion, 58(4)

ISSN

0741-3335

Authors

Jacobsen, AS
Stagner, L
Salewski, M
[et al.](#)

Publication Date

2016-04-01

DOI

10.1088/0741-3335/58/4/045016

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Inversion methods for fast-ion velocity-space tomography in fusion plasmas

A S Jacobsen¹, L Stagner², M Salewski¹, B Geiger³, W W Heidbrink²,
S B Korsholm¹, F Leipold¹, S K Nielsen¹, J Rasmussen¹, M Stejner¹,
H Thomsen⁴, M Weiland³ and the ASDEX Upgrade team³

¹ Department of Physics, Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark

² University of California Irvine, Irvine, CA 92697, USA

³ Max-Planck-Institute for Plasma Physics, Boltzmannstr. 2, 85748 Garching, Germany

⁴ Max-Planck-Institute for Plasma Physics, Wendelsteinstrasse 1, 17491 Greifswald, Germany

E-mail: ajsen@fysik.dtu.dk and lstagner@uci.edu

Received 31 August 2015, revised 6 November 2015

Accepted for publication 19 November 2015

Published 19 February 2016



Abstract

Velocity-space tomography has been used to infer 2D fast-ion velocity distribution functions. Here we compare the performance of five different tomographic inversion methods: truncated singular value decomposition, maximum entropy, minimum Fisher information and zeroth- and first-order Tikhonov regularization. The inversion methods are applied to fast-ion D_α measurements taken just before and just after a sawtooth crash in the ASDEX Upgrade tokamak as well as to synthetic measurements from different test distributions. We find that the methods regularizing by penalizing steep gradients or maximizing entropy perform best. We assess the uncertainty of the calculated inversions taking into account photon noise, uncertainties in the forward model as well as uncertainties introduced by the regularization which allows us to distinguish regions of high and low confidence in the tomographies. In high confidence regions, all methods agree that ions with pitch values close to zero, as well as ions with large pitch values, are ejected from the plasma center by the sawtooth crash, and that this ejection depletes the ion population with large pitch values more strongly.

Keywords: fast ions, tomography, sawtooth crash, fast-ion D-alpha spectroscopy

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

1. Introduction

Traditional fast-ion diagnostics and analysis procedures provide only incomplete information about the 2D fast-ion velocity distribution function. Usually, only a 1D function of the velocity distribution function can be measured. Here we use fast-ion D_α (FIDA) spectroscopy measurements that measure the intensity of Doppler-shifted D_α light as function of wavelength [1, 2]. Using velocity-space tomography it is possible to combine data from several such measurements to infer the 2D fast-ion velocity distribution function [3–10]. With this approach it should even be possible to combine measurements from different diagnostics which is beneficial as they are sensitive to different regions of velocity-space [6]. This velocity-space sensitivity is quantified by velocity-space sensitivity functions, also called weight functions, which have

been developed for FIDA [11, 12], collective Thomson scattering (CTS) [4], fast-ion loss detectors (FILD) [13], neutron emission spectrometry (NES) [14, 15] and gamma-ray spectroscopy (GRS) [16]. The weight functions, $w(x_1, x_2, \phi, E, p)$, relate a measurement, $s(x_1, x_2, \phi)$, to the fast-ion velocity distribution function, $f(E, p)$, where x_1 and x_2 define a range in the specific measurement variable, ϕ is the angle between the projection direction and the magnetic field, E is the fast-ion energy and p is the pitch defined as $p = \frac{v_{\parallel}}{v}$. Here v_{\parallel} is the ion velocity parallel to the magnetic field and v is the ion speed. p is defined positive in the co-current direction. In FIDA weight functions the measurement variable is the wavelength, λ , of the measured light:

$$s(\lambda_1, \lambda_2, \phi) = \iint w(\lambda_1, \lambda_2, \phi, E, p) f(E, p) dE dp. \quad (1)$$




Figure 1. Examples of FIDA weight functions calculated for ASDEX Upgrade discharge #31557 for two different ϕ -angles and a wavelength of 662 nm. The weight functions are shown on a linear scale. (a) $\phi = 80^\circ$, (b) $\phi = 10^\circ$.

Examples of FIDA weight functions are shown in figure 1. The shapes of the weight functions are determined by the angle, ϕ , between the line-of-sight of the diagnostic and the magnetic field in the measurement volume, as well as the wavelength range in which the signal is detected.

By discretizing equation (1) a linear system of equations is obtained:

$$S = WF. \quad (2)$$

S and W are normalized by the measurement uncertainties as described in [6]. S and F are vectors of length m and n , respectively, and W is an $m \times n$ matrix dubbed the transfer matrix. Calculating F from equation (2) is a mathematically ill-posed inverse problem. An ill-posed problem is one for which a unique solution might not exist. Furthermore, any solution might be extremely sensitive to small changes in the data. The sensitivity with respect to the data is suppressed by regularizing the problem. Many different inversion methods have been developed and applied in many scientific fields. Here we compare five inversion methods to measure fast-ion velocity distribution functions by tomographic inversion: truncated singular value decomposition, maximum entropy, zeroth- and first-order Tikhonov regularization and minimum Fisher information. These methods have previously been compared for position-space tomography in fusion plasmas [17, 18]. Here we test these methods for velocity-space tomography. In our comparison we use a transfer matrix describing FIDA measurements taken simultaneously in five views at ASDEX Upgrade in discharge #31557. This large number of FIDA views makes ASDEX Upgrade particularly suitable for tomographic inversions of fast-ion velocity distribution functions. This diagnostic set-up (i.e. this transfer matrix) is used, firstly, with synthetic measurements to calculate inversions for known velocity distribution functions and, secondly, with real five-view FIDA measurements taken just before and just after a sawtooth crash. The synthetic measurements enable us to quantify the performance of the

different methods for assumed measurement uncertainties since the true solution is known. The real measurements allow us to investigate the redistribution of fast ions due to a sawtooth crash resolved in 2D velocity space for ion energies above 20 keV.

The paper is organised as follows. The FIDA diagnostic set-up is described in section 2. Section 3 explains the different inversion methods. In section 4 we quantify the performance of the inversion methods by inverting synthetic data based on known test functions. In section 5 the methods are used to investigate the effect of a sawtooth crash on the central fast-ion population. The results are discussed in section 6 and conclusions are summarized in section 7.

2. ASDEX upgrade FIDA system

A FIDA diagnostic set-up measures Doppler-shifted deuterium Balmer-alpha light from the plasma. It is extensively used to diagnose fast ions at ASDEX Upgrade [19, 20], DIII-D [21, 22], NSTX [23], MAST [24, 25] and LHD [26]. The newly upgraded set-up at ASDEX Upgrade now consists of five different views, each with several lines of sight measuring at different radial locations. We use one line of sight from each view, each intersecting the beam path of neutral beam injector (NBI) Q3 in the plasma centre as shown in figure 2. The grey line is the Q3 NBI beam and each coloured line is a single FIDA line-of-sight. These originate from different positions in the plasma wall and intersect the NBI beam at approximately the same position. The measurement volume is placed slightly on the low-field side. Each view has a different angle between its line of sight and the magnetic field. Thereby, they probe different regions in velocity space as described by their weight functions [12]. In the plasma centre, the respective angles are 14° , 73° , 103° , 133° and 153° . A description of the upgraded FIDA system is found in [10].

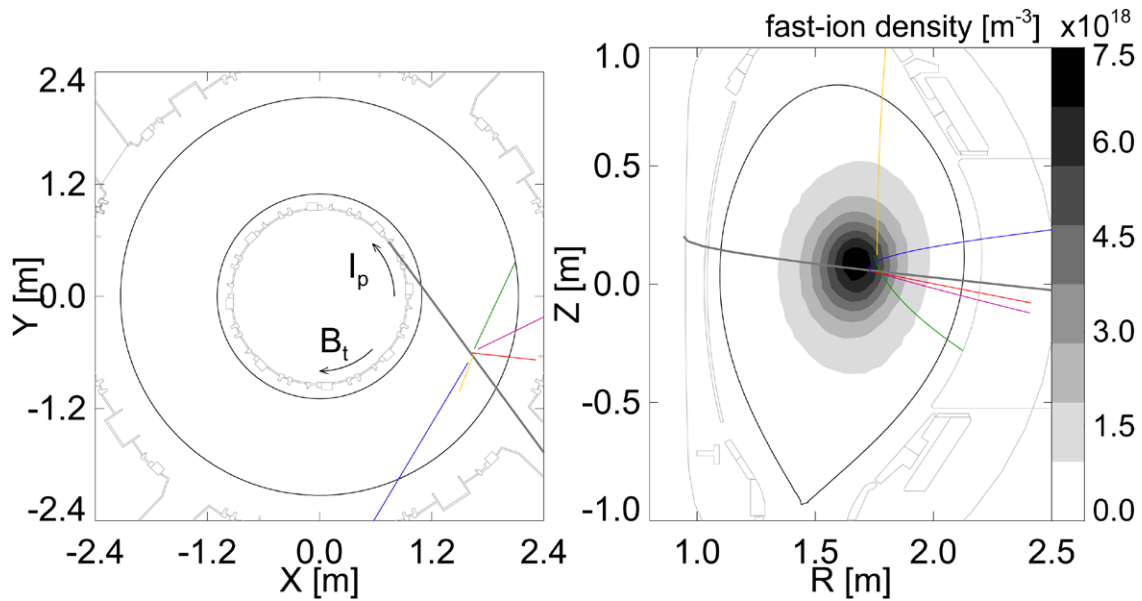


Figure 2. Sketch of the geometry of the FIDA diagnostic set-up. (a) Top view of the ASDEX Upgrade tokamak showing the NBI beam in grey and the FIDA lines-of-sight in colours. Only the lines-of-sight used here are shown. (b) Poloidal cross-section showing that the FIDA measurement volume used here is slightly on the low-field side of the ASDEX Upgrade tokamak.

3. Inversion methods

3.1. Singular value decomposition

Truncated singular value decomposition (SVD) has been used previously to calculate velocity-space tomographies in fusion plasmas [5–9]. The $m \times n$ transfer matrix, W , can as any matrix be written as the product of three matrices:

$$W = U\Sigma V^T, \quad (3)$$

where the columns of the $m \times m$ matrix U are the eigenvectors of the matrix WW^T and the columns of the $n \times n$ matrix V are the eigenvectors of the matrix $W^T W$ [27]. U as well as V are orthogonal matrices. Σ is a diagonal $m \times n$ rectangular matrix. The values in the diagonal are called the singular values. They are the square roots of the non-zero eigenvalues of both WW^T and $W^T W$ [28]. The values in the diagonal of Σ are ordered in a decreasing manner.

Given equation (1) the solution, F_{SVD} , is found as

$$F_{\text{SVD}} = W^+S, \quad (4)$$

where W^+ is called the *pseudoinverse* of W [28]. W^+ can be calculated using the SVD factorization:

$$W^+ = V\Sigma^+U^T, \quad (5)$$

where Σ^+ is a rectangular diagonal matrix with the reciprocals of the diagonal elements of Σ on the diagonal. Equation (3) can be written as a sum

$$W = \sum_{j=1}^r u_j \sigma_j (v_j^T), \quad (6)$$

where r is the number of non-zero singular values, u_j and v_j are the j 'th columns of U and V , respectively, and σ_j is the j 'th

singular value. v_j^T indicates the transpose of the j 'th column of V . F_{SVD} can likewise be expressed as a sum:

$$F_{\text{SVD}} = \sum_{j=1}^r \frac{(u_j^T S)}{\sigma_j} v_j. \quad (7)$$

Experimental data always contain some form of noise. Here we define

$$S = S_{\text{exact}} + e, \quad (8)$$

where S_{exact} is the idealized measurement without noise and e is the noise. Inserting equation (8) in equation (7) we get

$$F_{\text{SVD}} = \sum_{j=1}^r \frac{(u_j^T S_{\text{exact}})}{\sigma_j} v_j + \sum_{j=1}^r \frac{(u_j^T e)}{\sigma_j} v_j = F_{\text{exact}} + \sum_{j=1}^r \frac{(u_j^T e)}{\sigma_j} v_j, \quad (9)$$

where F_{exact} is the exact solution we seek and the last sum describes the effect of the noise. For very small singular values, the SVD solution can be completely dominated by the noise. To reduce its influence, a possibility is to truncate the sum after k terms. However, this makes it impossible to reconstruct F_{exact} completely. This method is called truncated SVD. Truncated SVD introduces the problem of choosing the optimum truncation level, k . Here we use the L-curve method to choose k [29]. In the L-curve method, we calculate a tomography for every truncation level. For each, the norm of the tomography ($R(x_\alpha)$) and the goodness of fit to the data ($\chi^2(x_\alpha)$) are calculated. When plotting the norm and the goodness of fit in a loglog plot, a plot in the shape of an L is obtained as seen in the left part of figure 3. The right part shows the associated curvature. The point with maximum curvature is the corner point. The corner is chosen as optimal because it represents a balance between fitting the data and regularizing the solution.




Figure 3. Example of how to chose the regularization strength based on the L-curve method.

3.2. Tikhonov regularization

The inverse problem posed in equation (2) can be formulated as a least squares problem, i.e. find the solution F which minimizes the norm of the residual:

$$\text{minimize}\{\|WF - S\|^2\}. \quad (10)$$

Well-posed problems can be solved using the normal equations:

$$F = (W^T W)^{-1} W^T S. \quad (11)$$

However, for ill-posed problems a small change in S can have a significant impact on F . In Tikhonov regularization, the ill-posed least squares problem is replaced by a closely related well-posed least squares problem

$$\text{minimize}\{\|WF - S\|^2 + \alpha\|LF\|^2\}, \quad (12)$$

where L is a regularization matrix of size $n \times n$ and α is a non-negative number determining the weight of the regularization term. As for truncated SVD, we determine the value of α using the L-curve method [30]. The Tikhonov solution, F_α , becomes

$$F_\alpha = (W^T W + \alpha L^T L)^{-1} W^T S. \quad (13)$$

The choice of regularization matrix determines the nature of the regularization. Common choices of L penalize the magnitude of f or its derivative to different orders. Therefore, Tikhonov regularization is also sometimes called linear regularization. The simplest regularization matrix is

$$L = I, \quad (14)$$

where I is the $n \times n$ identity matrix so that $L^T L = I$ and equation (13) becomes

$$F_\alpha = (W^T W + \alpha I)^{-1} W^T S. \quad (15)$$

This penalizes large absolute values of f and is called zeroth-order regularization.

First-order regularization penalizes large gradients. In 2D velocity space (v_{\parallel}, v_{\perp}), the penalty operator is

$$L^T L = \nabla_{v_{\parallel}}^T \nabla_{v_{\parallel}} + \nabla_{v_{\perp}}^T \nabla_{v_{\perp}}. \quad (16)$$

Here $\nabla_{v_{\parallel}}$ and $\nabla_{v_{\perp}}$ are matrix representations of finite difference operators. In (E, p) -coordinates, the velocity-space gradient is

$$\nabla F = \sqrt{2mE} (\nabla_E F) \hat{e}_E + \sqrt{\frac{m}{2E}} \sqrt{1-p^2} (\nabla_p F) \hat{e}_p. \quad (17)$$

The derivation of equation (17) is included in appendix A. In (E, p) -coordinates the penalty operator becomes

$$L^T L = 2mE \nabla_E^T \nabla_E + \frac{m}{2E} (1-p^2) \nabla_p^T \nabla_p. \quad (18)$$

3.3. Minimum Fisher information regularization

The principle of minimum Fisher information has been used to compute inversions in soft x-ray tomography in tokamak plasmas [17]. In reference [17] the minimum Fisher information principle is effectively built in as a Tikhonov penalty function. It can therefore be seen as a variant or extension of the general Tikhonov regularization method. The minimum Fisher information method penalizes large gradients divided by the function values. The normalization with the distribution itself means that the smoothing effect is strongest where the distribution has low values.

The minimum Fisher information method is implemented here as an iterative algorithm [17]. First a solution $F^{(1)}$ is found using Tikhonov regularization with a first-order linear penalty function. In the subsequent iterations, the penalty function in $(v_{\parallel}, v_{\perp})$ -coordinates becomes

$$L^T L = \nabla_{v_{\parallel}}^T M^{(n)} \nabla_{v_{\parallel}} + \nabla_{v_{\perp}}^T M^{(n)} \nabla_{v_{\perp}}, \quad (19)$$

where

$$M_{i,j}^{(n)} = \frac{1}{F_i^{(n-1)}} \delta_{i,j} \quad \text{if} \quad F_i^{(n-1)} > 0 \quad (20)$$

$$M_{i,j}^{(n)} = M_{\max}^{(n)} \delta_{i,j} \quad \text{if} \quad F_i^{(n-1)} \leq 0. \quad (21)$$




Figure 4. Test velocity distributions functions as a function of energy and pitch of the ions. The functions are given in units of [ions/keV cm⁻³]. (a) Gaussian. (b) Bi-Maxwellian. (c) NBI.

$M_{\max}^{(n)}$ is the largest $M^{(n)}$ for $F_i > 0$. In (E, p) -coordinates the penalty function becomes

$$L^T L = 2mE \nabla_E^T M^{(n)} \nabla_E + \frac{m}{2E} (1 - p^2) \nabla_p^T M^{(n)} \nabla_p. \quad (22)$$

In each iteration, the corresponding Tikhonov solution with the appropriate minimum Fisher information penalty function is found. We find that the solution converges after only a few iterations.

3.4. Maximum entropy regularization

The last inversion method we have implemented is maximum entropy regularization. In the case of maximum entropy regularization it is assumed that the object that is to be reconstructed from data is positive $f(E, p) \geq 0$. The other inversion methods can produce unphysical negative values which are then set to zero. The specific formulation of maximum entropy regularization adopted here can be found in references [17, 31]. Maximum entropy regularization can be formulated as a minimization problem of the form

$$\text{minimize} \left\{ \frac{1}{2} \|WF - S\|^2 + \alpha H \right\}. \quad (23)$$

where α is a free parameter controlling the strength of our assumptions similar to the free parameter introduced in Tikhonov and minimum Fisher information regularization. We determine the optimal value of α using the L-curve method [30]. H is the Shannon information entropy given by

$$H = - \sum_{i=1}^N (F_i - m_i - F_i \ln(F_i/m_i)). \quad (24)$$

The entropy H is minimized when $F_i = m_i$. Thus m_i is called the default model as it is the value F_i will take when there is no information or data influencing it. While the default model is usually set to be constant in phase-space to prevent biasing of the solution, we may choose to set the default model to be given by a theoretical model. For this work, the default model is set to be constant. The solution of this minimization problem, called the maximum entropy solution, is found using a general non-linear optimization library [32–34].

4. Tomographies using synthetic measurements

In this section we calculate tomographies using synthetic data obtained using equation (1) and known distribution functions. Inversions of synthetic spectra calculated from known distribution functions enable us to compare the performance of the inversion methods by quantitative figures of merit since we know the true solution.

4.1. Test velocity-distribution functions

Three different velocity distributions will be investigated in this analysis. A Gaussian distribution, a bi-Maxwellian distribution and a simulated NBI-distribution from TRANSP/NUBEAM [35]. The three distributions are shown in figure 4. We choose these three distribution functions as they pose different challenges to the inversion methods. The Gaussian distribution represents a localized source of fast particles typical for the peaks at the injection energies for neutral beam heating. The bi-Maxwellian is a wide function covering the entire pitch range. Here the challenge is to recreate the large-scale undulation. Lastly, we study a distribution function typical for neutral beam injection as simulated by TRANSP. This is an important test case as it should be very similar to the distribution functions in experiments with NBI heating. The challenge here is the structural complexity on both small and large scales. The NBI distribution used in this study is typical for a DIII-D discharge rather than for ASDEX Upgrade which does not have counter-injection beams.

4.2. Modelling of measurement noise

The photon noise of FIDA light scales approximately with the square root of the signal. However, in the absence of FIDA light the photon noise is dominated by bremsstrahlung setting a lower limit on the noise level. These two effects are modelled as

$$S_{\text{noisy}} = S_{\text{exact}} + k \langle \sqrt{S_{\text{exact}}} \rangle \eta, \quad (25)$$

where S_{noisy} is the noisy spectrum, S_{exact} is the exact noise-free spectrum, $\langle \rangle$ denotes the mean and k is a scaling constant that allows us to vary the noise level. η denotes a normally-distributed random deviate with mean zero and standard deviation




Figure 5. Examples of the average noise levels in the synthetic spectra calculated using the NBI test distribution and equation (25) for $k = 0.1, 0.5$ and 0.9 . The width of the spectra corresponds to the standard deviation of the noise for the given k -value.

$\max(e_{\min}, \sqrt{S_{\text{exact}}})$, i.e. the larger value of the bremsstrahlung level and $\sqrt{S_{\text{exact}}}$. By varying the noise level we can investigate how robust the methods are against noise. Figure 5 shows examples of the standard deviation of the synthetic spectra calculated using the NBI test distribution for $k = 0.1, k = 0.5$ and $k = 0.9$. The noise level of actual FIDA measurements depends on the plasma parameters. A k value in the range 0.3 – 0.5 represents the noise level in a typical discharge. When calculating the synthetic spectra, wavelengths blocked by impurity emission in ASDEX Upgrade discharge #31557 are included as the impurity emission is discharge-specific and here we wish to compare the inversion methods for a generic discharge.

4.3. Variance and bias of the solution

Using equation (25) noisy spectra can be calculated. In order to determine how the noise propagates from the spectra to the tomographies, 25 noisy spectra are calculated for every value of k . A tomography is calculated for each spectrum. Thereby, an ensemble of tomographies is obtained for each k -value. The propagation of the uncertainties from the measurements to the tomographies is found by calculating the variance of the ensemble of tomographies for each velocity-space coordinate. This variance corresponds to the propagation of the uncertainties through the regularized inverse of the transfer matrix as explained in [6].

A forward model can itself contain uncertainties. The forward model here is given by the FIDA weight functions. These are calculated numerically based on profiles of several nuisance parameters. The weight functions are most sensitive to the ion temperature and rotation velocity, the electron temperature and density and the effective charge Z_{eff} . Hence we consider the impact of these parameters on the tomographic inversion results for the different regularization methods. The uncertainties in the bulk plasma parameters lead to uncertainties in the

weight functions, δW . Assuming a Gaussian error distribution of the bulk plasma parameters, we calculate the uncertainty in the forward model (i.e. in the weight functions) by sampling a population of weight functions calculated varying one nuisance plasma parameter at a time and keeping the other parameters fixed. The total variance of the weight function is then obtained by summing up the variances obtained from each plasma parameter. The corresponding error, e_m , from the forward model error is

$$e_m = \delta W F_{\text{true}}. \quad (26)$$

e_m depends on the (often unknown) true distribution function. However, if an estimate of F_{true} can be obtained, e_m can be estimated. The combined uncertainty due to uncertainty in the forward model and measurement uncertainty is then

$$e_{\text{data}} = \sqrt{e_{\text{ph}}^2 + e_m^2}. \quad (27)$$

where e_{ph} is the photon noise. In the case of synthetic measurements, e_{ph} is calculated using equation (25). F_{true} is approximated by the sum of a Maxwellian for the bulk ions and a TRANSP/NUBEAM simulation to estimate the NBI ions. The Maxwellian is calculated using measured ion temperature and density. The forward model error is included in the analysis using real measurements but not for the synthetic measurements.

As mentioned above, uncertainties in the measurements propagate through to the tomography. However, this effect is attenuated by the regularization methods as the tomographies would otherwise be completely dominated by the noise in the measurements. The regularization itself introduces an error in the tomography though. Here we calculate this as the bias of the tomography. It is given as the difference of the mean of the ensemble of calculated tomographies for a given k and the true test distribution:

$$\text{bias} = F_{\mu} - F_{\text{true}}, \quad (28)$$




Figure 6. Tomographies of the Gaussian distribution from figure 4(a) in units of [ions/keV cm⁻³] based on synthetic measurements using various inversion methods and noise levels. The noise level k is defined in equation (25).

Table 1. Parameters of the Gaussian test distribution.

	True	SVD	T0	T1	MFI	ME
μ_E [keV]	50	49.37 ± 0.22	49.55 ± 0.19	48.57 ± 0.21	48.45 ± 0.06	50.78 ± 0.07
σ_E [keV]	10	15.16 ± 0.32	15.34 ± 0.27	16.61 ± 0.30	11.24 ± 0.08	8.87 ± 0.10
μ_p [-]	0	0.008 ± 0.005	-0.001 ± 0.004	-0.001 ± 0.004	-0.011 ± 0.001	0.035 ± 0.002
σ_p [-]	0.25	0.325 ± 0.007	0.329 ± 0.006	0.317 ± 0.006	0.228 ± 0.02	0.242 ± 0.003

where F_μ is the mean of the calculated tomographies. Thus the bias also depends on F_{true} .

We define a measure of the total uncertainty in the tomography as the mean squared error, MSE, given by

$$\text{MSE} = \text{variance} + \text{bias}^2. \quad (29)$$

4.4. Figures of merit

We define two different figures of merit. The first is the total mean squared error. This is a measure of the total uncertainty in the obtained tomography. It is calculated by summing the MSE of every pixel. For this figure of merit the smallest values correspond to the best performance of the inversion method.

The second figure of merit is the ratio of the inferred fast-ion density to the true fast-ion density which is calculated as the integral of the tomography normalized by the integral of the true distribution,

$$\frac{\int \int F_{\text{tomography}} dE dp}{\int \int F_{\text{true}} dE dp}. \quad (30)$$

Optimally, this figure of merit is one.

4.5. Inversion results

Figure 6 shows tomographies of the Gaussian distribution shown in figure 4(a) calculated with the different inversion

methods for various noise levels. All methods reconstruct the position of the Gaussian distribution well. The characteristic widths of the Gaussians are approximately right but tend to be slightly larger than in the original test distribution. Measurement noise enhances this trend. We further observe the appearance of jitter in the inversions throughout velocity space. The minimum Fisher information and maximum entropy regularization methods stand out from the other methods in that they resemble the original function the most and exhibits the least jitter. This suggests superior resolution performance of these methods. Table 1 contains the true center coordinates and width of the Gaussian distribution in both energy and pitch. Furthermore, it contains the values obtained from the $k = 0.5$ tomographies calculated using the five different methods. All methods find the center coordinates well. The minimum Fisher information and maximum entropy methods produce significantly more peaked distributions which is seen in their ability to match the true width of the Gaussian better.

Figure 7 shows the reconstructions of the bi-Maxwellian distribution function. The large-scale shape of the distribution is reproduced by all five inversion methods. The pitch angle symmetry with respect to $p = 0$ is reproduced well and the larger perpendicular temperature compared with the parallel temperature is reflected in the large fast-ion densities for pitches close to zero. The first-order Tikhonov, minimum Fisher information and maximum entropy methods reproduce

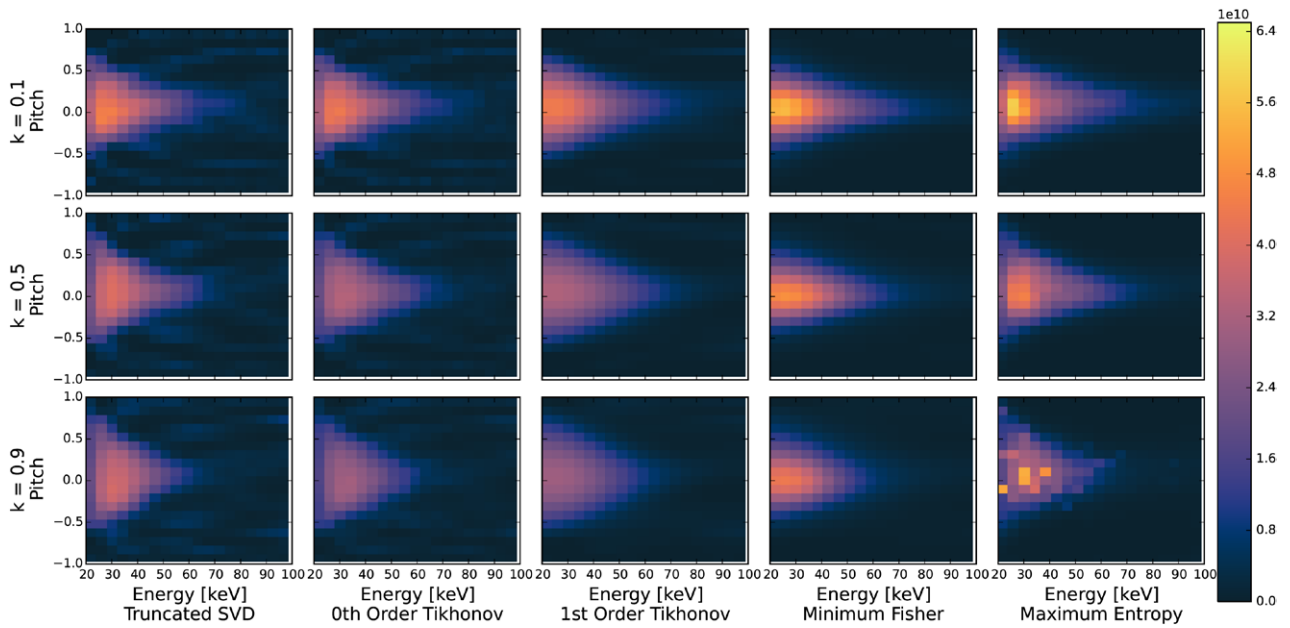


Figure 7. Tomographies of the bi-Maxwellian from figure 4(b) in units of [ions/keV cm⁻³] based on synthetic measurements using various inversion methods and noise levels. The noise level k is defined in equation (25).

Table 2. Parameters of the bi-Maxwellian test distribution.

	True	SVD	T0	T1	MFI	ME
E_{\parallel} [keV]	3	5.12 ± 0.19	5.34 ± 0.18	4.98 ± 0.09	2.94 ± 0.06	4.09 ± 0.13
E_{\perp} [keV]	20	24.36 ± 0.73	26.26 ± 0.73	23.79 ± 0.35	22.51 ± 0.32	24.73 ± 0.61

the distribution particularly well. Table 2 contains the true parallel and perpendicular temperatures used in calculating the bi-Maxwellian and the values obtained from the $k = 0.5$ tomographies. For the bi-Maxwellian distribution, the minimum Fisher information method most closely recreates the true values.

Figure 8 shows reconstructions of the NBI distribution function for various noise levels and inversion methods. This fast-ion distribution function is typical for neutral beam injection with two co-current beams with injection energies at 80 keV and 70 keV and one counter-current beam with an injection energy of 70 keV. Therefore, this distribution function is a more difficult test case than previously presented reconstructions of distribution functions which are more typical for a single NBI beam. The overall shape of the NBI distribution function is well reproduced by all five inversion methods. The protrusion at pitches of about 0.7 originates from the co-current beam injection, and the weaker protrusion at pitches of -0.7 from the counter-current beam injection. All reconstructions show that the full energy beam injection peak for co-current injection (positive pitch) is at larger energies than that for counter-current injection (negative pitch). The first-order Tikhonov, minimum Fisher information and maximum entropy regularization results in smooth tomographies. This makes the overall shape of the function with protrusions at positive and negative pitches stand out most clearly. The local maxima due to the beam injection peaks at full, half and third energies are recreated by the maximum entropy method in the case of low noise ($k = 0.1$). They are also visible in the SVD

and zeroth-order Tikhonov tomographies at low noise. For larger noise levels, none of the methods are able to resolve more than one peak.

Figure 9 shows the behaviour of the performance parameters as a function of noise level for the tomographies of the three test functions. Figures 9(a), (c) and (e) show the total mean squared error. The mean squared error increases for larger noise levels for all inversion methods and test distributions. The minimum Fisher information regularization method has the lowest mean squared error for all test distributions. Figures 9(b), (d) and (f) show the density ratios calculated using equation (30). The general trend is that the methods produce a lower density ratio for large error levels. Thus, for very large noise levels the absolute values of an inferred density obtained from a reconstruction might be unreliable. For the Gaussian test distribution, the minimum Fisher information and maximum entropy methods are very good at recreating the correct density. The other three methods overestimate the amount of ions present. This is also the case for the bi-Maxwellian distribution but not to the same extent. For the NBI test distribution the spread in densities is smaller than for the other cases.

The uncertainties of the tomographies of the beam distribution as defined in section 4.3 are shown in figure 10 for a noise level of $k = 0.5$ in equation (25). Here we disregard the model uncertainty for simplicity. The top row shows the square root of the variance of the tomographies. Compared with the values of the tomographies in figure 8, the uncertainties are about one order of magnitude smaller, and smallest for




Figure 8. Tomographies of the beam distribution from figure 4(c) in units of $[\text{ions/keV cm}^{-3}]$ based on synthetic measurements using various inversion methods and noise levels. The noise level k is defined in equation (25).

first-order Tikhonov and minimum Fisher information regularization. The middle row shows the bias. Negative values denote regions where too few ions are placed, positive values denote regions where too many ions are placed. The beam peaks are seen in the bias, especially for first-order Tikhonov, minimum Fisher information and maximum entropy regularization as these are only able to resolve the spiky nature of the peaks for low noise levels. The last row shows the square root of the mean squared error. The main contribution to the uncertainty is the bias for this regularization level.

4.6. Tomographies using two to four views

ASDEX Upgrade is currently the only machine with five FIDA views whereas FIDA systems on other machines have fewer views. To test the influence of the number of views on the performance of the different inversion methods, we repeat the analysis using two, three and four views. Tomographies of the NBI distribution function calculated for $k = 0.5$ are shown in figure 11. Even using only two views, it is possible to recreate the overall shape of the NBI test distribution. Adding additional views improve the quality of the tomographies of all inversion methods, in agreement with the findings for truncated SVD in [5]. Furthermore, it is evident that the best results are obtained using the first-order Tikhonov, minimum Fisher information and maximum entropy regularization methods for any number of views.

5. Tomographies of a measured sawtooth crash

A sawtooth crash is a periodic plasma instability which can occur when the central safety factor drops below one. It changes the magnetic field topology and has been observed to redistribute particles and energy from the center of the plasma. It has furthermore been observed on several machines

that passing fast ions are redistributed more strongly compared to trapped ions [9, 36, 37]. Here we use the five different inversion methods to investigate the effect of a sawtooth on the central fast-ion population in ASDEX Upgrade. Figure 12 shows time traces from AUG discharge #31557. The sawtooth crashes are evident in the central electron density as well as the central electron and ion temperatures. Figure 13 shows tomographies based on experimental data. They are calculated using the different inversion methods applied to FIDA spectra measured just before and after the sawtooth crash in ASDEX Upgrade discharge #31557 at 2.25 s.

Figure 14 shows the uncertainties of the tomographies of the pre-crash distribution. Again it is seen that the bias dominates the uncertainty. We stress again that the bias depends on the unknown true solution, for which we here use the TRANSP model.

Common for all regularization methods, the inferred fast ion density drops significantly during the sawtooth crash. By comparing the absolute values of the tomographies with the uncertainties, we can identify the velocity-space regions where we can be confident in the tomography. Figure 15 shows the calculated tomographies normalized with $\sqrt{\text{MSE}}$ for the five different inversion methods. To calculate the bias after the sawtooth crash, the Kadomtsev model as implemented in TRANSP is used to model the effect of the sawtooth crash on the fast ions. The parts of velocity space where the values are large correspond to regions where we are confident in the results and regions with low values correspond to uncertain regions. It is seen that the part of velocity space at the full energy peak at 60 keV is very uncertain for all inversion methods. This is because the methods are not able to resolve the peak for the given regularization level.

To further investigate the velocity-space dependence of the change in the fast-ion distribution function, we calculate the relative change:

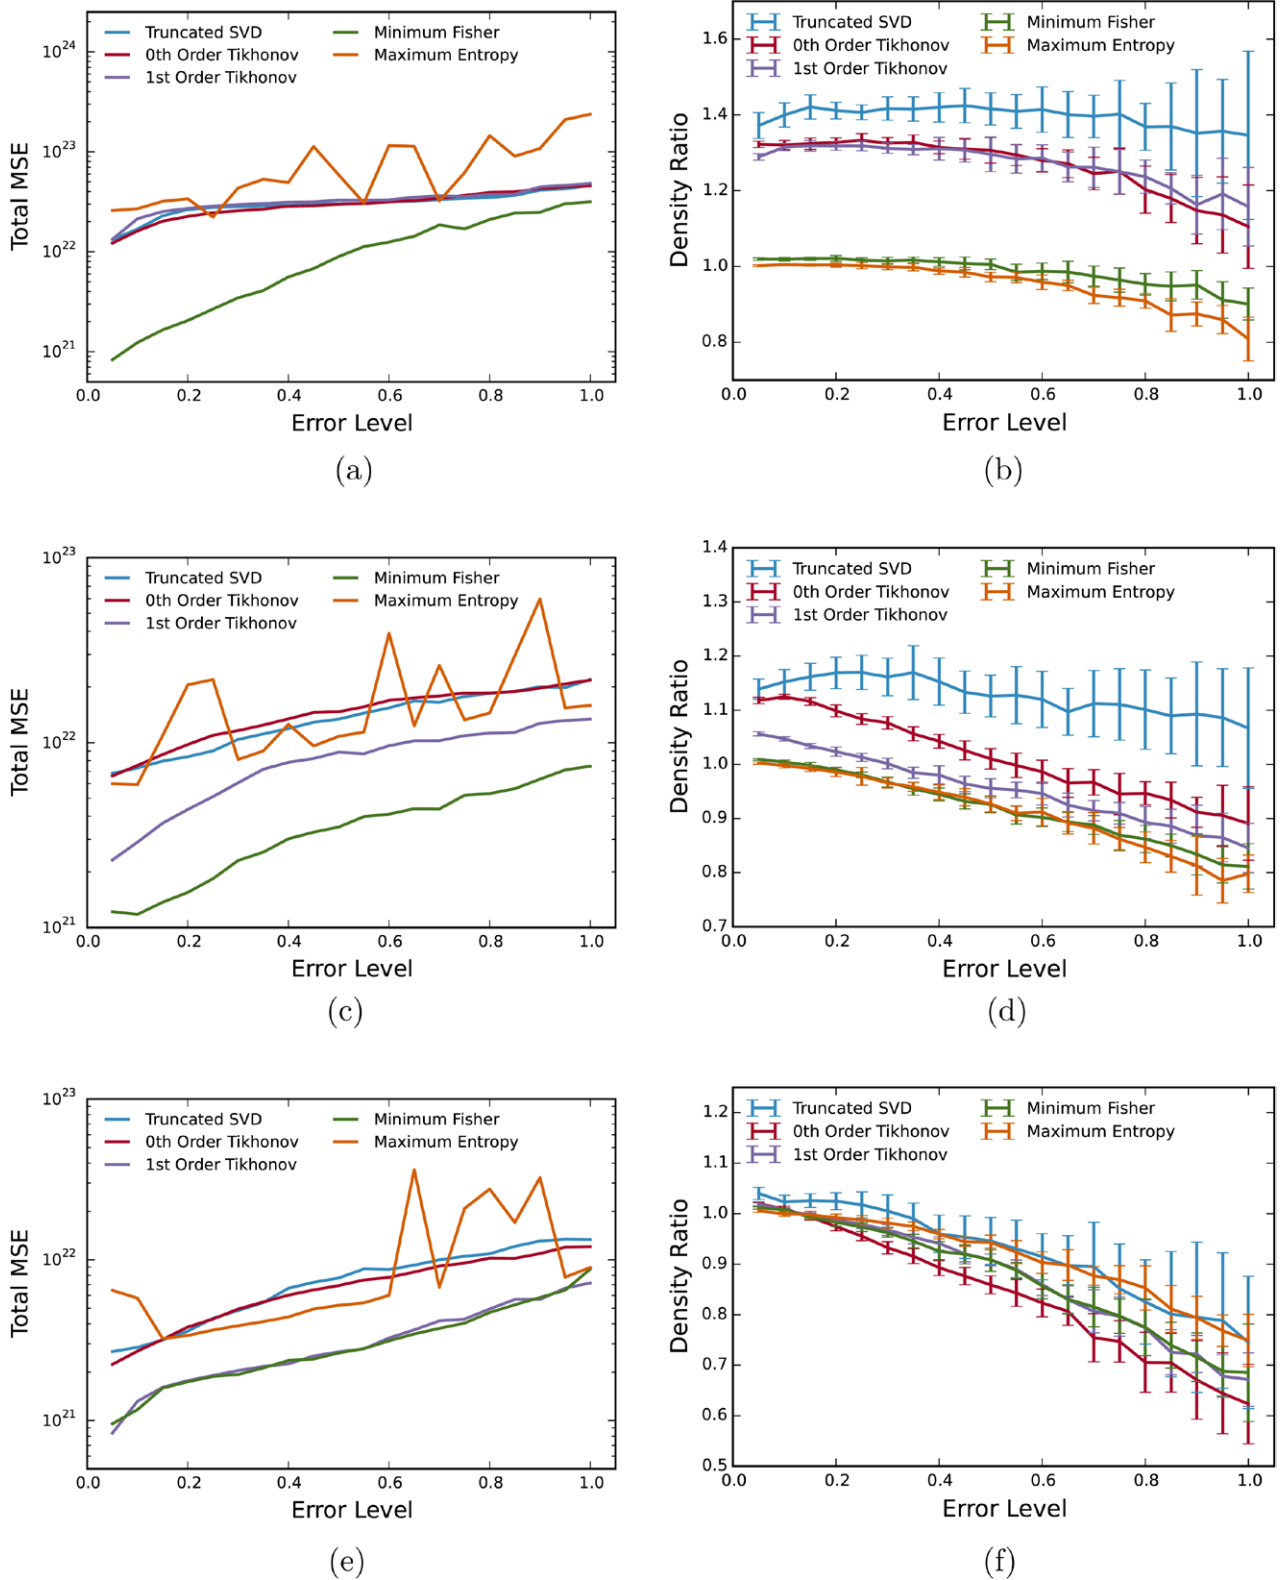


Figure 9. Figures of merit of the reconstructions of the test distributions. The left column shows the total mean squared error. The right column shows the density ratio. (a) Total MSE, Gaussian. (b) Density ratio, Gaussian. (c) Total MSE, bi-Maxwellian. (d) Density ratio, bi-Maxwellian. (e) Total MSE, NBI. (f) Density ratio, NBI.

$$(\Delta F)_{\text{rel}} = \frac{F_{\text{after}} - F_{\text{before}}}{F_{\text{before}}}. \quad (31)$$

The relative change is calculated for every regularization method and plotted in figure 16. The top row shows the

relative change as a function of energy and pitch. The bottom row shows the uncertainties of the relative change. The uncertainties of the relative change are calculated by generating an ensemble of relative changes and then calculating its variance. Note that here only the variance is included in the uncertainty,




Figure 10. Uncertainties for the tomographies of the beam distribution in units of $[\text{ions}/\text{keV cm}^{-3}]$. All uncertainties are calculated for a noise level of $k = 0.5$.




Figure 11. Tomographies of the NBI distribution in units of $[\text{ions}/\text{keV cm}^{-3}]$ calculated using two, three and four FIDA views. All are calculated for $k = 0.5$.

i.e. similar to the top row of figure 14 and the corresponding variance of the distribution after the crash.

The velocity-space dependence of the relative change is especially clear in the first-order Tikhonov and the minimum Fisher information figures as the amount of jitter in these tomographies is significantly smaller compared to the other methods. Both first-order Tikhonov and minimum Fisher information suggest that ions with large pitch values are redistributed more compared to ions with pitch close to zero. This trend is also confirmed by the singular value decomposition, zeroth-order Tikhonov and maximum entropy in the regions where the tomographies are reliable. The unreliable regions are here shown as those with large standard deviation compared with amplitudes of the tomographies. Similar trends

were observed previously using singular value decomposition [9] and a variant of a first-order Tikhonov [10] where different regularization levels were chosen rather than set by the L-curve method. Figure 17 shows the ratio of the post-crash distribution to the pre-crash distribution integrated over energy as a function of pitch for all five inversion methods. Thus it is a measure of the pitch dependence of the change in the fast ion distribution function. For pitch values close to zero all inversion methods except maximum entropy predict a redistribution level of between 10% and 20%. For pitch values above 0.4 the redistribution level increases to between 30% and 40% as seen by all five inversion methods. For negative pitch values where very few ions are present it is not possible to determine the amount of redistribution.




Figure 12. Time traces of AUG discharge #31557. (a) Toroidal magnetic field, total injected NBI power and the plasma current. (b) Ion and electron temperatures and electron density at $\rho_p = 0.1$.




Figure 13. Tomographies before and after a sawtooth crash calculated using the different regularization methods in units of $[\text{ions/keV cm}^{-3}]$.




Figure 14. Measures of uncertainties using the different regularization methods in units of $[\text{ions/keV cm}^{-3}]$.

6. Discussion

In order to calculate the true bias of a given tomography as defined in equation (28), it is necessary to know the true distribution. This makes it impossible to calculate the true bias of a tomography from experimental measurements. Here, we have used a TRANSP distribution and the Kadomtsev model

to generate an estimate of the true distribution. However, in other cases it might not be possible to calculate a good quantitative estimate. In these cases the best one can do is estimate a qualitative bias based on the general behaviour of a given method and regularization level. On the other hand, uncertainties based solely on the propagation of measurement uncertainties through a given regularization process only represents




Figure 15. Tomographies of the ion velocity distribution normalized with $\sqrt{\text{MSE}}$ before (top row) and after (bottom row) the sawtooth crash.

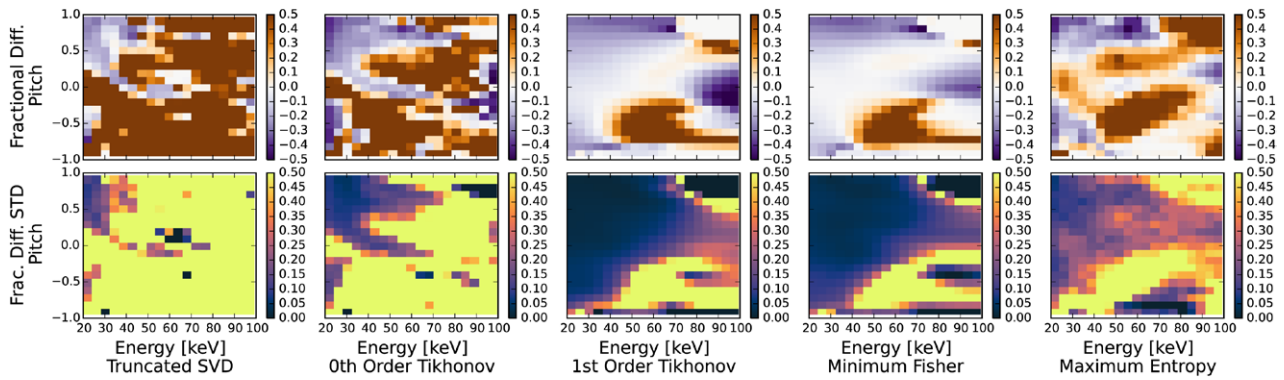


Figure 16. Relative change of the fast-ion velocity distribution function.




Figure 17. Ratio of the fast-ion velocity-space distribution functions before and after the crash integrated over energy shown as a function of pitch.

the spread of obtainable solutions, and thus can be misleading since they can be made almost arbitrarily small, simply by over-regularizing.

It is seen that when the noise level is not too large, the first-order Tikhonov, minimum Fisher information and maximum entropy regularization methods can reconstruct the overall shape of the true distribution function very well. However,

the first-order Tikhonov and minimum Fisher information methods lack capability to resolve very fine and detailed features. For large noise levels, the maximum entropy has a tendency to produce spiky tomography. The maximum entropy method was often observed to generate solutions with a large variance. Truncated SVD and zeroth-order Tikhonov can resolve fine details, especially for measurements with low

noise levels. However, they often produce features in wrong parts of velocity space.

It is seen that the absolute values of a derived quantity such as the fast-ion density depend on the noise level in the data. However, we find that the ratio of such quantities is less sensitive to the specific noise level and amount of regularization. Hence we can make statements about changes in such quantities with greater confidence than about the absolute values themselves since systematic errors introduced by the inversion methods will tend to cancel. For example, the bias in the tomographies tends to be similar before and after a sawtooth crash, and hence it partly cancels in the relative change.

7. Conclusion

We have compared the performance of five different inversion methods for velocity-space tomography. In order to estimate the confidence in the presented analysis, uncertainties of the tomographies are defined and calculated taking into account the photon noise, uncertainties in the forward model as well as uncertainty introduced by the inversion methods themselves. It is found that for the regularization level used here, the bias introduced by the inversion methods is the major contribution. The performance is tested using synthetic data calculated using a realistic transfer matrix from the five-view FIDA-system at ASDEX Upgrade. It is found that the first-order Tikhonov and minimum Fisher information regularization methods which penalize steep gradients as well as the maximum entropy method perform best for realistic test functions. The uncertainty analysis allows us to identify confidence regions in velocity space, and regions where the tomographies are not reliable for the given data and regularization level. Furthermore, the various methods are applied to actual FIDA measurements obtained in ASDEX Upgrade discharge #31557 just before and just after a sawtooth crash. Using velocity-space tomography it is possible to investigate the velocity-space dependence of the fast-ion redistribution in regions where we are confident in the tomography. We find that sawtooth crashes at ASDEX Upgrade affect ions with large pitch values more than ions with pitch close to zero.

Acknowledgments

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. The source code for the inversion methods is available upon request to lstagner@uci.edu.

Appendix A. Derivation of velocity-gradient in (E, p) -coordinates

To calculate the velocity-space gradient in (E, p) -coordinates, the gradient is transformed from $(v_{\parallel}, v_{\perp})$ -coordinates to $(E,$




Figure A1. The relations between the unit vectors $\hat{e}_{v_{\parallel}}$, $\hat{e}_{v_{\perp}}$, \hat{e}_E and \hat{e}_p .

p)-coordinates. First-order regularization in $(v_{\parallel}, v_{\perp})$ -coordinates can be achieved by setting

$$L^T L = \nabla_{v_{\parallel}}^T \nabla_{v_{\parallel}} + \nabla_{v_{\perp}}^T \nabla_{v_{\perp}}, \quad (\text{A.1})$$

where $\nabla_{v_{\parallel}}$ and $\nabla_{v_{\perp}}$ are finite difference matrix representations of the first-order differential operators. These have to be transformed to (E, p) -coordinates. It is apparent that the velocity-space gradient in (E, p) -coordinates has similarities to the real-space gradient in polar coordinates. The relations between the unit vectors $\hat{e}_{v_{\parallel}}$, $\hat{e}_{v_{\perp}}$, \hat{e}_E and \hat{e}_p are illustrated graphically in figure A1. The velocity-space gradient of f is

$$\nabla f = (\nabla_{v_{\parallel}} f) \hat{e}_{v_{\parallel}} + (\nabla_{v_{\perp}} f) \hat{e}_{v_{\perp}} = a(\nabla_E f) \hat{e}_E + b(\nabla_p f) \hat{e}_p, \quad (\text{A.2})$$

where a and b are unknowns that must be calculated from the Jacobian. ∇_E and ∇_p are velocity-gradients along \hat{e}_E and \hat{e}_p respectively. Writing \hat{e}_E and \hat{e}_p as functions of $\hat{e}_{v_{\parallel}}$ and $\hat{e}_{v_{\perp}}$ gives

$$\hat{e}_E = \cos(\theta) \hat{e}_{v_{\parallel}} + \sin(\theta) \hat{e}_{v_{\perp}} = p \hat{e}_{v_{\parallel}} + \sqrt{1-p^2} \hat{e}_{v_{\perp}}, \quad (\text{A.3})$$

$$\hat{e}_p = \sin(\theta) \hat{e}_{v_{\parallel}} - \cos(\theta) \hat{e}_{v_{\perp}} = \sqrt{1-p^2} \hat{e}_{v_{\parallel}} - p \hat{e}_{v_{\perp}}, \quad (\text{A.4})$$

where the relation $p = \cos(\theta)$ has been used. The gradient in energy is now found by dotting equation (A.2) with \hat{e}_E :

$$\begin{aligned} (\nabla f) \cdot \hat{e}_E &\Rightarrow (\nabla_{v_{\parallel}} f) p + (\nabla_{v_{\perp}} f) \sqrt{1-p^2} = a(\nabla_E f) \\ &= a \left((\nabla_{v_{\parallel}} f) \frac{\partial v_{\parallel}}{\partial E} + (\nabla_{v_{\perp}} f) \frac{\partial v_{\perp}}{\partial E} \right), \end{aligned} \quad (\text{A.5})$$

To calculate the partial derivatives, the relations between v_{\parallel} , v_{\perp} , E and p are needed:

$$v_{\parallel} = p \sqrt{\frac{2E}{m}} \quad (\text{A.6})$$

$$v_{\perp} = \sqrt{1-p^2} \sqrt{\frac{2E}{m}}. \quad (\text{A.7})$$

The partial derivatives are:

$$\frac{\partial v_{\parallel}}{\partial E} = \frac{p}{\sqrt{2mE}} \quad (\text{A.8})$$

$$\frac{\partial v_{\perp}}{\partial E} = \frac{\sqrt{1-p^2}}{\sqrt{2mE}}. \quad (\text{A.9})$$

Inserting equations (A.8) and (A.9) in equation (A.5) gives

$$(\nabla_{\parallel} f)p + (\nabla_{\perp} f)\sqrt{1-p^2} = a \left((\nabla_{\parallel} f) \frac{p}{\sqrt{2mE}} + (\nabla_{\perp} f) \frac{\sqrt{1-p^2}}{\sqrt{2mE}} \right). \quad (\text{A.10})$$

Equation (A.10) is fulfilled for

$$a = \sqrt{2mE}. \quad (\text{A.11})$$

Similarly, b can be found by dotting equation (A.2) with \hat{e}_p :

$$\begin{aligned} (\nabla f) \cdot \hat{e}_p &\Rightarrow (\nabla_{\parallel} f)\sqrt{1-p^2} - (\nabla_{\perp} f)p = b(\nabla_p f) \\ &= b \left((\nabla_{\parallel} f) \frac{\partial v_{\parallel}}{\partial p} + (\nabla_{\perp} f) \frac{\partial v_{\perp}}{\partial p} \right). \end{aligned} \quad (\text{A.12})$$

The partial derivatives are:

$$\frac{\partial v_{\parallel}}{\partial p} = \sqrt{\frac{2E}{m}}, \quad (\text{A.13})$$

$$\frac{\partial v_{\perp}}{\partial p} = -\frac{p}{\sqrt{1-p^2}} \sqrt{\frac{2E}{m}}. \quad (\text{A.14})$$

Inserting equations (A.13) and (A.14) in equation (A.12) gives

$$(\nabla_{\parallel} f)\sqrt{1-p^2} - (\nabla_{\perp} f)p = b \left((\nabla_{\parallel} f) \sqrt{\frac{2E}{m}} - (\nabla_{\perp} f) \frac{p}{\sqrt{1-p^2}} \sqrt{\frac{2E}{m}} \right). \quad (\text{A.15})$$

Equation (A.15) is fulfilled for

$$b = \sqrt{\frac{m}{2E}} \sqrt{1-p^2}. \quad (\text{A.16})$$

Thus, the velocity-space gradient in energy-pitch coordinates becomes

$$\nabla f = \sqrt{2mE} (\nabla_E f) \hat{e}_E + \sqrt{\frac{m}{2E}} \sqrt{1-p^2} (\nabla_p f) \hat{e}_p. \quad (\text{A.17})$$

References

- [1] Heidbrink W W, Burrell K H, Luo Y, Pablant N A and Ruskov E 2004 Hydrogenic fast-ion diagnostic using Balmer-alpha light *Plasma Phys. Control. Fusion* **46** 1855–75
- [2] Heidbrink W W 2010 Fast-ion $D\alpha$ measurements of the fast-ion distribution (invited) *Rev. Sci. Instrum.* **81** 10D727
- [3] Egedal J and Bindslev H 2004 Reconstruction of gyrotropic phase-space distributions from one-dimensional projections *Phys. Plasmas* **11** 2191
- [4] Salewski M *et al* 2011 On velocity space interrogation regions of fast-ion collective Thomson scattering at ITER *Nucl. Fusion* **51** 083014
- [5] Salewski M *et al* 2012 Tomography of fast-ion velocity-space distributions from synthetic CTS and FIDA measurements *Nucl. Fusion* **52** 103008
- [6] Salewski M *et al* 2013 Combination of fast-ion diagnostics in velocity-space tomographies *Nucl. Fusion* **53** 063019
- [7] Salewski M *et al* 2014 Measurement of a 2D fast-ion velocity distribution function by tomographic inversion of fast-ion D-alpha spectra *Nucl. Fusion* **54** 023005
- [8] Salewski M *et al* 2015 Doppler tomography in fusion plasmas and astrophysics *Plasma Phys. Control. Fusion* **57** 014021
- [9] Geiger B *et al* 2015 Fast-ion transport and neutral beam current drive in ASDEX Upgrade *Nucl. Fusion* **55** 083001
- [10] Weiland M, Geiger B, Jacobsen A S, Reich M, Salewski M and Odstrcil T 2016 Enhancement of the FIDA diagnostic at ASDEX Upgrade for velocity space tomography *Plasma Phys. Control. Fusion* **58** 025012
- [11] Heidbrink W W, Luo Y, Burrell K H, Harvey R W, Pinsker R I and Ruskov E 2007 Measurements of fast-ion acceleration at cyclotron harmonics using Balmer-alpha spectroscopy *Plasma Phys. Control. Fusion* **49** 1457–75
- [12] Salewski M *et al* 2014 On velocity-space sensitivity of fast-ion D-alpha spectroscopy *Plasma Phys. Control. Fusion* **56** 105005
- [13] Heidbrink W W *et al* 2014 Confinement degradation by Alfvén-eigenmode induced fast-ion transport in steady-state scenario discharges *Plasma Phys. Control. Fusion* **56** 095030
- [14] Jacobsen A S, Salewski M, Eriksson J, Ericsson G, Hjalmarrsson A, Korsholm S B, Leipold F, Nielsen S K, Rasmussen J and Stejner M 2014 Velocity-space sensitivity of the time-of-flight neutron spectrometer at JETa *Rev. Sci. Instrum.* **85** 11E103
- [15] Jacobsen A S, Salewski M, Eriksson J, Ericsson G, Korsholm S B, Leipold F, Nielsen S K, Rasmussen J and Stejner M 2015 Velocity-space sensitivity of neutron spectrometry measurements *Nucl. Fusion* **55** 053013
- [16] Salewski M *et al* 2015 Velocity-space observation regions of high-resolution two-step reaction gamma-ray spectroscopy *Nucl. Fusion* **55** 093029
- [17] Anton M, Weisen H, Dutch M J, von der Linden W, Buhlmann F, Chavan R, Marletaz B, Marmillod P and Paris P 1996 X-ray tomography on the TCV tokamak *Plasma Phys. Control. Fusion* **38** 1849–78
- [18] Craciunescu T, Bonheure G, Kiptily V, Murari A, Tiseanu I and Zoita V 2009 A comparison of four reconstruction methods for JET neutron and gamma tomography *Nucl. Instrum. Methods Phys. Res. A* **605** 374–83
- [19] Geiger B, Dux R, McDermott R M, Potzel S, Reich M, Ryter F, Weiland M, Wunderlich D and Garcia-Munoz M 2013 Multi-view fast-ion D-alpha spectroscopy diagnostic at ASDEX Upgrade *Rev. Sci. Instrum.* **84** 113502
- [20] Geiger B *et al* 2014 Fast-ion transport in the presence of magnetic reconnection induced by sawtooth oscillations in ASDEX Upgrade *Nucl. Fusion* **54** 022005
- [21] Heidbrink W W, Van Zeeland M A, Grierson B A, Muscatello C M, Park J M, Petty C C, Prater R and Zhu Y B 2012 Initial measurements of the DIII-D off-axis neutral beams *Nucl. Fusion* **52** 094005
- [22] Pace D C *et al* 2013 Energetic ion transport by microturbulence is insignificant in tokamaks *Phys. Plasmas* **20** 056108
- [23] Bortolon A, Heidbrink W W and Podestà M 2010 A tangentially viewing fast ion D-alpha diagnostic for NSTX *Rev. Sci. Instrum.* **81** 10D728
- [24] Michael C A *et al* 2013 Dual view FIDA measurements on MAST *Plasma Phys. Control. Fusion* **55** 095007
- [25] Jones O M, Michael C A, McClements K G, Conway N J, Crowley B, Akers R J, Lake R J and Pinches S D 2013 Fast-ion deuterium alpha spectroscopic observations of the effects of fishbones in the Mega-Ampere Spherical Tokamak *Plasma Phys. Control. Fusion* **55** 085009
- [26] Ito T *et al* 2010 Fast ion charge exchange spectroscopy adapted for tangential viewing geometry in LHD *Rev. Sci. Instrum.* **81** 1–5

- [27] Press W H, Teukolsky S A, Vetterling W T and Flannery B P 2007 *Numerical Recipes: the Art of Scientific Computing* (Cambridge: Cambridge University Press)
- [28] Strang G 1988 *Linear Algebra and its Applications* 3rd edn (London: Thomson Learning)
- [29] Hansen P C, Jensen T K and Rodriguez G 2007 An adaptive pruning algorithm for the discrete L-curve criterion *J. Comput. Appl. Math.* **198** 483–92
- [30] Hansen P C 1998 *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion* (Philadelphia: SIAM)
- [31] Gull S F and Skilling J 1984 Maximum entropy method in image processing *Communications, Radar and Signal Processing, IEE Proceedings F* **131** 646–59
- [32] Giffin A and Caticha A 2007 Updating probabilities with data, moments *Bayesian Inference, Maximum Entropy Methods in Science and Engineering* vol **954** pp 74–84
- [33] Wächter A and Biegler L T 2006 On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming *Math. Program.* **106** 25–57
- [34] Lubin M and Dunning I 2015 Computing in operations research using julia *INFORMS J. Comput.* **27** 238–48
- [35] Pankin A, McCune D, Andre R, Bateman G and Kritiz A 2004 The tokamak Monte Carlo fast ion module NUBEAM in the national transport code collaboration library *Comput. Phys. Commun.* **159** 157–84
- [36] Nielsen S K *et al* 2011 Dynamics of fast ions during sawtooth oscillations in the TEXTOR tokamak measured by collective Thomson scattering *Nucl. Fusion* **51** 063014
- [37] Muscatello C M, Heidbrink W W, Kolesnichenko Y I, Lutsenko V V, Van Zeeland M A and Yakovenko Y V 2012 Velocity-space studies of fast-ion transport at a sawtooth crash in neutral-beam heated plasmas *Plasma Phys. Control. Fusion* **54** 025006