EUROFUSION WPMST1-PR(16) 15937

M Salewski et al.

High-definition velocity-space tomography of fast-ion dynamics

Preprint of Paper to be submitted for publication in Nuclear Fusion

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.

21 April 2016

Fast-ion dynamics studied by photon-starved velocity-space tomography

**M Salewski[1], B Geiger[2], AS Jacobsen[2], WW Heidbrink[3]
SB Korsholm[1], F Leipold[1], J Madsen[1], D Moseev[4], SK Nielsen[1],
M Nocente[5,6], T Odstrčil[2], J Rasmussen[1], L Stagner[3],
M Stejner[1], M Weiland[2] and the ASDEX Upgrade team[2]**

[1] Technical University of Denmark, Department of Physics, Kgs. Lyngby, Denmark
[2] Max-Planck-Institut für Plasmaphysik, Garching, Germany
[3] University of California Irvine, Irvine, CA 92697, USA
[4] Max-Planck-Institut für Plasmaphysik, Greifswald, Germany
[5] University of Milano Bicocca, Department of Physics, Milano 20126, Italy
[6] Istituto di Fisica del Plasma, Consiglio Nazionale delle Ricerche, Milano 20125, Italy

E-mail: msal@fysik.dtu.dk

Abstract. Velocity-space tomography is notoriously a photon-starved method due to limited optical access and signal-to-noise ratio of fast-ion D_α (FIDA) spectroscopy and the strive for high-resolution images. Here we make up for this lack of data by using prior information: The target velocity space is restricted using null measurements, phase-space densities are imposed to be non-negative, and numeric simulation are optionally used as prior information. The latter approach allows the inference of the velocity-space distribution of any discrepancies between measurements and simulation even for only two simultaneous FIDA views. Using the new methods, we study velocity-space dynamics of fast ions in sawtoothing plasmas using high-resolution tomographic movies and show the formation of neutral beam injection peaks at full, half and one-third energy.

1. Introduction

Ion distribution functions in magnetized fusion plasmas are often described by a population of fast ions from fusion reactions or auxiliary heating and a population of thermal ions. The latter is by definition described by a Maxwellian distribution and thus fully determined by a temperature T_i, a density n_i and a drift velocity $\mathbf{v}_i$. As these parameters are approximately constant on a flux surface, bulk-ion measurements often refer to flux surface measurements of the fundamental parameters T_i, n_i and $\mathbf{v}_i$ [1–4]. Fast-ion velocity distribution functions are much more complex. They are not flux functions due to the large drift excursions of fast ions, and they are in general 6D functions in phase space. Nevertheless, the strong magnetization in fusion plasmas allows us to reduce the velocity space to 2D due to rotational

symmetry about magnetic field vector. Usually fast-ion measurements do not refer to measurements of the fundamental high-dimensional fast-ion distribution functions but rather to measurements of derived quantities such as the measured 1D spectra in fast-ion D_α spectroscopy (FIDA) [5–7], collective Thomson scattering (CTS) [8–11], neutron emission spectrometry (NES) [12–14] or γ-ray spectrometry (GRS) [15–18]. This is in contrast to bulk-ion measurements where the fundamental parameters T_i, n_i and $\mathbf{v}_i$ are inferred from the spectra and presented as measurements whereas the spectra themselves are of secondary interest.

Velocity-space tomography seeks to make up for this shortcoming and allows us to infer 2D fast-ion velocity distribution functions, which are of fundamental interest, from measured spectra, which are of secondary interest [19–30]. The 2D velocity distribution function is localized in a small measurement volume which we take to be one point in position space. The fast-ion measurements depend on phase space in a complicated way illustrated by so-called weight functions which have been formulated for FIDA [6, 31], neutral particle analyzers (NPA) [6], CTS [20], fast-ion loss detectors [32], NES [33, 34] and GRS [35, 36]. To exploit the rich information about fast ions contained in the measurements, we need to consider hundreds of data points, e.g. spectral bins, together with the corresponding weight functions. Further, the measurements also depend on nuisance parameters such as bulk-ion densities or temperatures.

Velocity-space tomography provides a way to process this wealth of information at once. It provides a 2D image that is straightforward to interpret, that is the best useful fit to hundreds of simultaneous measurements by any available diagnostic, that shows the fundamental quantity of interest rather than derived quantities and that accounts for nuisance parameters. We have previously shown that this approach reveals, for example, velocity-space redistribution patterns of sawtooth crashes [25–30]. The tomography approach also helps comparisons of fast-ion measurements and numeric simulations. Traditionally this is done by comparing the measurements with simulated measurements in units particular to the diagnostic, e.g. the spectral density of the measured neutron or photon fluxes (D_α, γ, or mm-wave). Velocity-space tomography allows us to use 2D velocity space as an alternative and tangible meeting ground between theory and observation which is the same for all diagnostics. This is highlighted by the recent demonstrations of combined inversion of FIDA and CTS measurements [22, 28, 30].

Until now velocity-space tomography has relied on the standard inversion methods singular value decomposition (SVD), the maximum entropy method and variants of the Tikhonov regularization [23, 26, 27]. Given enough measurements at high signal-to-noise ratio, these inversion methods work well. However, velocity-space tomography in fusion plasmas is notoriously photon-starved since the optical access to tokamak plasmas is limited and the signal-to-noise ratio is often low compared with many other tomography applications. In conflict with the limited amount of data, we strive for high-resolution images requiring the inference of many unknowns. Inversions based on standard methods have therefore been plagued by artifacts, for example phase-space densities at energies larger than the injection energy in neutral beam heated plasmas

or negative phase-space densities [23]. Artifacts can be attenuated by installing further fast-ion diagnostics [21, 26] but often this is difficult due to economical or physical constraints. Here we make up for the lack of data in velocity-space tomography by using various types of prior information as additional constraints on the velocity distribution function. This not only improves results for the well-equipped five-view FIDA diagnostic at the tokamak ASDEX Upgrade (soon six-view), but it also allows the use of inversion techniques for systems with as few as two views as is common for other tokamaks. The attenuation of artifacts due to use of prior information further allows an increased resolution of the images.

We discuss prior information in velocity-space tomography in section 2 and the choice of the regularization strength in section 3. Substantial improvements brought about by prior information are shown in sections 4 and 5. In section 6 we study fast-ion velocity distributions in neutral beam heated plasma and dynamics in sawtoothing plasma using high-resolution movies. Sections 7 and 8 present a discussion of possible applications and conclusions.

2. Prior information for photon-starved velocity-space tomography

Velocity-space tomography entails the solution of an ill-posed problem in which we seek F^* solving the matrix equation

$$WF^* = S \tag{1}$$

where W and S are known [21]. F^* is the fast-ion velocity distribution function discretized in n pixels, S holds m fast-ion measurements, and W is an $m \times n$ matrix composed of weight functions. S and W are normalized by the uncertainties of the measurements [22]. Noise in the measurements makes the rows of the matrix equation 1 inconsistent, so that there is no solution irrespective of the choice of n ($n = m$, $n < m$, or $n > m$). One might have hoped that one could instead solve the related minimization problem

$$F^* = \arg\min_F \|WF - S\|_2, \tag{2}$$

but this is also useless as the matrix W is ill-conditioned and hence the solution F^* is not stable. This means that small perturbations in S can lead to large perturbations in F^* which is therefore dominated by random jitter. Nevertheless, we can construct a related well-conditioned problem by imposing additional requirements that reflect prior assumptions about the solution providing useful and stable solutions. This is called regularization. A popular regularization method in plasma physics is the Tikhonov regularization [27, 37, 38] in which we solve the minimization problem

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2. \tag{3}$$

In equation 3, F^* minimizes the sum of the residual of the original ill-posed problem (upper row) and the norm of the additional requirement on the solution (lower row).

Common choices are to require that the norm of F^* or the norm of the gradients in F^* be small. The penalty matrix L is then respectively the identity matrix or a matrix operator effecting a finite difference approximation of a gradient. These two choices are respectively motivated by the prior assumptions that solutions have moderate amplitudes and that they are smooth. The free regularization parameter λ balances the relative importance of the residual of the original problem and the norm of the regularization constraint. This reflects the balance of how well the solution should fit the noisy data and how much much it should obey the regularization constraint. Equation 3 shows that for small λ's the residual of the original problem dominates whereas for large λ's the norm of the regularization constraint dominates. The problem now is to select λ leading to a useful and stable solution. In the following we discuss how various types of prior information can be used to further improve inversions for moderately photon-starved systems with five FIDA views and even make the use of inversion techniques possible for highly photon-starved systems with two FIDA views. Our exposition will focus on Tikhonov regularization, but our methods are applicable to other regularization methods. We use FIDA measurements at ASDEX Upgrade as an example and will discuss applications to other diagnostics and tokamaks and stellarators in section 7.

2.1. Null measurements

FIDA measures Doppler-shifted D_α-light emitted when a fast deuterium neutral is formed in a charge-exchange reaction from a fast deuterium ion and then decays from the third to the second excited state [5]. In many experiments, there is a measurable upper limit to the observed Doppler shifts. Parts of a spectrum where no FIDA light is observed are referred to as null measurements. The wavelength ranges of null measurements are related to velocity-space regions through weight functions [31]. We refer to such weight functions as null-measurement weight functions. A null measurement suggests that the velocity space covered by the null-measurement weight function contains so few ions that they cannot be detected against the noise floor. (If we had a noise-free FIDA null measurement, we could conclude that the phase-space volume covered by the corresponding null-measurement weight function w_0 does not contain any ions at all.) If treated on an equal footing with detections of FIDA light, null measurements already contribute strongly to the reconstruction of the large-scale shape of the velocity distribution function since they tend to decrease the reconstructed phase-space densities in the velocity space covered by the null-measurement weight functions. For that reason the FIDA system at ASDEX Upgrade has been upgraded to measure red- and blue-shifted light in all spectra so that the absence of FIDA light could be measured [26]. Still, inversions are plagued by artifacts in velocity-space regions covered by null-measurement weight functions where the phase-space densities should be negligible [23].

Here we here remedy such artifacts by analyzing the measurements in two stages. In the first stage we identify regions in velocity space where null measurements suggest

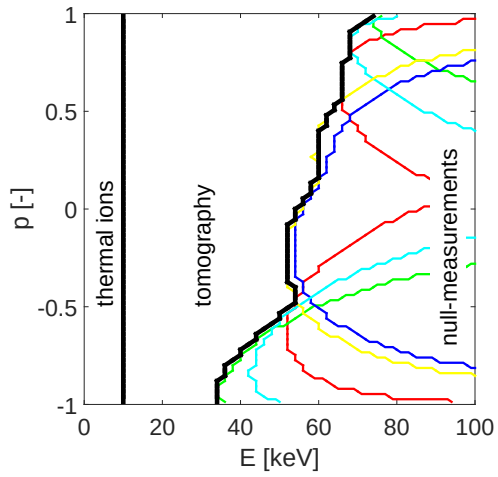

Figure 1. The coloured lines circumscribe the regions covered by null-measurement FIDA weight functions for ASDEX Upgrade discharge #31557. These null-measurements suggest that no fast ions exist with energies above the black line at 30-70 keV. The vertical black line at 10 keV is defined as the border between fast and thermal ions. The phase-space densities between the black lines are found by tomographic inversion.

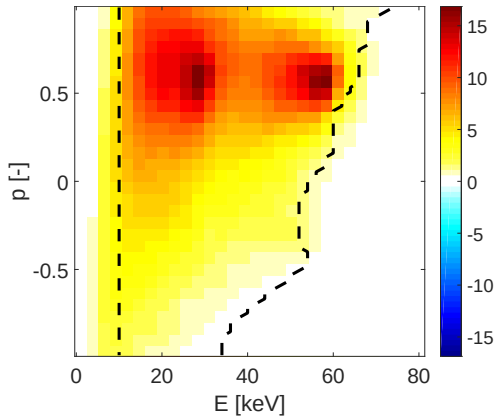

Figure 2. Comparison of a TRANSP simulation of the fast-ion velocity distribution function $[10^{10}\mathrm{keV}^{-1}\mathrm{m}^{-3}]$ and the target velocity space identified in figure 1. The simulation agrees with the FIDA measurements that the phase-space densities to the right of the dashed black line at 30-70 keV are negligible.

phase-space densities below the detection limit of the diagnostic according to

$$\int w_0 F_{true} dE dp = S < \epsilon \tag{4}$$

where w_0 is the null-measurement weight function, F_{true} is the unknown true distribution function, S is the measured signal, and ϵ is the noise floor. As customary, w_0 and F_{true} are given in energy-pitch coordinates (E, p), where E is the energy of the fast ions and p is the pitch defined as $p = \frac{v_\parallel}{v}$. $v_\parallel$ is the fast-ion velocity antiparallel with respect to the magnetic field and v is the fast-ion speed. The non-negativity of w_0 and F_{true} allows us to conclude from a certain null-measurement that F_{true} must be zero in the

regions covered by the null-measurement weight function. If noise is present, the strict argument does not hold but we can still assert that any phase-space densities in these regions are below the detection limit of the diagnostic and therefore negligible. This stage does not require the solution of an inverse problem and is thus very reliable if null-measurements can be told apart from measurements of small FIDA intensities. In the second stage we solve the tomography problem neglecting the phase-space densities in the regions covered by null-measurement weight functions:

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2 \quad \text{subject to} \quad F^*(E_0, p_0) = 0. \tag{5}$$

E_0 and p_0 are the energies and pitches covered by the null-measurement weight functions. This condition determines an upper boundary of the velocity space that we seek to reconstruct as illustrated in figure 1. As weight functions are not bounded in energy, the target velocity space has also been judiciously restricted at some energy in previous work. A common choice is 10-50 keV above the injection energy of 60 keV or 90 keV as phase-space densities at larger energies are supposed to be small. Null-measurement weight functions allow us to determine this upper boundary in energies from the measurements in an optimal way that does not allow artifacts in the null-measurement region and that minimizes the number of pixels in the inversion. This upper boundary is a strong function of pitch (figure 1). Neoclassical simulations by TRANSP [39] agree very well with the shape of the null-measurement velocity space as illustrated in figure 2.

For FIDA measurements the null-measurement approach is most applicable to experiments with neutral beam injection. In experiments with ion cyclotron resonance heating (ICRH) such FIDA null-measurements may not be observable. The null-measurement idea requires that presence and absence of FIDA signal can be defined based on the measurements. This is often the case if only neutral beam heating is used. Nevertheless, we note that null measurements are often observable in NES and GRS measurements in plasmas with third harmonic ICRH at JET as these diagnostics can easily detect MeV-range ions. Impurity radiation masking part of the null measurements is usually not a nuisance since the velocity space covered by the lost null measurements is often covered by other null measurements.

2.2. Non-negativity

One of the advantages of maximum entropy regularization is that it does not allow negative phase-space densities. The SVD and variants of the Tikhonov regularization (one of which inspired by Fisher information) do allow negative phase-space densities and in fact inversions often contained regions with small negative phase-space densities. Negative phase-space densities have usually been ignored [7, 21–24, 27–30]. In a recent study negative phase-space have been strongly attenuated by additional constraints implemented as fictitious measurements corresponding to pixel-sized weight functions covering the negative patches in the inversions [26]. These fictitious measurements were given enough weight to iteratively force the phase-space densities to negligible, yet still

negative, values. Here we simply impose the constraint that the solution be non-negative and solve the minimization problem

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2 \quad \text{subject to} \quad F^* \geq 0. \tag{6}$$

The implementation as constraint has the advantage that no assumptions about the negative regions need to be made. In fact, regions with negative amplitudes in unconstrained inversions can have positive amplitudes in inversions that are constrained to be non-negative. We will in the following also impose non-negativity and null-measurement constraints together:

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ 0 \end{pmatrix} \right\|_2 \quad \text{subject to} \quad \begin{matrix} F^*(E_0, p_0) = 0 \\ F^* \geq 0 \end{matrix} . \tag{7}$$

2.3. Prior information from a numeric simulation

In previous studies of velocity-space tomography based on three to five FIDA views or a mix of FIDA and CTS views, it was possible to find useful and stable inversions and hence provide measurements of the fast-ion velocity distribution function and derived quantities such as the fast-ion density [25, 27]. However, if fewer views are available or if the signal-to-noise ratio is unfavourable, this is sometimes out of reach. Here we propose a new goal of the tomography approach for such highly photon-starved situations. Rather than inferring the full 2D velocity distribution function, we only seek to infer the source of any discrepancy between a numeric simulation and experiment in velocity space. Firstly, this is a simpler task that is possible with fewer views as we will show. Secondly, deviations from a simulation is often what we are actually interested in. For example, the TRANSP code accounts for neoclassical transport and likely provides a good picture of the distribution function in MHD quiescent discharges. In MHD active discharges, we often observe discrepancies from such neoclasssical simulations. It is thought that a so-called anomalous transport exists on top of the neoclassical transport which still serves a baseline. In this case we can use the tomography approach to locate the source of any disagreement to this neoclassical simulation in velocity space rather than to measure the full 2D fast-ion velocity distribution function. To do this we penalize differences to the simulated velocity distribution function F_{sim} in absolute values or gradients and solve the minimization problem

$$F^* = \arg\min_F \left\| \begin{pmatrix} W \\ \lambda L \end{pmatrix} F - \begin{pmatrix} S \\ \lambda L F_{sim} \end{pmatrix} \right\|_2 \quad \text{subject to} \quad \begin{matrix} F^* \geq 0 \\ F^*(E_0, p_0) = 0 \end{matrix} . \tag{8}$$

For very small λ's the inversion is underegularized and is dominated by random jitter as in equations 3, 5 and 7. For very large λ's the solution is overregularized, but here the inversion then approaches the numeric simulation as equation 8 shows. The problem is now again to select λ assigning an appropriate balance between the measurements and the simulation which we will discuss in section 3. The solution F^* represents our estimate of the 2D velocity distribution function considering the simulation and

the measurements. As our goal is to locate discrepancies between simulation and measurements in velocity space, we calculate

$$\Delta F^* = F^* - F_{sim} \tag{9}$$

which shows the modifications to the simulation suggested by the measurements. Examples will be shown in section 5.

3. Choice of the regularization parameter

The number of measurement data points and the number unknowns in velocity-space tomography is fairly small compared with many other tomography problems. Quick automatic inversion of FIDA spectra based on a library of approximate weight functions after each plasma discharge is certainly possible and will be implemented in future work. For that purpose the regularization parameter λ must be computed automatically from the data. In traditional tomography, two popular methods to do this among many other alternatives are the L-curve method [40, 41] and the generalized cross-validation (GCV) method [Wahba 1990]. The L-curve method has been applied to velocity-space tomography previously [27]. The GCV asserts that an arbitrary measurement S_i should be predicted well by the regularized solution based on the other measurements in S and that orthogonal transformations of S should not affect the choice of λ [Wahba 1990]. Then λ minimizes the GCV function

$$GCV = \frac{\|WF - S\|_2^2}{(\mathrm{trace}(I - WW^\dagger))^2} \tag{10}$$

where I is the identity matrix and $W^\dagger$ is the regularized inverse matrix producing the regularized solution $F = W^\dagger S$.

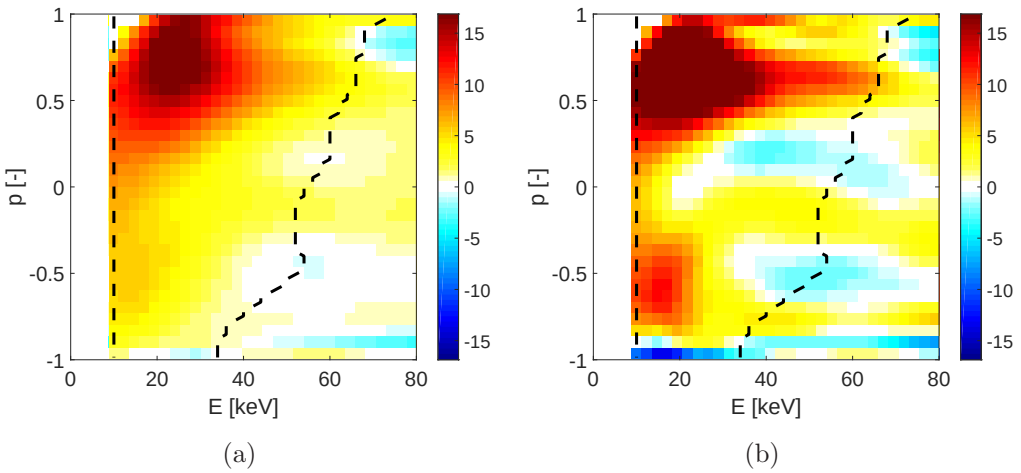

Figure 3. Choice of regularization parameter by two different methods applied to unconstrained first-order Tikhonov regularization (a) L-curve. (b) GCV.

Figure 3 shows two unconstrained first-order Tikhonov inversions of FIDA measurements in discharge #31557. The level of regularization was respectively selected

by the L-curve method and the GCV method. The discharge has been discussed previously [27]. The plasma was heated by neutral beam injection with the 2.5 MW source S3 at 60 keV. Figure 2 shows a TRANSP simulation of the velocity-distribution function in this discharge. Maxima appear at 60 keV and between 20 keV and 30 keV, corresponding to the full injection energy and between half and one-third injection energy. The velocity-space resolution of the inversion is not high enough to resolve the individual peaks at 20 keV and 30 keV.

The unconstrained first-order Tikhonov inversion with L-curve regularization parameter selection in figure 3(a) shows usual characteristics for this method. The inversion is smooth. The overall shape including the anisotropy and the location of the merged 20 keV and 30 keV beam injection peaks is close to our expectation. However, there are also limitations of the unconstrained inversion. First, the peak at full injection energy (60 keV) does not appear. Inversions of synthetic measurements have shown that sharp peaks tend to be attenuated, especially for first-order Tikhonov regularization. Second, there are substantial phase-space densities at energies well in the null-measurement region. Here we plot up to 20 keV above the full injection energy. As already mentioned, neoclassical transport theory as computed by TRANSP/NUBEAM also predicts almost zero phase-space densities in the null-measurement region (figure 2). This further corroborates that the inferred phase-space densities are artifacts. These unexpected features of the inversions are very likely artifacts since they also appear erroneously in inversions of synthetic data based on the TRANSP simulation [21,23,27]. Third, small patches of negative fast-ion phase-space densities appear, in this case in the null-measurement region. Negative phase-space densities are unphysical artifacts.

The L-curve and GCV methods select different balances between data fitting and smoothing. The GCV method tends to regularize less than the L-curve method and roughly produces a regularization level as was sometimes judiciously selected [7,23,26]. The full energy beam injection peak (60 keV) does not appear but there is a clear ridge of large phase-space densities between the expected locations of the beam injection peaks. The ridge does not extend to energies larger than 60 keV as expected. The form of the ridge is consistent with presence of fast ions that are slowing down due to collisions with electrons, i.e. they lose energy without significant pitch angle scattering. This is reasonable as this ridge appears above the critical energy of about 40 keV. Below the critical energy collisions with thermal ions become important leading to pitch-angle scattering as apparent in figure 3b. However, the unphysical negative regions are larger than for the L-curve method. Both methods are fairly robust but do not always work. The L-curve method can be applied irrespective of the incorporated prior information. Figures 3a and 3c show first-order Tikhonov inversions with and without non-negativity constraint, respectively. However, the GCV method requires the existence of the inverse matrix, and so it is not directly applicable to the non-negative least square problem formulation or maximum entropy methods.

In velocity-space tomography, the regularization parameter has up to now been set by judicious choice, by the L-curve method, and here by the GCV method. The choice

of regularization parameter or the of a method to calculate it remains an open problem, which we leave for future work, as one never knows a priori which method works best. Neither the L-curve nor the GCV method is consistently superior, and hence we use both methods in the following and judiciously chose the regularization parameter based on both methods.

4. Inversions from measurements using prior information

In this section we show that prior information imposed as side constraints based on non-negativity and null measurements improves inversions substantially. The effect of prior information based on numeric simulation is presented in section 5. This separation reflects the difference in the approaches. The inversions in this section are purely based on measurements and represent an estimate of the 2D fast-ion velocity distribution functions based on measurements alone. The inversions in section 5 represent and estimate based on measurements and simulations or, more interestingly, how a given simulation needs to be modified to match best with the measurements.

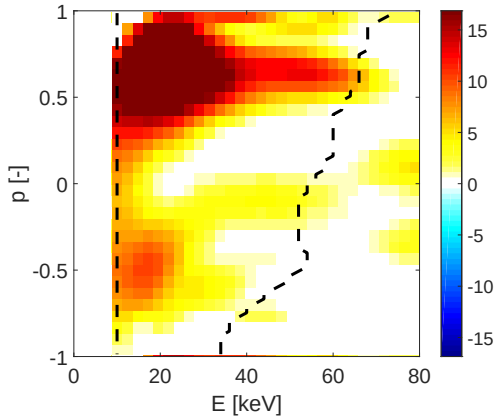

Figure 4. First-order Tikhonov inversion with non-negativity constraint.

Figure 4 shows a first-order Tikhonov inversion with non-negativity constraint according to equation 6. The regularization parameters selected by the GCV method. The GCV method cannot directly be applied together with the non-negativity constraint as the inverse matrix $W^{\dagger}$ does not exist. Here we use the regularization parameter λ from the unconstrained regularization (figure 3b), and then solve the non-negative least-square problem with that λ. The non-negativity constraint improves the inversions without any visible disadvantages compared with the unconstrained inversions in figure 3. By virtue of the constraint, there are no negative regions. Erroneous phase-space densities in the null-measurement region are decreased compared with the unconstrained inversions. The full energy beam injection peak appears.

Figure 5 shows a first-order Tikhonov inversion with the null-measurement constraint according to equation 5. The inversion shows a clear peak at full injection energy and a clear ridge consistent with ions slowing down due to collisions with

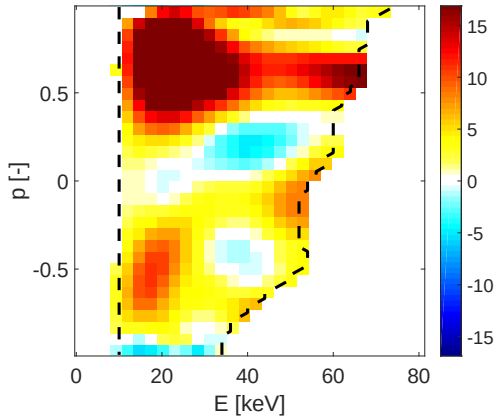

Figure 5. Tomographic inversions by first-order Tikhonov regularization using the null-measurement weight function constraint $F^*(E_0, p_0) = 0$. F^* can have either sign.

electrons. By virtue of the constraint, no fast ions are erroneously placed in the null-measurement region. Regions of erroneously negative phase-space densities are present at this regularization level. It is unclear if the two peaks appearing at negative pitches are real or if they are artifacts.

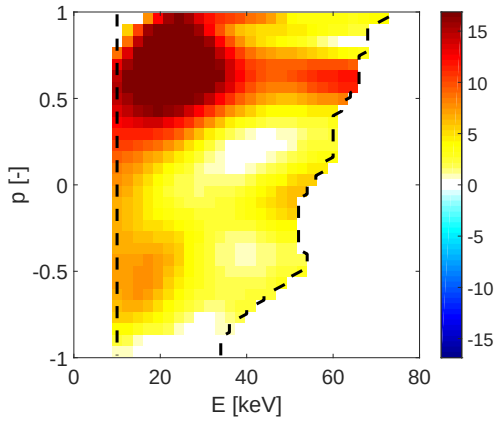

Figure 6. Tomographic inversions by first-order Tikhonov regularization using the null-measurement weight function constraint ($F^*(E_0, p_0) = 0$) and the non-negativity constraint ($F^* \geq 0$).

Finally, figure 6 shows the inversion with non-negativity and null-measurement constraints according to equation 7. By construction there are no artifacts in the null-measurement region and there are no negative phase-space densities. A peak at full injection energy appears and a ridge connects the beam injection peaks. There are small peaks at negative pitches at the same positions as in figure 5 but these peaks are strongly attenuated compared with figure 5. A comparison of figures 5 shows that the negative regions in figure 5 are not necessarily zero. The non-negativity constraint not only removes negative regions but also tends to attenuate peaks. The non-negativity constraint regularizes the inversion by attenuating high-frequency components of the inversion that enhance negative and positive extrema. Overall, figure 6 shows no clear

artifacts and additionally has a better capability to reconstruct the full injection energy peak and the ridge consisting of ions colliding with electrons.

5. Inversions using numerical simulations as prior information

A goal of inversion techniques is to locate differences between theory and observation in velocity space to obtain clues on which physics is not adequately described in the simulation or which systematic error confounds the measurements. For this goal it is advantageous to use the numeric simulation as prior information and reconstruct modifications of this simulation as implied by the measurements. Figure 7 illustrates two different ways to locate differences between measurements and simulation in velocity space. Figure 7(a) shows an assumed true distribution function F_{true} which we now seek to reconstruct. In an experiment F_{true} is never known. Here we use a modified TRANSP simulation. We have divided the phase-space density in the square shown in figure 7(b) by two. This selective reduction in phase-space density is a model for an assumed anomalous transport phenomenon localized in velocity space which is not modelled in the TRANSP simulation F_{sim}. Figure 7(b) shows the difference

$$\Delta F_{true} = F_{true} - F_{sim} \tag{11}$$

which is negative in the square and zero elsewhere. The goal is now to reconstruct ΔF_{true}, given synthetic measurements of F_{true} in five FIDA views and F_{sim}. We add 5% Gaussian noise to each synthetic measurement which is a realistic noise level for FIDA measurements at ASDEX Upgrade. ΔF^* is calculated according to equation 9 without and with using the TRANSP simulation as prior information, i.e. respectively using equations 7 and 8. Figure 7(c) shows ΔF^* for the case without the TRANSP simulation as prior. ΔF^* is dominated by large negative values at the beam injection peaks reflecting the difficulty to reconstruct the beam injection peaks. The reconstructed difference ΔF^* does not resemble the true difference ΔF_{true}. Figure 7(d) presents ΔF^* for the case where the TRANSP simulation has been used as prior information (equation 8). In this case the approximate location of the discrepancy between the true distribution and the TRANSP simulation is found based on experimentally accessible quantities. However, artifacts are present and the amplitudes in the reconstruction are too low by a factor 2-3. The success of this approach does not depend strongly on the position of the discrepancy in velocity space for the five-view FIDA system as we illustrate in figure 8. Each case shows the true location of the discrepancy and the reconstruction using the simulation as prior information. The approximate location of the discrepancy is in each case identified as the region with the lowest amplitudes.

In figure 9 we reconstruct the six cases from figure 8 using just two FIDA views. Even with only two views, the approximate location of the square is reconstructed in each case, even though the region of negative phase-space densities is significantly extended compared to the five-view case. The cases where the discrepancy is placed near $p = 1$ are very well reconstructed even for only two views. The reconstruction of ΔF^* does not

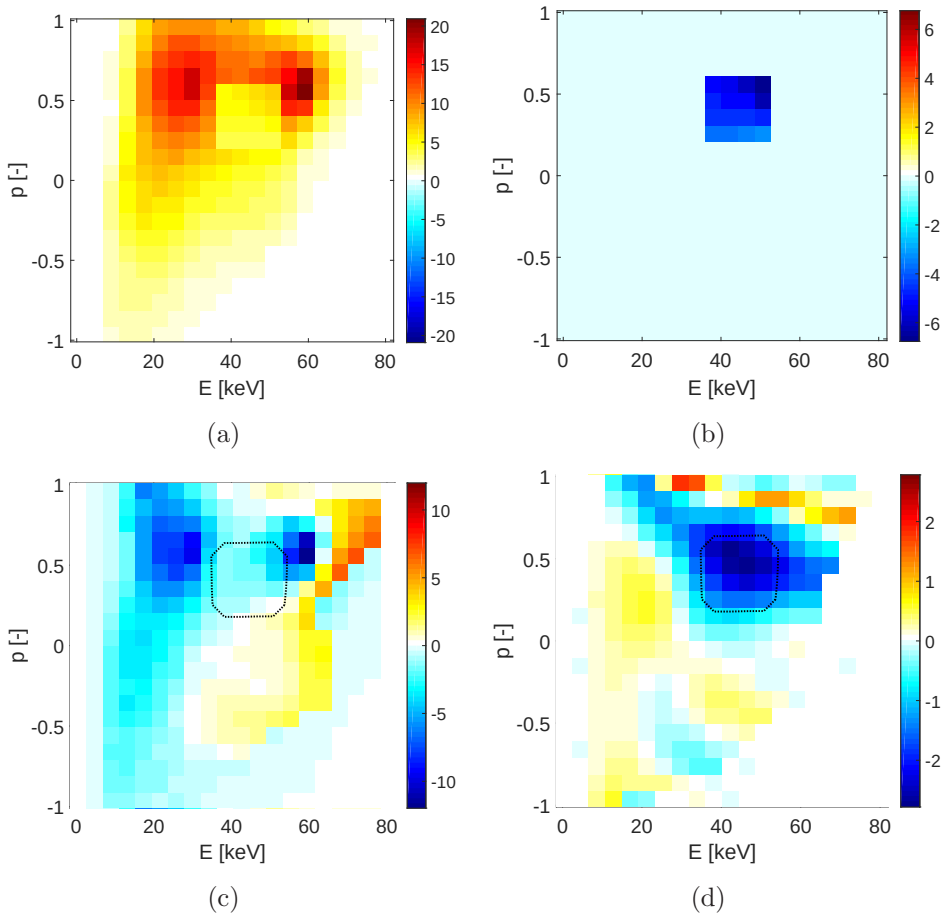

Figure 7. FIX MISSING BOXES! (a) F_{true} which is a modified TRANSP simulation with selective ejection of particles. (b) $\Delta F = F_{true} - F_{sim}$. (c)+(d) Inversion $\Delta F^* = F^* - F_{sim}$ based on synthetic measurements of F_{true} with 5% Gaussian noise for a five-view FIDA system not using (c) and using (d) the TRANSP simulation as prior information.

require as many measurements as the reconstruction of F^* since the simulation provides cogent prior information about the approximate basic shape. Hence this alternative approach to tomographic reconstruction should be highly useful for FIDA systems with fewer than five viewing directions as is common on many machines. DIII-D has three FIDA views [42, 43] and MAST [44, 45], NSTX [46], EAST [47], and LHD have two views [48, 49]. LHD has additionally a CTS view [50, 51].

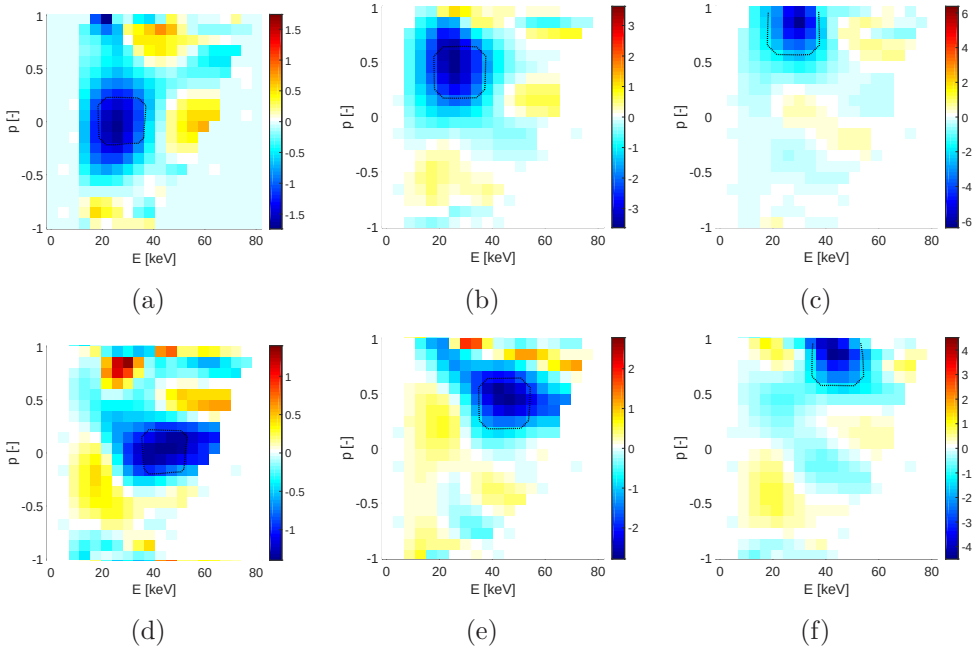

Figure 8. FIX MISSING BOXES! Reconstruction of $\Delta F^* = F^* - F_{sim}$ based on synthetic measurements of F_{true} with 5% Gaussian noise for a five-view FIDA system using the TRANSP simulation as prior information. In each subfigure the black line shows the true location of the discrepancy.

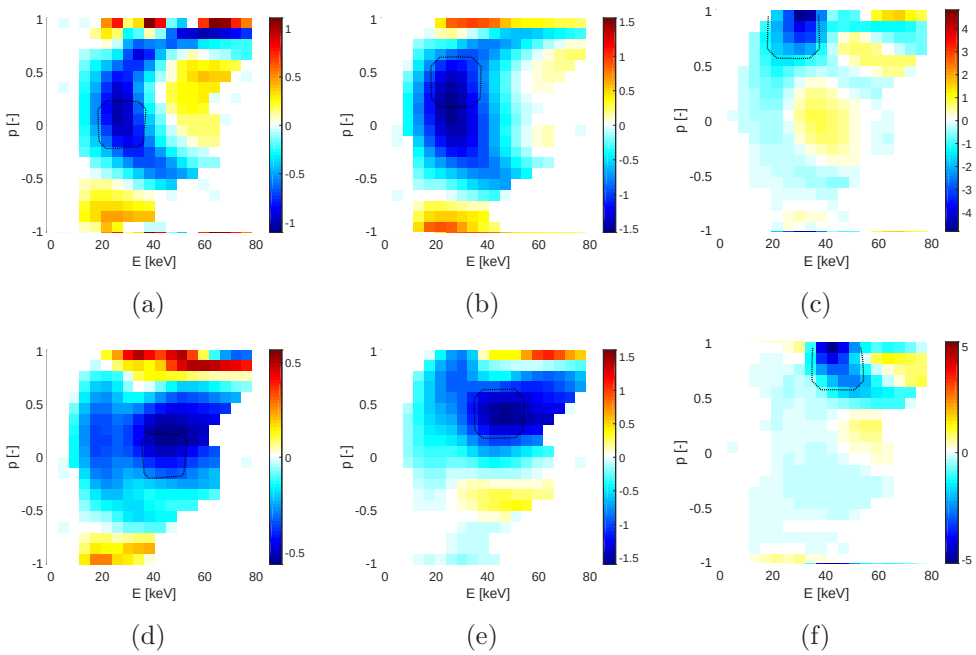

Figure 9. FIX MISSING BOXES! As for figure 8, but only two FIDA views are used.

6. Studies of beam injection and sawtooth dynamics

Prior information allows us to increase the resolution of the inversions (30×30 pixels) to study the presence of neutral beam injection peaks and the dynamics of sawteeth. Beam injection peaks are studied by SVD which is usually best to reconstruct fine-scale features at the expense of the appearance of some jitter. Figure 10 shows inversions using the SVD with and without null-measurement constraint for the reference discharge 31557. The SVD reconstructs the full energy beam injection peak and merges the half and one-third energy peaks as previously observed [23]. Negative phase-space densities appear as we could not enforce non-negativity in the SVD. As observed previously, the negative regions are mostly close to or in the null-measurement region, and the amplitudes are fairly small. In the inversion with the null-measurement constraint, the full energy injection peak is strongly enhanced and several patches with negative fast-ion densities disappear. The ridge of enhanced fast ion densities connecting the beam injection peaks as observed in the inversions with first-order Tikhonov and minimum Fisher regularization does not appear.

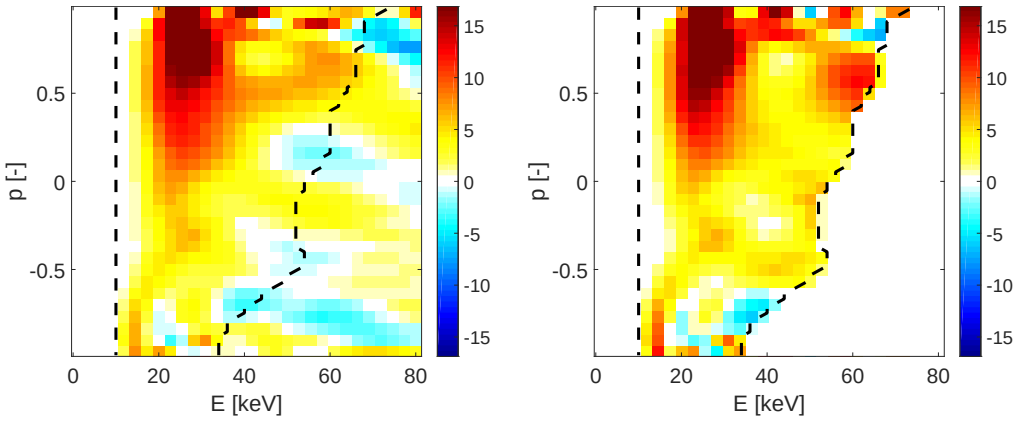

Figure 10. Tomographic inversions by SVD. (a) No constraint. (b) Null-measurement weight function constraint: $F^*(E_0, p_0) = 0$.

Figure 11 shows an inversion by SVD with null-measurement constraint in the similar discharge 32323. Here we increased the resolution to 60×60 pixels. In this case all three beam injection peaks appear at the correct energies at 20 keV, 30 keV and 60 keV and pitch. The extent in pitch direction is fairly small in the 30 keV and 60 keV peak and much larger in the 20 keV peak. This is reasonable as the 20 keV peak is below the critical energy so that ion-ion collisions lead to increased pitch angle scattering. In discharge 31557 discussed above only two peaks appear at this resolution. The first-order Tikhonov and minimum Fisher information inversions show only two peaks as well. Even though three peaks do not often appear in inversions, the correct and well-understood location and large-scale coherence suggests that the peaks are not artifacts but supported by the FIDA measurements.

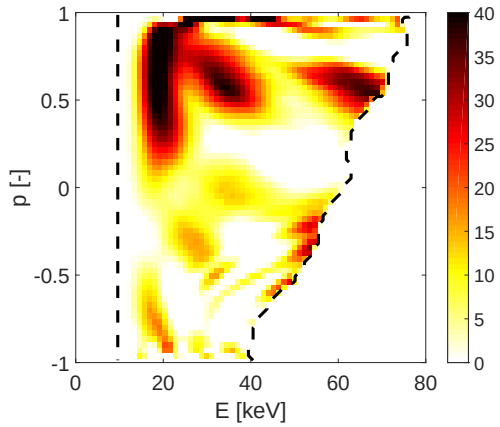

(a)

Figure 11. Tomographic inversion by SVD using the null-measurement weight function constraint: $F^*(E_0, p_0) = 0$.

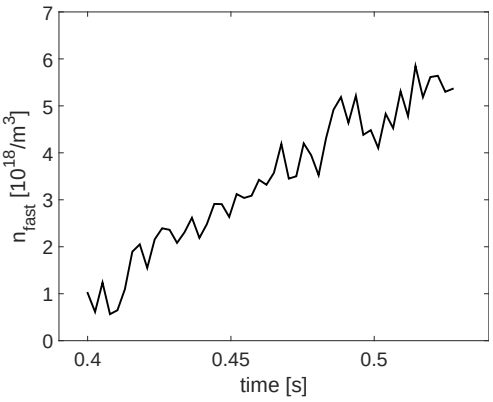

Figure 12. Increase in the fast-ion density in ASDEX Upgrade discharge #33138 right after the NBI is switched on. The fast ion density is obtained integration of the tomographic inversion in velocity space above energies of 10 keV.

In discharge 33138 FIDA measurements were made during the start-up phase of the NBI to study the appearance of the beam injection peaks. Figure 12 shows the increasing fast-ion density during the start-up phase which is obtained by integration of the velocity distribution functions. The appearance of the beam injection peaks is studied in a movie provided as supplemenary material. FIDA measurements in five views have been made with a time resolution of 2.5 ms. Figure 13 shows selected frames. The beam injection peaks at full, half, and one-third injection energiy reliably appear in the movie, and the growth can be studied.

Discharge 32323 had sawtooth activity. The improved robustness of the inversions due to the null-measurement and non-negativity constraints allows us to study the dynamics of the 2D fast-ion velocity distribution function in a sawtoothing plasma by a tomographic inversion movie. Previous inversions of FIDA and CTS measurements focussed on the redistribution due to the crash [25–30]. In this discharge only four views

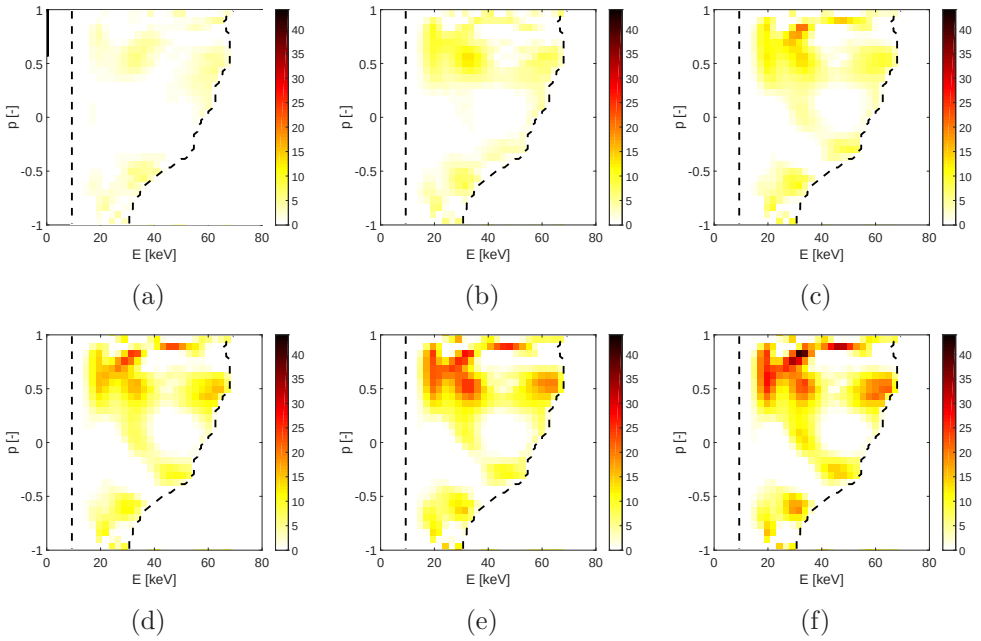

Figure 13. Frames from the NBI start-phase movie attached as supplementary material a selected times in ASDEX Upgrade discharge #33138.

had a time resolution of 2.5 ms whereas the fifth view had a time resolution of 25 ms as the camera did not allow faster data acquisition in this discharge, and hence we do not use this view for the movie. The same level of regularization is used throughout the movie to avoid sudden jumps. The movie is provided as supplementary material. The movie shows a continuous increase of the fast-ion density which then suddenly drops at the sawtooth crash. The movie shows how the distribution function responds to the continuous fuelling by the neutral beam injector which provides particle sources at 20 keV, 30 keV and 60 keV. Right after the crashes the 60 keV peak almost disappears. It gradually builds up as the fast-ion density increases during the sawtooth cycle. The ridge between the beam injection peaks also becomes stronger during the sawtooth cycle.

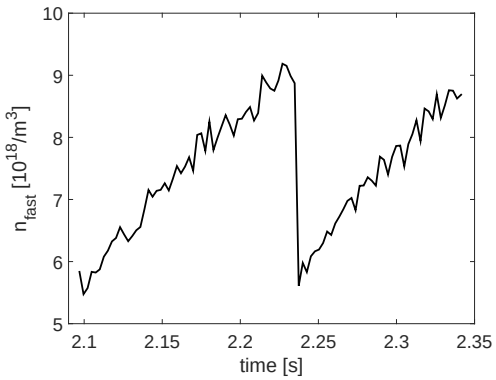

Figure 14. Sawteeth in the fast-ion density in ASDEX Upgrade discharge #32323. The fast ion density is obtained integration of the tomographic inversion in velocity space above energies of 10 keV.

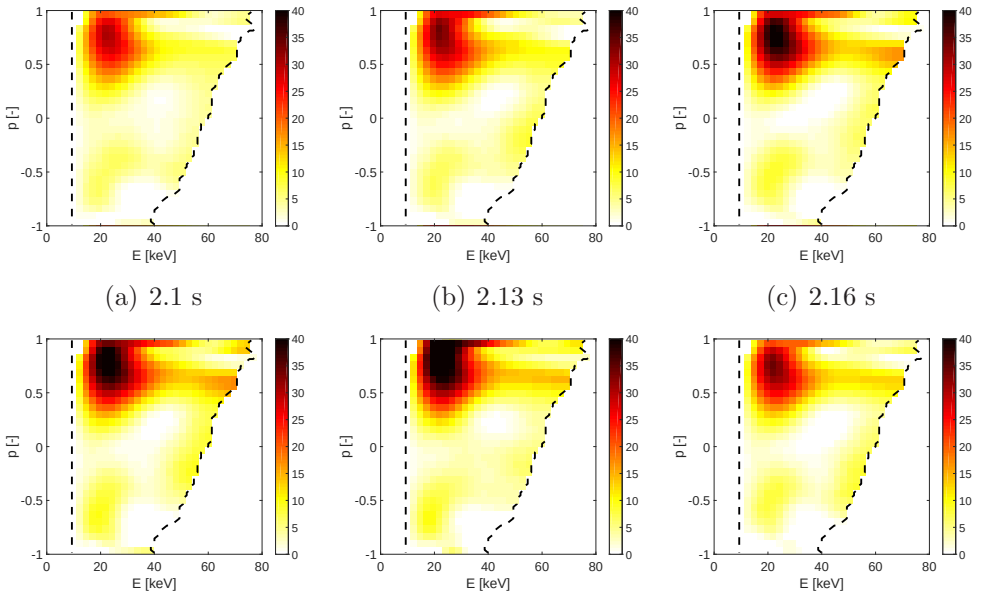

(a) 2.1 s (b) 2.13 s (c) 2.16 s

(d) 2.19 s (e) 2.22 s (f) 2.25 s

Figure 15. Frames from the sawtooth crash movie attached as supplementary material a selected times covering one sawtooth cycle in ASDEX Upgrade discharge #32323.

7. Discussion

To date all velocity-space tomography studies have been made at ASDEX Upgrade based on three to five simultaneously acquired FIDA spectra or four simultaneously acquired FIDA spectra and one CTS spectrum. The installed hardware at ASDEX Upgrade now allows simultaneous measurements with six FIDA views, two CTS views, one NPA view and one NES view though this multidiagnostic tomography still needs to be demonstrated. FIDA systems installed on other machines have three or fewer FIDA views. Examples are three FIDA views at DIII-D, two FIDA views at NSTX, two FIDA views at MAST, two FIDA views at EAST, and two FIDA views and a CTS view at LHD. No matter how rich a suite of fast ion diagnostics is available, the velocity-space tomography problem will remain photon-starved as increases in the number of measurement data will be matched by a corresponding increase in the resolution of the inversion. Efficient use of prior information is here shown to help extract substantially more information from the measurements. As a consequence tomographic inversion techniques could become useful at the other machines with just two or three FIDA or CTS views.

An important future task of the velocity-space tomography is to measure MeV-range fast-ion velocity distribution functions as we approach the era of burning reactor-grade plasmas. Unfortunately, FIDA spectroscopy is not thought to work well for measurements of MeV-range ions since the charge-exchange cross sections are low at such energies, the neutral beam penetration is often poor due to high densities of reactor-grade plasmas, and intensive bremsstrahlung is expected. Nevertheless, MeV-

range ions can be diagnosed by NES, GRS or CTS. Velocity-space tomography using different diagnostics [22] has already been demonstrated using FIDA and CTS as example [28]. A near-term goal is now to demonstrate tomographic inversion of GRS and NES measurements at JET, so as to allow velocity-space tomography in the upcoming deuterium-tritium (DT) campaign. Several high resolution detectors based on different measurement principles are installed at JET, and further upgrades are foreseen for the DT campaign. Simultaneous high-resolution NES and GRS measurements along two sightlines have recently been demonstrated [14]. As NES and GRS weight functions have recently been formulated [33–36], measurements of MeV-range ion velocity distribution functions based on NES and GRS can now be attempted. The effective use of prior information presented here will pave the way to make velocity-space tomography available to study MeV-range ions in burning plasmas.

8. Conclusions

Velocity-space tomography is notoriously a highly photon-starved enterprise as we would like to infer the 2D velocity distribution function in high resolution and hence need to infer as many unknown parameters as the measurement data supports. At the same time the optical access to tokamaks and hence the number of simultaneous measurements are limited. Furthermore, the signal-to-noise ratio of the measurements is lower than in many other tomography applications. Even if we use four to five FIDA views installed at ASDEX Upgrade, or any upgrade beyond, the velocity-space tomography problem will remain photon-starved. Here we make up for the lack of measurements by using additional prior information. The inversions are substantially improved by using two constraints: 1. The non-negativity of phase-space densities. 2. The negligible phase-space densities in velocity-space regions that do not emit detectable FIDA radiation leading to null measurements. Due to the new methods the three neutral beam injection peaks at full, half and one-third energy appeared at the expected locations in the inversion, and their appearance during the start-up phase of a neutral beam injector could be studied. We could further study the dynamics of the fast-ion velocity distribution function in a sawtoothing plasma in unprecedented detail in a tomographic inversion movie that efficiently summarizes about 50.000 data points.

Lastly, we propose an alternative approach to velocity-space tomography. We infer the 2D fast-ion distribution function considering the measurements as well as a simulation. In this case tomographic inversion uses the simulation as prior information. If the measurements and the simulations are inconsistent, the most likely velocity-space distribution of the discrepancies can then be found by subtracting the simulation from the inversion with the simulation as prior information. This could not be achieved with the pure velocity-space tomography approach, which does not use the simulation, even with five FIDA views. This alternative approach is also promising in severely photon-starved cases as we demonstrated using only two FIDA views and opens up for the use of velocity-space tomography methods at many machines.

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.

References

[1] Grierson B A *et al* 2012 *Physics of Plasmas* **19** 056107
[2] Grierson B A *et al* 2012 *Rev. Sci. Instrum.* **83** 10D529
[3] Stejner M *et al* 2013 *Plasma Phys. Control. Fusion* **55** 085002
[4] Stejner M *et al* 2015 *Plasma Phys. Control. Fusion* **57** 062001
[5] Heidbrink W W *et al* 2004 *Plasma Phys. Control. Fusion* **46** 1855–1875
[6] Heidbrink W W *et al* 2007 *Plasma Phys. Control. Fusion* **49** 1457–1475
[7] Geiger B *et al* 2015 *Plasma Phys. Control. Fusion* **57** 014018
[8] Salewski M *et al* 2010 *Nucl. Fusion* **50** 035012
[9] Moseev D *et al* 2011 *Plasma Phys. Control. Fusion* **53** 105004
[10] Nielsen S K *et al* 2015 *Plasma Phys. Control. Fusion* **57** 035009
[11] Rasmussen J *et al* 2015 *Plasma Phys. Control. Fusion* **57** 075014
[12] Hellesen C *et al* 2010 *Nucl. Fusion* **50** 084006
[13] Hellesen C *et al* 2013 *Nucl. Fusion* **53** 113009
[14] Eriksson J *et al* 2015 *Nucl. Fusion* **55** 123026
[15] Kiptily V G *et al* 2010 *Nucl. Fusion* **50** 084001
[16] Tardocchi M *et al* 2011 *Physical Review Letters* **107** 205002
[17] Nocente M *et al* 2012 *Nucl. Fusion* **52** 063009
[18] Nocente M *et al* 2015 *Nucl. Fusion* **55** 123009
[19] Egedal J and Bindslev H 2004 *Physics of Plasmas* **11** 2191
[20] Salewski M *et al* 2011 *Nucl. Fusion* **51** 083014
[21] Salewski M *et al* 2012 *Nucl. Fusion* **52** 103008
[22] Salewski M *et al* 2013 *Nucl. Fusion* **53** 063019
[23] Salewski M *et al* 2014 *Nucl. Fusion* **54** 023005
[24] Salewski M *et al* 2015 *Plasma Phys. Control. Fusion* **57** 014021
[25] Geiger B *et al* 2015 *Nucl. Fusion* **55** 083001
[26] Weiland M *et al* 2016 *Plasma Phys. Control. Fusion* **58** 025012
[27] Jacobsen A S *et al* 2016 *Plasma Phys. Control. Fusion* **58** 045016
[28] Jacobsen A S *et al* 2016 *Plasma Phys. Control. Fusion* **58** 042002
[29] Jaulmes F *et al* 2016 *Nucl. Fusion* **56** at press
[30] Rasmussen J *et al* 2016 *Nucl. Fusion* **56** at press
[31] Salewski M *et al* 2014 *Plasma Phys. Control. Fusion* **56** 105005
[32] Pace D C *et al* 2012 *Rev. Sci. Instrum.* **83** 073501
[33] Jacobsen A S *et al* 2014 *Rev. Sci. Instrum.* **85** 11E103
[34] Jacobsen A S *et al* 2015 *Nucl. Fusion* **55** 053013
[35] Salewski M *et al* 2015 *Nucl. Fusion* **55** 093029
[36] Salewski M *et al* 2016 *Nucl. Fusion* **56** 046009
[37] Anton M *et al* 1996 *Plasma Phys. Control. Fusion* **38** 1849–1878
[38] Craciunescu T *et al* 2009 *Nucl. Instrum. Meth. Phys. Res. A* **605** 374–383
[39] Pankin A *et al* 2004 *Comp. Phys. Comm.* **159** 157–184
[40] Hansen P C 1992 *SIAM Review* **34** 561–580
[41] Hansen P C and OLeary D P 1993 *SIAM J. Sci. Comp.* **14** 1487–1503

[42] Heidbrink W W *et al* 2012 *Nucl. Fusion* **52** 094005
[43] Pace D C *et al* 2011 *Plasma Phys. Control. Fusion* **53** 062001
[44] Michael C A *et al* 2013 *Plasma Phys. Control. Fusion* **55** 095007
[45] Jones O M *et al* 2014 *arXiv:1401.6864*
[46] Bortolon A *et al* 2010 *Rev. Sci. Instrum.* **81** 10D728 ISSN 1089-7623
[47] Huang J *et al* 2014 *Rev. Sci. Instrum.* **85** 11E407
[48] Ito T *et al* 2010 *Rev. Sci. Instrum.* **81** 10D327
[49] Ito T *et al* 2012 *Plasma and Fusion Research* **5** S2099
[50] Kubo S *et al* 2010 *Rev. Sci. Instrum.* **81** 10D535
[51] Nishiura M *et al* 2014 *Nucl. Fusion* **54** 023006