SUMMARY REPORT OF THE 4TH IAEA TECHNICAL MEETING ON FUSION DATA PROCESSING, VALIDATION AND ANALYSIS (FDPVA)

S. M. Gonzalez de Vicente
International Atomic Energy Agency, Department of Nuclear Sciences and Applications
Vienna, Austria
Email: S.M.Gonzalez-De-Vicente@iaea.org

D. Mazon
CEA, IRFM
Saint-Paul-lez-Durance, France
Email: Didier.MAZON@cea.fr

M. Xu
Southwestern Institute of Physics
Chengdu, Sichuan, China
Email: minxu@swip.ac.cn

S. Pinches
ITER Organization
Saint-Paul-lez-Durance, France
Email: Simon.Pinches@iter.org

M. Churchill
Princeton Plasma Physics Laboratory
Princeton, NJ, United States
Email: rchurchi@pppl.gov

A. Dinklage
Max Planck Institut fuer Plasmaphysik
Greifswald, Germany
Email: dinklage@ipp.mpg.de

R. Fischer
Max Planck Institut fuer Plasmaphysik
Garching, Germany
Email: rainer.fischer@ipp.mpg.de

A. Murari
Consorzio RFX, CNR, ENEA, INFN, Università di Padova
Padova, Italy
Email: andrea.murari@istp.cnr.it

P. Rodriguez-Fernandez
MIT Plasma Science and Fusion Center, Boston, USA
Email: pablorf@psfc.mit.edu

J. Stillerman
MIT Plasma Science and Fusion Center, Boston, USA
Email: jas@psfc.mit.edu

J. Vega
CIEMAT
Madrid, Spain
Email: jesus.vega@ciemat.es

G. Verdoolaege
Ghent University, Department of Applied Physics
Gent, Belgium
Email: geert.verdoolaege@ugent.be

Abstract

The objective of the fourth Technical Meeting on Fusion Data Processing, Validation and Analysis was to provide a platform during which a set of topics relevant to fusion data processing, validation and analysis are discussed with the view of extrapolating needs to next step fusion devices such as ITER. The validation and analysis of experimental data obtained from diagnostics used to characterize fusion plasmas are crucial for a knowledge-based understanding of the physical processes governing the dynamics of these plasmas. This paper presents the recent progress and achievements in the domain of plasma diagnostics and synthetic diagnostics data analysis (including image processing, regression analysis, inverse problems, deep learning, machine learning, big data and physics-based models for control) reported at the meeting. The progress in these areas highlight trends observed in current major fusion confinement devices. A special focus is dedicated on data analysis requirements for ITER and DEMO with a particular attention paid to Artificial Intelligence for automatization and improving reliability of control processes.

Keywords: Integrated Data Analysis, Data Validation, Bayesian Techniques, Neural Networks, Machine Learning, Disruption Predictors, Image Processing

1. Introduction

The fourth IAEA Technical Meeting on Fusion, Data Processing Validation and Acquisition (FDPVA, November 29 to December 6 2021), reviewed pre- and post-

processing, calibration and validation of measured nuclear fusion research data. The meeting was held remotely due to the global Covid-19 pandemic. The event was organized by IAEA (remotely hosted by Centre for Fusion Science, Southwestern Institute of Physics) and brought together more than 100 scientists and engineers working on instruments, methods and mathematical solutions for research in the field of nuclear fusion and plasma physics.

"We are entering in a more complex world of data analysis thanks to the huge number of measurement systems equipping present day tokamaks and for that reason artificial intelligence should be developed in a more systematic way to ease plasma discharge analyses" said Didier Mazon (Co-Chair of the International Programme Advisory Committee) in his introductory talk. Progress made in that direction was shown during the meeting. In particular, new developments in the following fields were discussed: Data analysis preparation for ITER and Software Tools for ITER diagnostics, Data Analysis for Fusion Reactor, Uncertainty Propagation of Experimental Data in Modelling Codes, Applications of probabilistic inference and statistics, Real time prediction of off-normal events, with particular attention to disruption and predictive maintenance; Image Processing, Deep Learning, Inverse Problems, Causality Detection in Time Series, Synthetic Diagnostics, Integration, Verification and Validation, Integrated Data Analysis, and Big data. Material can be found here: https://conferences.iaea.org/event/251/timetable/?view=standard_numbered_inline_minutes. Part of the material used in this summary paper is taken from this website.

Relevant tutorial sessions were held: Introduction to IMAS (S. Pinches, ITER); Introduction to Scientific Machine Learning and Deep Learning (M. Churchill, PPPL); and Introduction to Bayesian methods and their use in fusion data analysis (G. Verdoolaege, Univ. Gent). Video material of the tutorial sessions can be found here: https://nucleus.iaea.org/sites/fusionportal/_layouts/15/guestaccess.aspx?guestaccesstoken=SasuQp45gsvslwjrmqC%2fcBRbBFEzhzTs7MfdYo1BHgQ%3d&docid=2_1ad1e517894c54098b9e960a07189adf5&rev=1

2. Summary of the meeting sessions

This section briefly summarizes the 11 sessions that covered specific topics of interest for fusion data processing, validation and analysis, focusing on main highlights, progress and outcomes of the general discussions.

2.1. Data analysis preparation for ITER and software tools for ITER diagnostics (DAP) - S. Pinches

This session covered a broad range of topics and provided an opportunity for the fusion community to highlight particular areas of recent progress in the areas of data analysis and the development of software tools for ITER.

A physics area of special interest to ITER, given its mission to create a burning plasma dominated by alpha-particle heating, is that of energetic particle stability and this was particularly touched upon by work using deep neural networks to classify observed energetic particle driven Alfvén Eigenmodes (AE) in DIII-D[1]. The work provided a

good proof-of-principle test showing the capability of simple yet effective models in identifying AE based only on Electron Cyclotron Emission (ECE) measurements. Given the potential need for real-time control of such modes, e.g. to avoid deconfining alpha-particles in ITER and other future devices before they slow-down and pass on their energy, this motivates further work in this area.

Another area where real-time data analysis made use of neural networks was on EAST to predict the breakdown of the neutral beam injector (NBI) system[2]. Compared with the challenging work on predicting tokamak disruptions, the prediction of breakdown events for the NBI ion source was found to be relatively straightforward. Due to the short timescales associated with the breakdown and diagnostic response, the only practical implementation found was to use field-programmable gate arrays (FPGAs).

The creation of data analysis software benefits greatly from open-source software and one relevant example presented was that of ToFu, a python library that supports the creation of synthetic diagnostics and tomographic inversion[3]. It contains representations of various tokamaks including ASDEX Upgrade, ITER, SPARC and WEST and is interoperable with other tools for tokamak plasma tomography such as Tomotok[4].

Another development in support of experimental data analysis that was presented in this session was the Mori-Zwanzig projection operator method[5]. This is a statistical tool to analyze correlation among time-series data. In the work presented, it was applied to the interactions of turbulence and zonal flows and used to help derive physical insights.

One of the most well-known core-transport codes used in the tokamak fusion community is the TRANSP code [6] and it was reported that the code was undergoing a significant refactoring and modernization process. In particular, modules were being updated and made external, and interfaces were being changed to use the standard Interface Data Structures (IDSs) of ITER's Integrated Modelling & Analysis Suite (IMAS). Indeed, the intention was to eventually replace the current Plasma State used within TRANSP with the corresponding set of IMAS IDSs. As a first example, the Multi-Mode Model (MMM) for anomalous transport has been turned into a stand-alone library that uses IDSs for input and output.

Before ITER starts operating, the development and testing of synthetic diagnostics and analysis tools for ITER is based upon an extensive set (2000+) of ITER scenario simulations. Cataloguing these using a new simulation management tool, SimDB, was

the focus of another of the presentations within this session.
The requirements for SimDB were to make it easy to find simulations matching given criteria, to facilitate the acceptance and deprecation of datasets, to be able to validate datasets against prescribed rules, and to make it easier for users to fetch the identified datasets. These are satisfied with a self-documenting command line tool and a web-based dashboard that exposes the catalogue of simulations to users, neither of which are specific to ITER.
In effect, SimDB, supports simulation data becoming more FAIR (Findable, Accessible, Interoperable, Reusable), a recurrent theme throughout the Technical Meeting which was covered in around five dedicated presentations as well as a demonstration in the opening tutorial session.
One of these related presentations was also in this session and gave an update on the architecture for the implementation of a FAIR fusion data framework. This was as part of the EU's Fair4Fusion project[7] that aims to demonstrate the benefits of making experimental data from fusion devices more easily findable and accessible.

2.2. Data analysis for fusion reactor (DAT) - D. Mazon

The session on data analysis for fusion reactor focused on the fast characterization of plasma behaviour (states, profiles, edge behaviour) through different models and automated techniques, in view of reactor performance control which were summarized in this section.
The first technique presented was the Permutation Entropy (PE). PE is an information-theoretic quantity that measures the complexity of time series. This measure has been successfully implemented in different science branches e.g. medicine (detection of epileptic electroencephalogram) and economics (characterizing complexity changes in stock market). In practice, PE reduces the description of complexity to a single number through the probability distributions of ordinal patterns (permutations) in consecutive data. The main reasons of introducing this method for plasma studies lie in the simplicity of the method, making it extremely fast and robust to compute. This method is fast given that it is based on sorting algorithms rather than traditional distance calculations. It is also robust since it is an ordinal method resulting in invariance against transformations preserving ordinal rankings between measurements. In sliding window analysis of a single information channel, a change of PE can indicate a bifurcation of the system state. Therefore, the PE approach is applied on large data sets of highly sampled plasma data in an automated procedure.

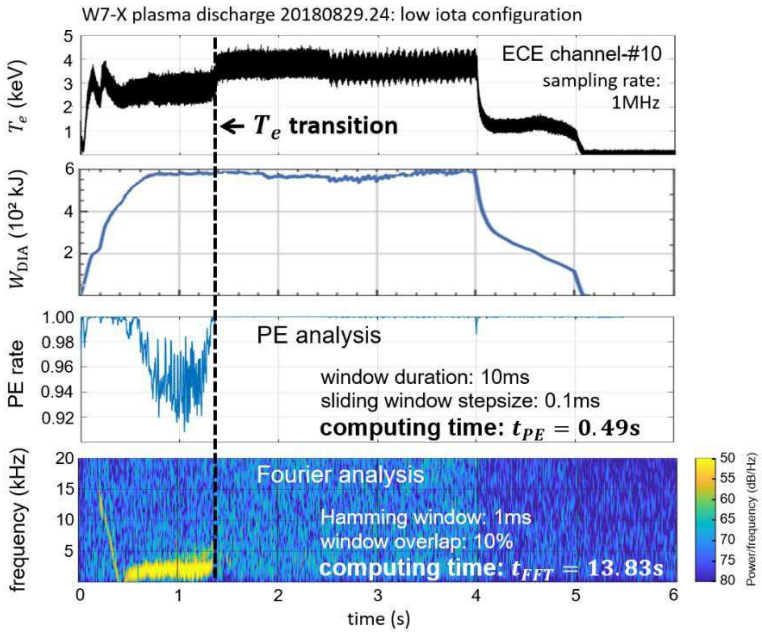

Figure 1. Comparison between Fourier and PE analysis for plasma state detection

Fast characterization of plasma states through PE analysis of bulk data from W7-X plasmas was reported in the session. The specific case applying PE presented was the analysis of a multi-variate, highly sampled time series from an electron cyclotron emission radiometer and a soft X-ray diagnostic. The bulk processing was employed to investigate the parameter dependencies such as different heating powers and densities. Spatio-temporal changes of the plasma states were detected from emissivity changes resulting in significant alterations of the PE in individual data channels, see Figure 1. The reason for the sensitivity of PE was identified (a posteriori) to be correlated with the occurrence of low frequency emissivity fluctuations, which cease when a spontaneous transition to high core-electron temperatures occurred. A T_e transition was detected and localized close to the plasma center. Also, a counter-acting re-arrangement of temperature and apparent decreasing density was observed, preserving the total amount of energy. These results are in accordance with previously unrevealed changes in plasma profiles. The identification of spontaneous plasma transition periods was validated by spectrogram analysis.

While visual inspection of the (noisy) data allows one to conclude state changes, the time to identify the bifurcation is much reduced when automated analyses with PE were conducted. This acceleration in the processing time allowed analysis of a large amount of data and to detect systematic changes in the plasma state in a set of experiments. This suggests that a complex measurement such as PE can support in-situ monitoring of plasma parameters and for novelty detection in plasma data. PE is therefore proposed as a method for big-data-processing of plasma data. Moreover, the acceleration in processing time offers provides results fast enough to induce control actions even on the time scale of the experiment.

multi-fluid plasma and neutrals interactions code SOLPS- ITER [8] was used to demonstrate model predictive control (MPC) of key variables in the tokamak plasma edge. Though SOLPS-ITER provides state-of-the-art simulation of the scrape-off-layer (SOL), it takes upwards of weeks-months time for the convergence of a steady-state

solution at an ITER configuration [9]. This computational expense makes SOLPS-ITER predictions unavailable for real-time analysis to allow mitigation measures against plasma facing component damage due to excessive heat flux loads from plasma escaping core confinement in a fusion reactor. Following the theoretical two-point model formalism[10], connecting upstream SOL conditions with downstream divertor target parameters, an interpretable reduced model of the plasma boundary in response to neutral gas puff actuation that mediates detachment power dissipation was proposed.

The Sparse Identification of Nonlinear Dynamics (SINDy) [11] was deployed to model point measurements from SOLPS-ITER as a coupled system of ODEs with respect to the level of controlled actuation. SINDy promotes sparsity in the selection of terms from a candidate library of functions using a regularizer on the optimization function to ensure the simplest description of a given system. Figure 2 shows the application of the procedure to output time series from SOLPS-ITER of the outboard midplane separatrix electron density and outboard divertor target separatrix electron temperature for a DIII- D configuration. Perturbations from a fixed point steady-state are obtained through a scan of gas puff rates in order to excite a range of dynamics for feature selection of the machine learning algorithm. Rolling cross-validation for this offline demonstration was used to determine the viability of the extracted system of equations. Starting from $t_0 = 0.18$s, the reduced model is trained over an incrementally increasing interval. As each candidate model is obtained, an out-of-sample prediction is calculated over the rest of the testing data. An error threshold is applied over the eleven deviations from the SOLPS-ITER simulation, which when cross-triggers a restart of the training routine.

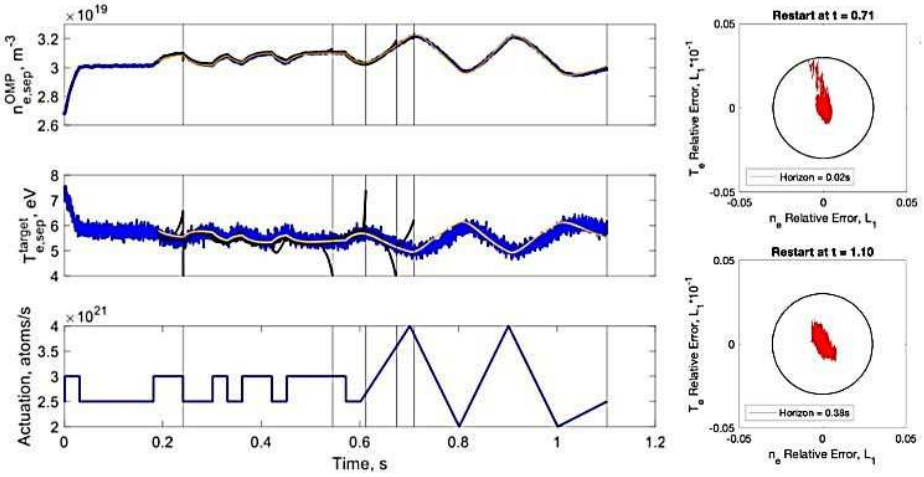

Figure 2. Demonstration of rolling cross-validation in the extraction of SINDy reduced model

In Figure 2, the final selected model is shown in left hand panels in gold against the reference simulation in blue. Rejected models are shown in black and terminated at the restarted training demarcation in gray. For this data set, a coupled linear system between the outboard midplane separatrix electron density and outboard divertor target electron temperature was obtained. The right hand panels of Figure 2 show the running deviation (in red) of the last two candidate models with the final system of equations achieving a prediction horizon from last restart to the end of the available testing data of time prediction horizon (tph) = 0.38s. The SINDy procedure was shown to be deployable in real-time with each model extraction taking only 8 ms per iteration.

Gaussian Process Regression (GPR) was also presented. It is a Bayesian method for inferring profiles based on input data. The technique is increasing in popularity in the fusion community due to its many advantages over traditional fitting techniques, as it includes intrinsic uncertainty quantification and demonstrates robustness to over-fitting. Most fusion researchers to date have utilized a different GPR kernel for each tokamak regime. This requires a Machine Learning (or simpler) methods to first predict the regime, choose the right kernel for that regime, and then use that kernel. The disadvantage of this method is that it requires an additional step, and it is unclear how well it will behave if the plasma enters a new, unexpected regime. "The methodology aims at developing a general kernel for all regimes (including radially-varying hyperparameters), utilizing heavy-tailed likelihood distributions to automatically handle data outliers and using GPflow for full Bayesian inference via Markov chain Monte Carlo to sample hyperparameter distributions. A single GPR method that is robust across many different tokamak regimes and a wide range of data inputs and quality was presented[1]". Through the choice of a piecewise kernel, the length scales in the pedestal and the rest of the profile can be optimized separately to fit the whole profile very well, as shown by Figure 3.

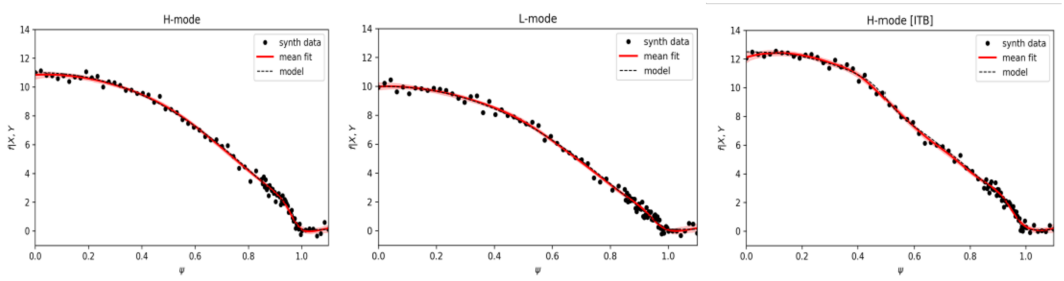

Figure 3. Evidence of profile fitting in different plasma mode

Additionally, the use of a student-t likelihood function allows for fitting the data even when outliers are present thanks to the possible heavy tails of the distribution. Likewise, if there are no outliers, the student-t degree-of-freedom parameter is optimized to a large value, where the function approaches a gaussian. In this way, the error of the fit can remain small and avoid having the mean pulled askew by the outlying data.

Finally, it was demonstrated that digital twins are capable of predicting plasma evolution ahead of plasma progression within a Tokamak and are crucial tool required for real-time plasma intervention and control. Considering speed and scale required, quite often these have to be purely data conditioned models as opposed to being physics

conditioned, making data selection a vital component of model efficacy. However, as we move to the exascale regime, the amount of data generated tends to choke the data pipelines, introducing latency to the model. It might also be the case that some of the data available might be redundant and creating imbalances within the training dataset. In this work it was demonstrated that a machine learning pipeline mapped out in hyperspace the distributions of the plasma behaviors within a specific campaign. The embedding created through dimensionality reduction within the pipeline was then used as the sampling space for the training dataset for a Convolutional LSTM that mapped the control signals to diagnostic signals in a sequential manner. Primarily experiment with MAST data were performed with the control signals being plasma current, toroidal magnetic field, plasma shape, gas fueling and auxiliary heating. The diagnostics of interest were the core density and temperature as measured by the Thomson scattering diagnostic. With initial focus on a single experimental campaign (M7), it was demonstrated that the predictive model trained on all available data is capable of achieving a mean squared error of 0.0285. However, the pipeline demonstrated that by using a distance based informed sampling method to gather only 10 percent of the dataset a comparable mean squared error of 0.0293 can be achieved. It was further demonstrated that the robustness of the pipeline by extending the model to operate within the space of the M9 campaign in addition to the M7 campaign. This work showed that a predictive model trained on all of the available data across both campaigns achieves a mean squared error of 0.0279, while the one sampled using the knowledge garnered from the cluster representations (mapped individually across each campaign) achieves an L2 error of 0.0282, while only relying on 10 percent of the dataset.

2.3. Applications of probabilistic inference and statistics (API) - G. Verdoolaege

This session was devoted to the methods and applications of probabilistic inference and statistics. The scope included a variety of data science activities, such as parameter estimation, model comparison, uncertainty quantification and propagation, etc.

Several presentations, including the invited one, concerned applications of Bayesian probabilistic inference. Herein, the advantage of Bayesian methods was exploited, as they provide a framework for rigorous analysis of error propagation and integrated treatment of heterogeneous sources of data. Moreover, various machine learning techniques have their roots in Bayesian methods, hence providing motivation for explicitly formulating the assumptions and approximations that go into the analysis. On the other hand, Bayesian inference can require significant computational resources, either for approximating (marginal) posterior distributions or for sampling from them. Hence, in applications dealing with complex forward models, notably those involving modeling codes, or when targeting real-time applications, specialized techniques need to be considered that can speed up the inference process.

This leads to an important application of Bayesian methods, i.e. the inference of parameters in modeling codes, like transport coefficients, from experimental data. For each calculation of the forward model, the (transport) code has to be run, so probabilistic inference often requires high-performance computing. The particular case that was discussed at the meeting concerned the inference of particle transport coefficients in tokamak plasmas[12]. The inference was based on spectroscopic measurements during impurity injection experiments using laser blow-off at Alcator C-

Mod and DIII-D, combined with 1.5D impurity transport modeling and radiation forward modeling using the Aurora toolbox[13]. Aurora takes into account the influence on the impurity ionization balance of charge exchange reactions with neutrals. Nested sampling on a high-performance computing platform was used to determine the optimal

spline model for the profile of transport coefficients and to perform transport coefficient sampling. This allowed identifying discrepancies with results from neoclassical and turbulence codes in case of flat or hollow impurity profiles.

Fusion data is other common application area of Bayesian inference. Recent developments were presented at the meeting of a platform, called “Retina”, toward integrated data analysis (IDA) using Bayesian methods at the HL-2A tokamak, with a view to IDA for the new HL-2M tokamak. This has so far been applied to estimation of density and temperature profiles from diagnostic measurements of electron cyclotron emission, Thomson scattering and reflectometry[14]. Tomographic inversion of the emissivity profile from soft X-ray spectroscopy and bolometry was also demonstrated, as well as the reconstruction of the plasma current profile from magnetic coil measurements. As a prior distribution, Gaussian processes were used in a radial or poloidal cross-section, allowing great flexibility of the inferred profiles. In specific cases, where the forward model is linear or can be linearized without great loss of fidelity, the inference is very fast, potentially opening up real-time applications. Indeed, in those cases the posterior distribution is multivariate Gaussian with mean and covariance available in closed form.

Inference of parameters of plasma filaments in the edge and SOL of MAST was also discussed at the meeting. To do this, signals of ion saturation currents were obtained from a reciprocating Langmuir probe and Bayesian probability was used for fitting the filament dynamics. Particular attention was paid to the nature of the background signal, which was seen to be influenced by small filaments that are only weakly constrained by the data. In a next step, the distribution of filament characteristics was explored and the filament distribution was seen to be well described by a Poisson distribution.

An interesting development that was presented at the meeting: simulation- based (or likelihood-free) inference by means of neural networks (see e.g.[15]). This encompasses a number of techniques for probabilistic inference that attempt to circumvent the issue of computationally heavy forward models involving simulations. Traditionally, techniques based on sampling and rejection were used, referred to as “approximate Bayesian computation”, which still require significant computational resources. However, with the recent revolution in the domain of machine learning, it has become feasible to learn the likelihood or even the full posterior distribution by means of a neural network model. This scheme was applied to inference of SOL transport coefficients from fluid simulations using the UEDGE code as a black-box simulator. The particular technique used to approximate the posterior is known as a “normalizing flow”, which refers to learning a series of Jacobians in transforming from a space in which the distribution is multivariate normal. In an alternative approach, known as “amortization”, a neural network learns the mapping between prior and posterior for generic experimental data. This scheme, which is foreseen in future work, has the advantage of allowing fast statistical inference during experiments.

A well-known technique from the control community that is firmly rooted in Bayesian probability is the Kalman filter. It is used for data assimilation based on a system model and a series of measurements of quantities to be controlled. In an application at LHD, an ensemble Kalman filter and smoother was used for prediction of density and

temperature profiles, based on the TASK3D integrated modeling code as a system model. The method was implemented in the ASTI data assimilation system and results of numerical experiments were shown, aiming at control of the central electron and ion temperatures, using electron cyclotron heating power as a control parameter[16]. The technique will be used for control of real plasmas at LHD and possibly other devices.
A common issue when analyzing databases including a significant number of quantities that span a broad variety of plasma or machine conditions, is related to finding meaningful structure in such complex data. In particular, it is important to ensure sufficient robustness of the model and the fitting methods, especially in higher-dimensional databases that are difficult to probe for structure. Examples in fusion are the scaling laws fitted to multi-machine data, like the energy confinement scaling. Following a recent revision of the well-known IPB98(y,2) scaling, based on a revised version of the international global H-mode confinement database, it was seen that the dependence on machine size was reduced considerably in the new ITPA20 scaling[17]. This is an important cause for the lower confinement time prediction for ITER by the new scaling. Ongoing work was presented at the meeting aimed at explaining the weaker scaling with major radius. By means of optimization techniques, the smallest subset of the new database was obtained that has the largest influence on the size scaling. This will allow characterizing the operational conditions exhibiting the weakest size dependence, which will provide crucial information toward confinement scaling in ITER.
It is clear that data-driven techniques in fusion can provide insights and tools that are complementary to physics-based studies. Fitting semi-empirical models like scaling laws is one instance where this complementarity has long been exploited. With the more recent introduction of powerful data science methods from statistics and machine learning, the wealth of opportunities has become even more apparent. At the same time, ample progress is still to be made in merging these two views on scientific discovery, particularly in fusion. An example discussed at the meeting concerned statistical model comparison for determining the main explanatory dimensionless variables in a power law model for the thermal diffusivity in LHD[18]. A database of diffusivities was compiled using the TASK3D-a code and a modified Akaike information criterion was invoked to extract the most informative variables. Thus, model comparison and feature extraction can contribute to elucidating the physics of transport processes. This interaction between statistics, machine learning and domain knowledge regarding the physics and technology of fusion devices is expected to only increase in the future.

2.4. Real time prediction of off-normal events, with particular attention to disruptions and predictive maintenance (RTP) - A. Murari

In metallic devices, the occurrence of disruptions is particularly difficult to predict because of the nonlinear interactions between various effects, such as neoclassical convection of impurities, centrifugal forces, rotation, profile hollowness and magnetohydrodynamic (MHD) modes, just to name a few. While efforts to develop physics based plasma simulators are continuing, data driven predictors, based on machine learning, remain an important fall back solution.

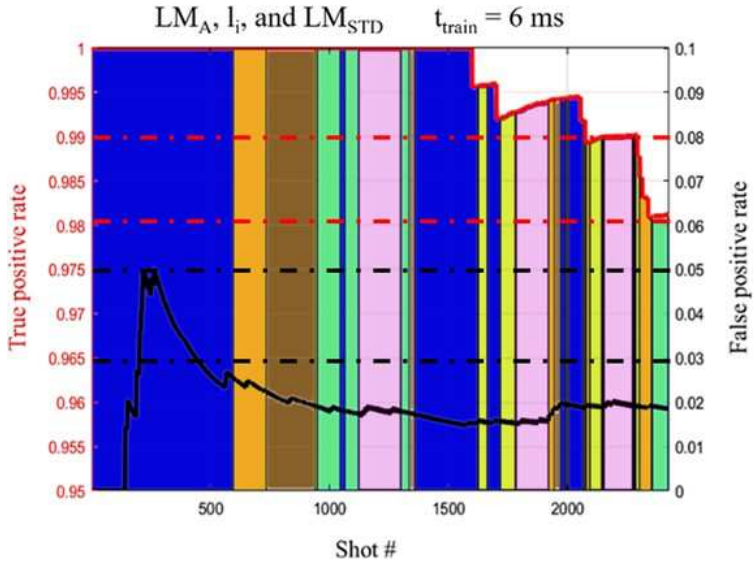

Figure 4. Performance of the adaptive predictor for mitigation on JET. Left axis (red): the success rate. Right axis (black): the false alarms rate. The horizontal axis reports the shot number. The different colours correspond to different decision functions, used to synthesise the outputs of the individual classifiers in the ensemble

Disruption predictors based on traditional machine learning have been very successful in present day devices but have shown some fundamental limitations in the perspective of the next generation of tokamaks, such as ITER and DEMO. In particular, even the most performing require an unrealistic number of examples to learn, tend to become obsolete very quickly and cannot easily cope with new problems. These drawbacks can all be traced back to the type of training adopted: *closed world training*. In the last years, it has been shown how the new approach of *open world training* can solve or at least significantly alleviate the aforementioned issues. Adaptive techniques, based on ensembles of classifiers, allow following the changes in the experimental programmes and the evolution in the nature of the disruptions[19][20]. This approach has been implemented with ensembles of classification and regression trees classifiers. Some of these predictors have achieved the best performances ever obtained in JET, in terms of both success rates and false alarms, as shown in Figure 4[21][22]. Exploiting unsupervised clustering, new predictors can autonomously detect the need for the definition of new disruption types, not yet seen in the past[23]. All the solutions can be implemented *from scratch*, meaning that the predictors can start operating with just one example of disruptive and one of safe discharge[24][25].

In the perspective of contributing to the safe operation of new large tokamaks, being able to transfer experience from one device to another would be also very beneficial. A procedure to deploy predictors trained on one device at the beginning of the operation of a different one has been developed[26]. The proposed tools were tested by training these classifiers using ASDEX Upgrade data and then deploying them on JET data of the first campaigns with the new ITER Like Wall[27][28]. The obtained results were very encouraging. After a transition learning phase, in which in any case the performances remained sufficiently high, the predictors managed to meet the ITER requirements for mitigation in terms of both success rate and false alarms. Promising improvements have also been achieved for prevention using, in particular,

information about the radiation profiles and visible cameras. The proposed techniques would be particularly valuable at the beginning of the operation of new devices, when experience is limited and not many examples are available. Implementation of different more advanced metrics to determine the distance of the operational points from the disruption boundary are under investigation[29].
The development of techniques to improve the interpretability of machine learning techniques, so that they can be used in support to theory formulation, is also progressing significantly[30][31][32][33]. A crucial aspect of these techniques is their integration within the plasma control systems and their implementation under deterministic conditions.
Finally, significant advances in machine learning methods, to perform unsupervised clustering of the disruptive phase of discharges to find common termination paths, have been achieved[34]. It should be noted that, at present, also solutions based on deep learning are trying to give answers to several problems such as feature extraction and transfer learning. Deep learning is being applied to specific diagnostics as well (for instance, magnetics or beam emission spectroscopy) to recognise disruption precursors.

2.5. Image processing (IMP) - J. Stillerman

The content of this session was about automated processing of IR camera images to detect hot spots. Groups from both CEA and W-7X expect to apply these techniques to their plasma control systems in real-time. The labor to acquire labeled data sets from existing videos is significant. An overall description of the problem, acquiring, labeling, analyzing and automating hot spot detection using IR cameras was provided. Further details on the deep learning platform in use at WEST were given. The pipeline consists of a Cascate R-CNN step to identify the hot spots on a frame-by-frame basis. This is followed by a classification step where the hot spots are categorized into the event ontology. Similar work was done on W7 -X. The pipeline in this case used background subtraction, max-tree classification, pruning, and then back to generated images for classification. There is not enough annotated data to apply deep learning techniques to the data.

2.6. Deep learning (DEL) - P. Rodriguez-Fernandez

The focus of the deep learning session has been twofold: learning from expensive simulation codes or learning from experimental measurements. This session included content involving the use of deep learning and other machine learning models for applications in plasma physics and fusion energy research.
The analysis of experimental data and the use of computational models to predict plasma behavior in magnetic confinement fusion devices are often hindered by the large computational cost of the analysis and modeling techniques. In many situations, the unfeasibility of measuring every possible parameter during experiments has also hampered our ability to interpret and thus predict plasma behavior. Techniques based on deep learning can be used to accelerate the interpretation, analysis, modeling and

prediction of plasma experiments. Databases of simulation results and diagnostic signals can be fed into deep-learning models that can provide predictions of plasma quantities (based either on models or on experimental expectations) extremely fast, even reaching real-time capabilities in some situations. The session was dedicated to such techniques to accelerate our understanding and our predictive capabilities of magnetic confinement fusion devices.

Deep learning models have been applied extensively to reproduce the output of simulations, with the goal of facilitating big-data validation of the physics assumptions that go into the models and improve their predictive capability. Furthermore, for control and real-time purposes, such reduced models can be used to guide tokamak discharges into high-performing and low-risk parameter spaces. During the deep learning session, techniques to improve transport modeling with codes such as the Trapped Gyro-Landau Fluid (TGLF) solver were discussed, and benefits of using big data validation techniques were highlighted. Regions of the parameter space can be identified where models fail to reproduce experiment, and correction factors based on plasma parameters can be applied to improve their predictive capability. A promising technique to improve transport modeling with the TGLF turbulence model was presented. The use of big data validation (over 200k simulation results and comparison to experimental data) gave insights on the region of the parameter space in which TGLF fails to reproduce experimental fluxes. With this information in hand, predictions from TGLF can be improved by multiplying the output fluxes by error factors expressed as a function of plasma parameters. Not only deep learning models can be used to study the validity and predictive power of reduced transport models, but also neural network models can be used to reproduce the nonlinear dynamics of turbulence. In particular, the dominant turbulence type, radial profiles, and time-evolution can be predicted and used to accelerate the convergence of nonlinear codes, such as the extended fluid code (ExFC), and reduce the overall computational time. Machine learning techniques to find such factors are promising. Neural network models to predict the turbulence type and radial profiles that result from fluid simulations with the ExFC code were presented. In particular, the use of recurrent neural networks (RNNs) is promising to predict future time slices of the simulation, and can be combined with the real ExFC code to reduce the overall computational time.

As another application of deep learning models, they can also be applied directly to experimental data to predict quantities and associated uncertainties, which are useful to optimize and guide plasma discharges. This was a topic of active discussion during the deep learning session. The use of machine learning to predict kinetic profile shapes and turbulence features in experimental tokamak discharges has been proven useful and has a two-fold application. On the one hand, the experimental information to be extracted from the plasma discharge can be extended, as more information is available from the diagnostic systems. Connected to this point, a technique to predict kinetic profiles shapes using neural networks was presented. Experimental information is input into the model and the time traces of kinetic profiles can be predicted. This has a direct application to inform the exploration of the parameter space and as a tool for physics operators to attain reliable and high performing plasma discharges. The use of deep

learning models to exploit the 2D capabilities of the BES diagnostic was presented. On the other hand, if the deep learning models are fast enough, they can be used to inform the exploration of the parameter space and as a tool for physics operators to attain reliable and high performing plasma discharges. Classification of tokamak discharges into confinement regimes and prediction of the onset of edge localized modes from diagnostic signals were clear examples where deep learning models can enhance our knowledge of tokamak plasma physics and inform discharge planning. The use of deep learning methods to accelerate the prediction and training of Bayesian models are also promising for real time applications. The session also featured a discussion of transfer learning, which is key to ensure the success of upcoming burning plasmas such as ITER and SPARC. In particular, the study of disruption prediction algorithms in current devices and their extrapolation to other machines is very important for this task. Work focused on the identification of confinement regimes in tokamaks using seq2seq models was shown. Automatic classification of these regimes greatly aids the labelling of experimental discharges, and can be used to gather insights in what triggers mode transitions. Hybrid deep neural networks for disruption predictions, with an emphasis placed on transferability were presented. Reliably predicting disruptions in ITER is key to its success and the question of whether the models developed in current devices can be readily transferred is important for this task. Finally, the fundamentals of the Minerva modeling framework that enables the implementation of physics models and uncertainties to infer plasma quantities were discussed. Leveraging neural networks to learn the Bayesian model joint probability distribution provides avenues to make computationally cheaper Bayesian inferences, that can eventually be employed in real time to reliably predict plasma parameters.

2.7. Inverse problems (INP) - M. Churchill

The session on inverse problems contained a fascinating array of research into using Bayesian analysis and other algorithms for extracting physics parameters from experimental diagnostics which have integral relations with these physics parameters. These problems are often ill-posed and require strong biases in the algorithms to invert, or accurate synthetic diagnostics to leverage in Bayesian analysis to extract the physics parameters. A common theme was leveraging machine learning, in particular neural networks, for various tasks within the workflows.

An example of using deep learning to perform approximate Bayesian inference was presented, applied to many line-integrated diagnostics to extract physics parameters, e.g. the X-ray imaging crystal spectrometer (XICS) on W7-X to extract ion and electron temperature. The Minerva framework was used to create synthetic diagnostics and a Bayesian model, and a deep neural network then trained on many synthetic samples to learn the inverse function mapping of diagnostic data to physics parameters. The benefit of using a deep neural network for approximate Bayesian inference is that the inference is much faster, in this case 100 microseconds, versus the 10 minutes needed for traditional Bayesian inference methods in Minerva.

A review of 2-D tomographic reconstruction algorithms used on the EAST tokamak for various diagnostics such as soft X-rays (SXR) showed that a newer method called Gaussian Process Tomography performs very accurately and fast for a number of

diagnostics, which in the past have used varying algorithms depending on the diagnostic. Various convolutional neural networks were also implemented, which showed good performance and even faster results. Application examples such as extraction of MHD mode structure analysis were demonstrated with these tomographic inversion methods. Bayesian experimental design principles were used in the design of the soft X-ray system on the Keda Torus eXperiment (KTX). This allowed determining information gain about targeted physics parameters of interest (e.g. radiation in the plasma edge) for given design parameters (e.g. number, location of sightlines). Bayesian methods were also applied for optimal settings for the tomographic reconstruction process.

A project labeled EFIT-AI was presented, which modernizes and applies machine learning to the popular EFIT code used for magnetic equilibrium reconstructions in tokamaks based on magnetics and other diagnostics. Modernizing the make system with CMake and parallelizing led to a 7x speedup, making higher-resolution grids more feasible to compute. A large dataset of magnetic equilibriums calculated by EFIT were gathered, to train fully-connected feed forward neural networks to replicate parts of the reconstruction process accurately. A flexible Gaussian Process Regression algorithm was made to fit plasma profiles across a range of plasma conditions, allowing to better automate kinetic EFIT runs. Finally, SVD-based model order reduction is being explored to capture the 3D perturbed equilibriums from an MHD code MARS-F.

Extracting temperature of material walls from IR camera measurements is critical for machine protection, but difficult due to reflections, instrument calibration, etc. A digital twin approach was taken to create a detailed, end-to-end simulation modeling all physical phenomenon, from source to optical response of instrument. A reduced photonic model assuming diffuse surface reflection was included, and an iterative solution to extract the temperature comparing the output of this synthetic IR model to prototype IR camera measurements was used, achieving excellent agreement, and on timescales that can be useful for real-time interpretation. Further work adding additional realism to the model and faster ways to extract the temperature (e.g. deep neural networks) are planned.

Reconstructing the electron density using line-integrated interferometer diagnostics on EAST is important for feedback control of density. A difficulty is identifying when incorrect diagnostic data is present in the many interferometer channels used to do the inversion. Neural networks were employed to learn the inverse mapping of interferometer channel data to density profiles, and an algorithm for accounting for bad channels was employed, often replacing these values with averages of neighboring channels to input into the neural network. It was shown that this method works very well, even when up to four channels of interferometer data are missing.

Integrated data analysis on the HL-2A device was performed employing magnetic coils and interferometers for plasma current tomography. A new Advanced Squared Exponential (ASE) prior was used in the Bayesian inference, showing better accuracy and robustness to noise compared to the previously used Conditional Autoregressive (CAR) prior. A neural network was trained to find reference discharges most suitable for a particular shot, which aids in the reconstruction process.

2.8. Causality detection in time series (CDT) - J. Vega

Causality is a crucial aspect of human understanding and therefore one would expect that it would play a major role in science and particularly in statistical inference. On the contrary, traditional statistical and machine learning tools cannot distinguish between correlation and causality. This lack of discrimination capability can have catastrophic consequences for both understanding and control, particularly in the investigation of complex systems. The field of so-called Observational Causality Detection (OCD) is devoted to refining techniques for the extraction of causal information directly from data. In the last years, a conceptual framework, based on the concept of intervention, has been developed to substantiate the statement that correlation is not causality. The translation of such a conceptual framework into mathematical criteria applicable to times series is progressing. The proposed tools can be classified into two major categories: those based on the analysis of the system dynamics in phase space (such as Convergent Cross Mapping and Recurrence Plots) and those relying on the statically and information theoretic properties of the data (such as transfer Entropy and Conditional Mutual Information). More recent techniques are based on neural networks of specific topologies.

In fusion devices, as in many other experiments, time series are the typical form of the signals produced by the measuring systems. The detection of causality between time series is therefore of great interest, since it can give a unique contribution to the understanding, modelling, and prediction of phenomena still not fully understood. However, detecting and quantifying the causal influence between complex signals remains a difficult task, not solved yet in full generality.

The next generation of Tokamaks and the future reactor will be operated relying much more on feedback than present day machines. The control of macroscopic instabilities, such as edge-localized mode (ELMs) and Sawteeth, will be essential. In this perspective, various pacing experiments have been indeed successfully carried out in many devices in the framework of scenario optimisation. In the case of external pulse perturbations, the idea consists of triggering the instabilities sufficiently often that their crashes do not reach excessive proportions. Unfortunately many details of their interactions with the plasma remain poorly understood. Since both instabilities are quasi periodic in nature, it is difficult to determine the efficiency of pacing schemes such as pellets, vertical kicks or ion-cyclotron resonance heating (ICRH) notches. Indeed, after each of these perturbations, if enough time is allowed to elapse, an ELM or sawtooth crash is bound to occur. Quite sophisticated data analysis methods have been devised to assess this aspect[35][36]. Their deployment to investigate ELM pacing with pellets and sawteeth triggering with ICRH modulation have provided very interesting results and have allowed determining the efficiency of these synchronisation experiments quite reliably[37][38]. Another intriguing detail is relative importance of phase and amplitude in frequency synchronization. A data analysis methodology for investigating this aspect has been also developed. The technique is based on the wavelet decomposition of the signals and information theoretic indicators, to determine the actual form of the interactions. In both JET and ASDEX Upgrade coherent results have been obtained. The main effect, in both ELMs pacing with pellets and sawteeth synchronization with ICRH modulation, is due to the influence of the amplitude of the

external perturbations. Some evidence of phase synchronization has been found, which could show the direction of future optimization of the interventions.
A new causality detection method based on Time Delay Neural Networks (TDNNs) has been recently developed. The architecture of TDNNs is sufficiently flexible to allow predicting one time series, on the basis of its past and the past of others. With suitable statistical indicators, it is possible to detect and quantify the mutual influence between signals. The proposed approach has also been tested varying the noise of the signals and the number of data to perform the analysis, in order to provide a comprehensive assessment of the limits and potentialities of TDNNs.

2.9. Synthetic diagnostics, integration, verification and validation (SYD) – A. Dinklage

Synthetic diagnostics, integration, verification and validation: concise simulations of measurements allow one to explore the performance and efficacy of instruments. This makes synthetic diagnostics a way to determine how requirements on instruments are met in large devices. Widespread applications reported in different session indicate the use of forward-models as more and more established approach in fusion data analysis. Ultimately, synthetic diagnostics allow one to assess of machine access, time resolutions and many more aspects to benefits for the development and engineering of diagnostics. The session was dedicated to synthetic diagnostics and discussions aim at re-using developed virtual instruments for future devices.
Error estimations in filament measurements using a synthetic probe were discussed. Such probes are used in the scrape-off layer of magnetic confinement fusion plasmas to measure flows relevant to particle and power exhaust. The specific case reported were filaments measured on W7-X. Those were compared with drift-plane simulations. The specific issue arising in the comparison stems from details of the filament shape that requires conditional averaging. The synthetic instrument offers to mimic conditionally averaged measurement by deriving samples of measurements from different filament positions of single simulation rather than repeating large numbers of simulations. A detailed assessment of this trick indicated some underestimation of the filament size in the experimental measurements of about 20% for the simplified approach. A synthetic scaling to correct for these errors allowed one to conclude that good agreement between experiments and simulations are even better than previously reported. Taking these deviations into account, the advantage of the method lies in the substantially accelerated analysis allowing one for a tractable analysis of large data sets. The work provided valuable insight into the inherent errors with probe measurements of filaments.
An update of forward modelling modules in Bayesian analyses for W7-X was given. Vignetting of camera views in X-ray tomography systems were needed to correct for systematic errors. Some cross validation of electron and ion temperature and density, respectively, were conducted. The computation time of the large and complicated network models was significantly sped by artificial neural networks to be trained on synthetic and experimental data.
Ways to estimate head-load distributions from Monte-Carlo samples using edge transport simulations with a synthetic camera were shown. Purpose of the comparison

is to unravel difference of simulated heat fluxes to observed temperatures. This comparison may reveal the role of anisotropic diffusion since this process couples fluxes from different geometrical domains in the divertor plasma structure or captures effects from counter-streaming flows. It was shown that the inclusion of the full 3D plasma feeding the flows needs to be taken into account consistently. As a specific outcome, the synthetic diagnostic workflow may reveal the specific identification of sources for locally enhanced loads.
The development of the ITER synthetic reflectometry diagnostic was discussed. The status report of the development is an insightful example how a synthetic instrument is integrated into IMAS. The development workflow is driven by the requirements for reflectometry measurements on ITER such as ICH coupling, the characterization of the L-H transition or advanced control. The assessment of high-field side reflectometry HFS reflectometry signals employs predictive scenarios and settings as data sources representing the machine description. Outcomes are simulated signals. A next step will involve the extension to different plasma scenarios and a comparison with experimental data.

2.10. IMAS/IDA integration (IDA) - R. Fischer

In present and future fusion devices huge amounts of measurements coming from many diagnostic systems have to be analyzed. Analysis of these data aims to extract the maximum possible information from the available diagnostics for plasma control and machine safety as well as for physics studies. A multitude of heterogeneous diagnostics provides redundant and complementary information for a variety of plasma parameters. Frequently, the analysis of data from one diagnostic relies on parameter estimates from complementary diagnostics. A joint analysis of interdependent diagnostics benefits from the simultaneous availability of complementary information.
IDA in the framework of Bayesian probability theory provides a method for a coherent combination of measured data from heterogeneous diagnostics as well as prior and modelling information[39]. The method relies on numerically-robust forward modelling (FM) of measured data from given physical parameters, also known as synthetic diagnostic, and on a Bayesian quantification of statistical and systematic (modelling) uncertainties. The probabilistic combination of measurements from various diagnostics as well as prior information from physical treatments results in a probability distribution describing the information obtainable from the various diagnostics and modelling studies. The results benefit from the amount of information provided and from the interdependencies between the diagnostics and between the parameters.
Based on more than 20 years of experience in applying IDA to various tokamak and stellarator devices, various diagnostic combinations and various parameter sets, a new implementation of the IDA approach was triggered by a newly founded Integrated Data Analysis and Validation (IDAV) specialist working group (SWG) within the ITPA Diagnostics Topical Group. The primary goal is to provide a general data analysis code package compatible with any fusion device, which is modular with respect to the choice of diagnostics. The feature set includes a combination of low- and high-fidelity forward models, flexible parameterization, low- to high fidelity priors and modelling information. Essential to the code package are various methods to estimate parameters

together with their uncertainties. The code is written in modern python programming language. A first test example combining synthetic data from Thomson scattering and electron cyclotron emission (ECE) diagnostics with the forward modelled data from the PFPO-1 ITER Toroidal Interferometry Polarimeter (TIP) illustrated the implementation. The ITER example benefited from reading the Interface Data Structures (IDS) for the TIP geometry and for the ITER equilibrium from the ITER Integrated and Modelling & Analysis Suite (IMAS)[40]. IMAS provides via the IDS standardized access to experimental and simulated data, the full description of the tokamak subsystems (diagnostic, heating system, etc.), the physical concepts describing the plasma and synthetic diagnostics for ITER. The development of synthetic diagnostics for ITER is essential to optimise the design of the diagnostics by modelling their performance in various scenarios, to develop the necessary control algorithms utilising them, and to perform specific physics studies, including IDA, for each phase of the ITER Research Plan. The work involves the standardized approach of IMAS with the Plasma Control System Simulation Platform (PCSSP) focused on controlling the plasma behaviour and optimising its performance. Developing synthetic diagnostics using the IMAS Data Model ensures portability and a more flexible use within different workflows, as well as supporting better traceability and reproducibility of the data generated, providing a robust modelling procedure. Various requirements on the performance of each model depending on its application were shown. A common requirement of the synthetic diagnostics is that they have to follow the IMAS standard, i.e. they have to exchange IDSs exclusively as input and output. Using the IMAS standard has permitted the development of a workflow that can generate synthetic diagnostic data from ITER scenario simulations, following the same strategy as the IMAS workflow for Heating and Current Drive sources[41]. Several examples of IMAS synthetic diagnostic models developed for Interferometry, Refractometry, Bolometry, Neutron flux monitors, and Visible Spectroscopy were shown. Ultimately, these models will be combined in an integrated approach to data analysis to deliver a robust interpretation of ITER experimental data. Next steps for ITER applications were identified as, for example, the compatibility and numerical efficiency of the IDA workflow with the IMAS synthetic diagnostics.

Before the start of ITER operation and the availability of experimental data, synthetic diagnostics can be used to simulate measurements for given plasma parameters from predictive simulations and the configuration of each diagnostic system. For the identification of the L-H transition in ITER PFPO campaigns predictive simulations use advanced core and edge transport solvers like ASTRA[42], JINTRAC [43] and SOLPS-ITER [8] [44]which results are stored in the IMAS Scenario Simulations database. These scenarios together with the synthetic diagnostics using the IMAS Machine Description database are used to produce simulated data to study the detection of the L-H transition. The diagnostics encountered comprise the CASPER (Hα) workflow for the visible spectrometer camera, the interferometer/polarimeter synthetic diagnostic and the ECE synthetic diagnostic.

For the ECE synthetic diagnostic a sophisticated forward model solving the radiation transport was employed abandoning the classical interpretation using the standard

black-body assumption. This is essential for the optically thin pedestal region of current H-mode plasmas, for the much hotter plasmas in future machines like ITER, SPARC and DEMO, for low-density scenarios as well as oblique ECE measurements and harmonic overlap. In these situations, the kinetic broadening of the ECE due to the relativistic mass increase and the Doppler shift can no longer be neglected and radiation transport effects need to be included in the interpretation of the ECE measurements. This also inhibits the direct inference of the electron temperature T_e as the measurements are no longer localized. Additionally, the ECE T_e information is entangled with the electron density ne which is resolved by combining IDA with a radiation transport code like ECRad [45]. Predictive ECE spectra for ITER and SPARC illustrated the necessity for radiation transport modelling.
The reflectometer diagnostics are expected for ITER in PFPO-2. A new efficient FM for the swept density reflectometry diagnostic was proposed. In contrast to the error-prone Abel inversion approaches, the use of a reflectometry FM allows one to use redundant overlapping frequency bands and relaxes the influence of poor SNR data gaps. Additionally, the analysis of reflectometer measurements benefits from a combined analysis with other density diagnostics. An independent scrape-of-layer density diagnostic like the lithium beam resolves the initialization problem for densities below the lowest measured cut-off density. Where the Abel inversion on noisy data only provides cut-off positions with potential density ambiguity, a Bayesian approach provides unambiguous density profiles including uncertainty measures. An integrated workflow for energetic particles stability was developed within IMAS[46].
The time-dependent workflow solved problems with the centralization of data from different codes and demonstrated the orchestration of the retrieval and storage of IDSs as well as their passing between the physics actors involved, namely the equilibrium code HELENA [47] and the linear gyro-kinetic stability code LIGKA[48]. The workflow allows for an automatic, time-dependent, reproducible and consistent stability analysis based on documented input and output.
An IMAS data processing workflow implemented at WEST allows for obtaining reduced databases with quasi-stationary plasma states and with time-dependent plasma parameters from the integration of information coming from several diagnostics.

2.11. Big data (BIG) - J. Stillerman

The big data session of the meeting covered a wide range of topics. The data from magnetic fusion experiments is large and growing quickly. How the community can maximize the benefits of this potentially overwhelming quantity of data was discussed. FAIR, a set of principles dealing with the sharing and documenting of data, was addressed. There was also a tutorial on this subject at the beginning of this Technical Meeting. Long-pulse or continuous experiments need to support streaming data analysis. Large data sets can benefit from applying machine learning techniques to automatically analyze and classify data. Tracking data provenance is a critical part of data analysis.
Two talks in the session discussed the FAIR data principles. That is that data should be FAIR - Findable, Accessible, Interoperable, and Reusable. A general overview of the FAIR4Fusion project was presented including motivations, challenges, and benefits.

This framework for data sharing is guided by the principle “as open as possible, as closed as necessary”. It includes both technical and administrative elements which are both required for data to be shared and exploited by the community. Three projects related to FAIR were presented. The IMAS integrated modeling & analysis suite was presented as a common data interchange format. Code containerization (docker etc…) can be used to create multi-step repeatable data analysis chains. The MAST experiment has created a data portal to provide data to the wider community.
One of the four FAIR principles is R - Reusable. In order for data to be reusable, a good understanding of its provenance is required. Two presentations addressed documenting the provenance of computed results. One applied the non-domain specific W3C-PROV to data from MAST and WEST. Standardizing provenance representations enables the FAIR I - interoperable principle. Documenting the full life-cycle of data, including all of the codes which are used to produce a result was presented. While not called out by the authors as FAIR related, having the whole analysis pipeline in source code control with a CI/CD pipeline has the same motivations and results. The Japanese `Fusion Cloud` allows collaborating scientists to leverage distributed computing to analyze data from existing and future fusion experiments. Their application of globally unique identifiers enables researchers to refer to and cite the data used, which is another FAIR R - Reusable principle. In the US, a framework called DELTA is used to stream data between pulses to off-site supercomputers. This timely data analysis provides experimenters and session leaders with actionable results in between plasma pulses.
Synthesizing actionable results from large diverse data sets is a very common activity in the magnetic fusion research community. Two presentations discussed applying machine learning to aid this. One was based on Clustering (MiniBatchKMeans) and threshold techniques to 'clean-up' frequency spectra so that underlying Broadband turbulence could be studied. The other applied a bidirectional LSTM neural network to compare diagnostic signals with actuator inputs, thereby classifying the discharge as 'normal' or 'off-normal'.

3. Round table discussion

A round-table discussion was held on the last day of the meeting, covering several topics that were treated or introduced at the workshop. The aim was to stimulate discussion toward future directions that the field should explore, driven by the needs that are currently perceived in this domain. Chaired by D. Mazon (CEA, France) and M. Xu (SWIP, China), the discussion, dealing with both methods and applications, led to several concrete proposals for the next edition of the workshop and more generally for development of fusion data processing, validation and analysis in the broader fusion community.
Before proceeding to the specific topics brought up during the discussion, it is interesting to look back on the recommendations that were made during a similar discussion at the end of the previous meeting in 2019. This reveals that good progress has been made on several fronts. For instance, there is a strong effort led by the ITER Organization for developing and maintaining synthetic diagnostics, which is essential for many other activities. Among those is the joint analysis of data from multiple

diagnostics within the framework of Bayesian inference ("data fusion"), which needs diagnostic forward models for likelihood distributions. In fact, one of the main outcomes from the previous meeting was the strong message and intention to organize joint activities of data fusion using Bayesian methods for ITER diagnostics. These techniques are sometimes referred to IDA in the fusion community, a term that will be used here as well. Since the last meeting, this recommendation has led to the start of a new Specialist Working Group on Integrated Data Analysis and Validation (SWG IDAV) in the framework of the ITPA Topical Group on Diagnostics. Concretely, two implementations of IDA are presently being considered at ITER, with contributions from various institutions.

On the other hand, the consolidation of expertise and standardization of tools for data analysis and validation remains an important point of attention, as was noted already in 2019. It is essential for methods and software to be benchmarked, and for the most promising ones to be transferred and adopted by the main (future) fusion devices around the world. The FDPVA community can play an important role in this regard, by organizing dedicated sessions at the biennial meetings and by joining various groups and communities (also from outside fusion) working on similar applications. This perspective was a connecting thread throughout the round-table discussion.

3.1. Anomaly detection

Detection of off-normal events, or anomalies, is an area where statistics and machine learning can provide a major benefit for fusion, particularly in the real-time setting. One of the main current applications is disruption prediction, but others are taking off as well, aimed at the detection of growth of plasma instabilities, hot spots on plasma-facing components or, more generally, condition monitoring of machine components.

From the discussion emerged the view that future editions of the workshop could stress more clearly the distinction between techniques that can be applied in real time, and those that presently are too computationally intensive for that. In particular, the potential of permutation entropy was noted as a computationally light-weight means to signal upcoming disruptions from time series data. It is a nonparametric technique that is based on sorting algorithms, which is amenable to real-time implementation (e.g. on FPGAs).

Furthermore, anomaly detection so far has mainly relied on time series data, whereas more recently space-resolved (profile) data have been exploited. This aspect could also be represented more prominently at future meetings.

Finally, as was noted at the previous meeting, the need remains for a concerted benchmarking of disruption prediction tools. The prerequisite is a multimachine database that links to the original time series or profile data, which presently does not exist. Tools for automated database building could contribute to this goal. In addition, a contest in the style of e.g. the Kaggle competitions [49] could be organized at a future meeting.

3.2. Surrogate modelling and reduced models

Surrogate modelling of complex fusion codes, notably using neural network models, is being increasingly widely adopted by the fusion community, as reflected by various presentations at the workshop. These methods can also drastically lower the

computational needs of Bayesian inference, by emulating a computationally demanding forward model or even the entire inference process.
One recent application is the extraction of reduced models from a database of plasma simulations (e.g. using SOLPS), by means of system identification tools, aimed at model predictive control. A challenge is to do this for time-dependent dynamics, rather than steady-state plasmas, but nevertheless the approach could offer substantial advantages in certain scenarios where speeding up code might be difficult.

3.3. Image processing

There is an ever growing trend in the area of fusion diagnostics to make use of cameras for machine protection, but also for physics analysis. Accordingly, there is a strong need to increase the introduction and development of (automated) image processing techniques. There is a lot to be gained by employing common tools and by sharing expertise, therefore these activities need to be well represented at the workshop. Expertise and tools developed in other scientific disciplines could be very useful in fusion as well. For instance, experts in processing and analysis of astronomical images (e.g. infrared) could be invited as speakers at the FDPVA workshop.

3.4. Equilibrium reconstruction

An important trend in magnetic equilibrium reconstruction is to accelerate methods for real-time application. Again, machine learning can play a role here, although some implementations are already real-time capable. Therefore, a session devoted specifically to the comparison of methods for real-time equilibrium reconstruction could be held at a future edition of the workshop.

3.5. Plasma simulators

Over the past few years the fusion community has seen the development of several plasma simulators (“flight simulators”) for preparing discharge scenarios and for plasma control purposes. An overview and comparison of the main efforts in this direction would be very interesting for the next meeting. The FDPVA community could contribute to the integration of reduced or emulated models into these plasma simulators, and validation against experimental data. In this context, the presentation at the past meeting of work from outside the fusion community regarding digital twins was strongly appreciated. This is an excellent example of how well-established methods and expertise from other disciplines could contribute to developments in data analysis and validation in fusion.

3.6. Synthetic diagnostics and IMAS

As mentioned in the introduction, great progress is being made by the ITER Organization in developing and maintaining synthetic diagnostics for ITER. Nevertheless, validation at other machines, preferably within the context of IMAS, is certainly necessary. This could be facilitated by incorporating additional data processing routines into IMAS, like more advanced interpolation routines. For instance, it could be investigated whether porting routines from OMAS to IMAS is feasible, all the while using the IMAS data dictionary.

3.7. Standardization, consolidation of methods and tools

Standardization of methods and tools is essential for comparison of results across experiments and devices. The FDPVA is particularly involved in estimation and propagation of uncertainties on measurements and simulations, with the aim of assessing measurement quality and for benchmarking purposes. Promotion by the FDPVA of a standard for expressing the uncertainty in measurements and codes is considered to stimulate good practice regarding definition and handling of uncertainties in the fusion community. An international standard for evaluating and expressing measurement uncertainty has been documented in the "Guide to the expression of uncertainty in measurement" (GUM) by the Joint Committee for Guides in Metrology (GUM08)[50]. It is proposed to invite at the next FDPVA meeting a metrology expert to provide an up-to-date review of the recommendations made in the GUM document. Consolidation of data analysis methods and tools developed in fusion is another key objective of the FDPVA. This is urgently needed both for ITER and for cross-device application, taking into account the right priorities (see e.g. Loarte20[51]). As was already mentioned above in the context of disruption prediction, it requires benchmarking using a standardized data set. A similar activity is already being maintained by a subgroup of the SWG IDAV in the ITPA Topical Group on Diagnostics. Another area where this would be very useful is in the development and application of synthetic diagnostics.
Furthermore, for the next edition it is proposed to organize a session on interpretation of uncertainties, with particular attention to approximation of uncertainties, e.g. for real-time purposes or strongly ill-posed problems.

3.8. Conclusions of the round – table session

The round-table discussion held at the 4th IAEA TM FDPVA can be summarized by the following conclusions and list of actions:

- Continue to push for initiatives aimed at standardization, benchmarking and transferal of methods and tools from the FDPVA field to ITER and additional fusion devices.
- Organize sessions at the next edition of the workshop on equilibrium reconstruction, on plasma simulators, and on interpretation and approximation of uncertainties.
- Invite speakers at the next edition from the field of metrology for uncertainty quantification and from astronomy for image processing.
- Organize a "challenge" at the next meeting to stimulate benchmarking of e.g. disruption predictors, synthetic diagnostics or plasma simulators.

4. Conclusions

The 4th IAEA Technical Meeting on Fusion Data Processing, Validation and Analysis (Online, November 29 to December 6, 2021) was greatly acknowledged by the community, being the reference event in the field of fusion. With the growth of massive measurement systems and data for future fusion reactors, data analysis is moving towards faster, more systematic and smarter directions. The recent highlights and progress in the following 12 sessions have been briefly summarized in this report.
The next edition i.e. the 5th IAEA Technical Meeting on Fusion Data Processing,

Validation and Analysis is expected to take place in 2023 in China being chaired by M. Xu and D. Mazon. More innovative and groundbreaking achievements in this field will be reported in this next edition.

Acknowledgements

The authors would like to acknowledge contribution from Shaun DE WITT, Shinya MAEYAMA, Azarakhsh JALALVAND, Jonathan HOLLOCOMBE, Didier VEZINET, Alexei PANKIN, Marcin PLOCIENNIK, Jinxin WANG, Juan GUERRERO ARNAIZ, Sebastian DE PASCUALE, Jarrod LEDDY, Vignesh GOPAKUMAR, Karel VAN DE PLASSCHE, Emi NARITA, James BUCHANAN, Takashi NISHIZAWA, Francesco SCIORTINO, Tianbo WANG, Yuya MORISHITA, Lynton APPEL, Randy CHURCHILL, Joseph HALL, Peizheng ZHANG, Masayuki YOKOYAMA, Michela GELFUSA, Riccardo ROSSI, David SMITH, Dhaval GADARIYA, Zongyu YANG, Diogo R. FERREIRA, Chengshuo SHEN, Raphael MITTEAU, Erwan GRELIER, Aleix PUIG SITJES, Mark BOYER, Francisco MATOS, Tom F. NEISER, Lian WANG, Lakshya MALHOTRA, Fengming XUE, Andrea PAVONE, Liqing XU, Yiming ZU, Lang LAO, Marie-Hélène AUMEUNIER, Ting LAN, Zijie LIU, Teddy CRACIUNESCU, Luca SPOLLADORE, Emmanuele PELUSO, Brendan SHANAHAN, Andreas LANGENBERG, Alexander KNIEPS, Valentina NIKOLAEVA, Mireille SCHNEIDER, Severin DENK, Dirk STIEGLITZ, Anna MEDVEDEVA, Alin POPA, Jorge MORALES, Ralph KUBE, Hideya NAKANISHI, Michal OWSIAK, Stasinos KONSTANTOPOULOS, Pär STRAND, Luigui SALAZAR, Chenguang WAN, Nathan CUMMINGS, Xiaojuan LIU, W. JIA and X. XIONG for the 4th IAEA Technical Meeting on Fusion Data Processing, Validation and Analysis.

References

[1] Jalalvand A, Kaptanoglu A A, Garcia A V, et al. Alfvén eigenmode classification based on ECE diagnostics at DIII-D using deep recurrent neural networks. Nuclear Fusion, 2021, 62(2): 026007.

[2] Wang J, Wang N, Liu Z, et al. Breakdown prediction of ion source at EAST-NBI using neural network. Plasma Physics and Controlled Fusion, 2021, 63(12): 125014.

[3] https://tofuproject.github.io/tofu/index.html

[4] Svoboda J, Cavalier J, Ficker O, et al. Tomotok: python package for tomography of tokamak plasma radiation. Journal of Instrumentation, 2021, 16(12): C12015.

[5] https://github.com/smaeyama/mzprojection

[6] https://transp.pppl.gov

[7] https://www.fair4fusion.eu

[8] Bonnin X, Dekeyser W, Pitts R, et al. Presentation of the new SOLPS-ITER code package for tokamak plasma edge modelling. Plasma and Fusion Research, 2016, 11: 1403102-1403102.

[9] Kaveeva E, Rozhansky V, Senichenkov I, et al. Speed-up of SOLPS-ITER code for tokamak edge modeling. Nuclear Fusion, 2018, 58(12): 126018.

[10] Stangeby P C. Basic physical processes and reduced models for plasma detachment. Plasma Physics and Controlled Fusion, 2018, 60(4): 044022.
[11] Brunton S L, Proctor J L, Kutz J N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the national academy of sciences, 2016, 113(15): 3932-3937.
[12] Sciortino F, Howard N, Reksoatmodjo R, Foster A, Hughes J, Marmar E, Miller M, Mordijck S, Odstrčil T, Pütterich T. Experimental inference of neutral and impurity transport in Alcator C-Mod using high-resolution x-ray and ultra-violet spectra. Nuclear Fusion, 2021, 61(12): 126060.
[13] Sciortino F, Howard N T, Marmar E S, Odstrcil T, Cao N M, Dux R, Hubbard A E, Hughes J W, Irby J H, Marzouk Y M, Milanese L M, Reinke M L, Rice J E, Rodriguez-Fernandez P. Inference of experimental radial impurity transport on Alcator C-Mod: Bayesian parameter estimation and model selection. Nuclear Fusion, 2020, 60(12): 126014.
[14] Pan W a, Wang T, Wang Z, Yang Y, Wu H, Geert V, Yang Z, Liu C, Guo W, Li B, Liu Z, Wu W, Xu M. Integrated Data Analysis on the Electron Temperature profile of HL-2A with Bayesian probability inference Method. Plasma Science and Technology, 2022, 24 055601
[15] Cranmer K, Brehmer J, Louppe G. The frontier of simulation-based inference [J]. Proceedings of the National Academy of Sciences, 2020, 117(48): 30055-62.
[16] Morishita Y, Murakami S, Yokoyama M, Ueno G. ASTI: Data Assimilation System for Particle and Heat Transport in Toroidal Plasmas. Computer Physics Communications, 2022, 108287.
[17] Verdoolaege G, Kaye S M, Angioni C, Kardaun O J, Maslov M, Romanelli M, Ryter F, Thomsen K, Contributors J, EUROfusion MST1 Team t. The updated ITPA global H-mode confinement database: description and analysis. Nuclear Fusion, 2021, 61(7): 076006.
[18] Yokoyama M, Yamaguchi H. Progress of statistical modelling of thermal transport of fusion plasmas. Nuclear Fusion, 2020, 60(10): 106024.
[19] Murari A, Vagliasindi G, Arena P, et al. Prototype of an adaptive disruption predictor for JET based on fuzzy logic and regression trees. Nuclear Fusion, 2008, 48(3): 035010.
[20] Murari A, Lungaroni M, Peluso E, et al. Adaptive predictors based on probabilistic SVM for real time disruption mitigation on JET. Nuclear Fusion, 2018, 58(5): 056002.
[21] Murari A, Lungaroni M, Gelfusa M, et al. Adaptive learning for disruption prediction in non-stationary conditions. Nuclear Fusion, 2019, 59(8): 086037.
[22] Murari A, Rossi R, Peluso E, et al. On the transfer of adaptive predictors between different devices for both mitigation and prevention of disruptions. Nuclear Fusion, 2020, 60(5): 056003.
[23] Murari A, Rossi R, Lungaroni M, et al. Stacking of predictors for the automatic classification of disruption types to optimize the control logic. Nuclear Fusion, 2021, 61(3): 036027.
[24] Dormido-Canto S, Vega J, Ramírez J M, et al. Development of an efficient real-

time disruption predictor from scratch on JET and implications for ITER. Nuclear Fusion, 2013, 53(11): 113001.
[25] Vega J, Murari A, Dormido-Canto S, et al. Adaptive high learning rate probabilistic disruption predictors from scratch for the next generation of tokamaks. Nuclear Fusion, 2014, 54(12): 123001.
[26] Murari A, Rossi R, Peluso E, et al. On the transfer of adaptive predictors between different devices for both mitigation and prevention of disruptions. Nuclear Fusion, 2020, 60(5): 056003.
[27] Paméla J, Romanelli F, Watkins M L, et al. The JET programme in support of ITER. Fusion Engineering and Design, 2007, 82(5-14): 590-602.
[28] Romanelli F, Laxåback M. Overview of JET results. Nuclear fusion, 2011, 51(9): 094008.
[29] Murari A, Peluso E, Lungaroni M, et al. Investigating the physics of tokamak global stability with interpretable machine learning tools. Applied Sciences, 2020, 10(19): 6683.
[30] Murari A, Boutot P, Vega J, et al. Clustering based on the geodesic distance on Gaussian manifolds for the automatic classification of disruptions. Nuclear Fusion, 2013, 53(3): 033006.
[31] Vega J, Murari A, Dormido-Canto S, et al. A linear equation based on signal increments to predict disruptive behaviours and the time to disruption on JET. Nuclear Fusion, 2019, 60(2): 026001.
[32] Rattá G A, Vega J, Murari A, et al. PHAD: a phase-oriented disruption prediction strategy for avoidance, prevention, and mitigation in JET. Nuclear Fusion, 2021, 61(11): 116055.
[33] Yang Z, Xia F, Song X, et al. In-depth research on the interpretable disruption predictor in HL-2A. Nuclear Fusion, 2021, 61(12): 126042.
[34] Vega J, Castro R, Dormido-Canto S, et al. Automatic recognition of plasma relevant events: Implications for ITER. Fusion Engineering and Design, 2020, 156: 111638.
[35] Murari A, Craciunescu T, Peluso E, et al. Detection of causal relations in time series affected by noise in tokamaks using geodesic distance on gaussian manifolds. Entropy, 2017, 19(10): 569.
[36] Craciunescu T, Murari A, Lerche E, et al. Image-based methods to investigate synchronization between time series relevant for plasma fusion diagnostics. Entropy, 2020, 22(7): 775.
[37] Murari A, Craciunescu T, Peluso E, et al. How to assess the efficiency of synchronization experiments in tokamaks. Nuclear Fusion, 2016, 56(7): 076008.
[38] Murari A, Craciunescu T, Peluso E, et al. On efficiency and interpretation of sawteeth pacing with on-axis ICRH modulation in JET. Nuclear Fusion, 2017, 57(12): 126057.
[39] Fischer R, Fuchs C J, Kurzan B, et al. Integrated data analysis of profile diagnostics at ASDEX Upgrade. Fusion science and technology, 2010, 58(2): 675-684.
[40] Imbeaux F, Pinches S D, Lister J B, et al. Design and first applications of the ITER integrated modelling & analysis suite. Nuclear Fusion, 2015, 55(12): 123006.

[41] Schneider M, Lerche E, Van Eester D, et al. Simulation of heating and current drive sources for scenarios of the ITER research plan. Nuclear Fusion, 2021, 61(12): 126058.
[42] Pereverzev G V, Yushmanov P N. ASTRA. Automated system for transport analysis in a tokamak. 2002. (https://w3.pppl.gov/hammett/work/2009/Astra_ocr.pdf)
[43] Romanelli M, Corrigan G, Parail V, et al. JINTRAC: a system of codes for integrated simulation of tokamak scenarios. Plasma and Fusion research, 2014, 9: 3403023-3403023.
[44] Wiesen S, Reiter D, Kotov V, et al. The new SOLPS-ITER code package. Journal of nuclear materials, 2015, 463: 480-484.
[45] Denk S S, Fischer R, Poli E, et al. ECRad: An electron cyclotron radiation transport solver for advanced data analysis in thermal and non-thermal fusion plasmas. Computer Physics Communications, 2020, 253: 107175.
[46] V. A. Popa, Max-Planck-Institut für Plasmaphysik, Garching (2021) IPP 2021-18
[47] Huysmans G T A, Goedbloed J P, Kerner, W, Isoparametric Bicubic Hermite Elements for Solution of the Grad-Shafranov Equation, International Journal of Modern Physics C Vol. 02, No. 01, pp. 371-376 (1991)
[48] Lauber P, Günter S, Pinches S D. Kinetic properties of shear Alfvén eigenmodes in tokamak plasmas[J]. Physics of plasmas, 2005, 12(12): 122501.
[49] Kaggle competitions [M]. URL: https://www.kaggle.com/competitions
[50] Evaluation of measurement data — Guide to the expression of uncertainty in measurement [M]. Joint Committee for Guides and Metrology. 2008. URL: https://www.bipm.org/utils/common/documents/jcgm/JCGM_100_2008_E.pdf
[51] Loarte A. Required R&D in Existing Fusion Facilities to Support the ITER Research Plan [M]. ITER Technical Report ITR-20-008. 2020. URL: https://www.iter.org/doc/www/content/com/Lists/ITER%20Technical%20Reports/Attachments/14/ITR_20_008_Required_RD_in_existing_fusion_facilities_to_support_the_ITER_Research_Plan.pdf