DISRUPTION PREVENTION VIA INTERPRETABLE DATA-DRIVEN ALGORITHMS ON DIII-D AND EAST

C. REA, K.J. MONTES, R.S. GRANETZ
Massachusetts Institute of Technology, Plasma Science and Fusion Center
Cambridge, MA, USA
Email: crea@mit.edu

W. HU, Q.P. YUAN, D.L. CHEN, B. SHEN, B.J. XIAO
Institute of Plasma Physics, Chinese Academy of Sciences
Hefei, Anhui, China

J.L. BARR, B. SAMMULI
General Atomics
San Diego, CA, USA

K.G. ERICKSON
Princeton Plasma Physics Laboratory
Princeton, NJ, USA

Abstract

The Disruption Prediction via Random Forest (DPRF) algorithm is currently implemented in both DIII-D and EAST PCS, and provides not only predictions of impending disruptions in real-time, but simultaneously identifies the drivers of the disruptivity, i.e. feature contributions – all in about 150 - 250 microseconds. This is the first demonstration that a machine learning-based algorithm can provide interpretable predictions in real-time on multiple devices. Providing indications of the ongoing disruption dynamics to the Plasma Control System (PCS) proves to be essential in actuating the proper response to avoid performance loss and disruptions deleterious consequences. On DIII-D, DPRF was upgraded including real-time calculations of profile-based indicators of temperature, density and radiation. Such peaking factors prove to be relevant metrics in impurity accumulation events leading to disruptions in scenarios close to ITER baseline, providing a warning more than 1s prior to disruption. On EAST, DPRF was trained on high density disruptions, and during closed-loop experiments, it has shown to be capable of predicting such cases, and trigger the mitigation system with relevant accuracy. Data-driven models will be integral part of device protection system for ITER and the next generation of devices, but only if capable to demonstrate optimal predictive capabilities and the possibility to reconcile the predictions with the underlying physics dynamics. Current results represent a step forward in data-driven control for scenario optimization and disruption avoidance for ITER and next generation devices. This work establishes the importance of developing tools capable of identifying and informing in real-time the PCS on the dangerous plasma parameters deviations to the disruptive space to enable the proper actuators' response.

1. INTRODUCTION

A complete understanding of pre-disruptive and disruptive phenomena in tokamak plasmas still represent a challenge to the fusion community and hinders the path to a viable solution for a nuclear reactor. Among the tasks that tokamak scientists need to address one should mention the need to develop, qualify, and validate strategies for mitigation and avoidance of disruptions that could guarantee uninterrupted reactor operations, by predicting, detecting and adequately mitigating 100% of disruptions. The community is actively working around these issues and trying to push forward both understanding and possible solutions, also thanks to the aid of machine learning methodologies. The need to carry out disruption research in all its aspects, from predictive capabilities to mitigation techniques, finds support in the recent US DPP-CPP strategic plan [CPP2020], where it is explicitly recommended a strong focus on the development and validation of specific metrics for Disruption Prediction, Avoidance, and Mitigation (DPAM), by exploiting existing tokamaks as transitioning devices to ITER. This is necessary to gain confidence and experience in burning plasma physics and its control, and to qualify tokamak devices as viable solutions for a Fusion Pilot Plant.

In this general context, we have worked to develop a portable and interpretable data-driven algorithm for disruption prediction, and this was included as part of the real-time algorithms suite in DIII-D and EAST Plasma Control Systems (PCS). The Disruption Prediction via Random Forest (DPRF) algorithm guarantees explainable predictions in real-time thanks to the feature contribution analysis [REA2019], which is able to identify the main drivers of the increasing risk of an impending disruption, i.e. disruptivity. DPRF was tested in closed-loop

mitigation experiments on both DIII-D and EAST tokamaks. On DIII-D, indicators of plasma temperature, density and radiation based on profile peaking factors [PAU2018, REA2020] were successfully implemented in real-time and used as input features for DPRF during recent experiments. On EAST, DPRF was trained on a set of disruptions in which the plasma electron density reached at least the 80% of the Greenwald density limit, and in post-mortem analysis it has shown a successful alarm rate of 92% and a false alarm rate of 9.9%.

This manuscript is organized as follows: in Section 2 we will describe the DPRF algorithm in some detail, together with the input features used for the real-time applications on both devices. As mentioned above, DPRF provides explainable predictions, therefore a short paragraph will be dedicated to summarizing the technique adopted to interpret the disruptivity, i.e. the feature contribution analysis. In Section 3 we discuss the dedicated real-time experiments on DIII-D, together with various peculiarities of the real-time implementation, while Section 4 summarizes the results obtained on EAST. Finally, a closing summary together with a description of the next steps are reported in Section 5.

2. DPRF ALGORITHM, INPUT FEATURES, AND INTERPRETABILITY

DPRF is based on the Random Forest machine learning algorithm [BRE2001], which estimates the probability of an impending disruption by developing a large number (typically hundreds) of independent, de-correlated base learners (decision tree models), thus collecting a parallel set of predictions. The final prediction from the ensemble is aggregated by averaging this very large number of models' predictions. A calibrated threshold on the disruptivity is then found to maximize the success rate and warning time. Technical details on how this procedure is implemented can be found in previous publications by the authors [MON2019]. Briefly, a multi-dimensional grid-search is carried out, while maximizing a chosen performance metric: in this case, a shot-by-shot accuracy on detecting disruptive sequences weighted by the number of misclassified samples[1]. EAST and DIII-D DPRF real-time implementations have different optimal hyper-parameters in terms of number of trees and layers. The Out-Of-Bag error during training was found to plateau to a minimum [BRE2001, REA2019] at 375 trees and a depth of 15 layers for the DIII-D dataset, and at 920 decision trees with 15 layers depth on EAST. The different model hyper-parameters reflect the differences in dataset size and disruption dynamics.

While it is necessary to provide enough warning time to predict an impending disruption, it is crucial to also diagnose what type of event in the disruptive chain is causing the deviation from the safe operational space. This allows the PCS to respond adequately to different instabilities by directing it to trigger different actuators. In this regard, 1D/2D profiles of plasma temperature, density and radiation contain highly predictive information that can be used for example to diagnose impurity accumulation events leading to radiative collapse and therefore disruptions. Reducing the plasma profiles to their peaking factors -- following what documented in [PAU2018, REA2020] -- handles the curse of dimensionality and represents a manageable and real-time compatible way for a data-driven algorithm to augment predictive capabilities. By incorporating these scalar features in DIII-D DPRF we aim at increasing the sensitivity of the model towards certain pre-disruptive scenarios, such as impurity accumulation events or unplanned H-to-L back-transitions.

On EAST, DPRF was tailored to predict high density disruptive scenarios and closed-loop performances were tested during recent mitigation experiments in 2020, where several disruptions were effectively mitigated via Massive Gas Injection (MGI) of Neon. The developed algorithm still presents some residual flaws as it shows to be prone to false alarms, e.g. when H-to-L back-transitions occur during the discharge. More work will also be required to implement definitions of profile peaking factors analogous to DIII-D but synthesized from EAST diagnostic systems, which are crucial to define a device-independent framework for properly diagnosing disruption precursors.

In Table 1, a list of the plasma descriptors available in real-time is provided. DIII-D and EAST share most (7) of the input features for DPRF, mostly taken from Rogowski loops, the interferometer and the available equilibrium reconstruction code. On EAST, the ParaEquilibrium (PEFIT) code reconstructs the plasma state in real-time via parallelized GPU calculations [HUA2016, HUA2020]: PEFIT takes approximately 500 μs to complete one equilibrium iteration, parameter calculation and data transfer. On DIII-D, rt-EFIT [FER1998] reconstructs an approximate solution to the Grad-Shafranov equilibrium by matching magnetic sensors and loop voltage measurements through iterated computations: DPRF is being sent all the quantities of interest at the end of the slow loop of rt-EFIT, every 1 ms.

---

[1] In particular, we implement $F_\gamma = \frac{(1+\gamma^2)TP}{(1+\gamma^2)TP+\gamma^2 FN+FP}$ score with γ=2. TP are the correctly classified disruptive samples, FN the misclassification on the disruptive classes, and FP the non-disruptive data wrongly classified as disruptive.

As mentioned above, one of the major differences between the available inputs consists in the four profile peaking factors that are available in real-time for DIII-D: Te and ne peaking are obtained from the real-time Thomson scattering system, by combining the core and the tangential system views to increase resolution in the core plasma region. Prad Core and X-Div peakings come from the real-time bolometer array: since a limited amount of channels are digitized in real-time, only a combination from the lower fan was chosen to ensure good coverage of lower single null plasma configurations. An accurate description of how these zero-dimensional features were engineered for DIII-D can be found in [REA2020].

TABLE 1. DPRF REAL-TIME INPUT SIGNALS AND THEIR DESCRIPTION. THE LAST COLUMN INDICATES WHETHER THE PLASMA FEATURE WAS USED ON DIII-D, EAST, OR BOTH DEVICES

Signal	Description	Available
Ip_error = (Ip-Ip_target)/Ip_target	Fractional error on the plasma current	BOTH
Vloop	Plasma loop voltage [V]	BOTH
n/nG	Greenwald density fraction	BOTH
Wmhd	Plasma stored energy [J]	BOTH
βn	Normalized ratio of plasma pressure to magnetic pressure	BOTH
κ	Plasma elongation	BOTH
li	Normalized internal inductance	BOTH
z_error = (z- z_target)/a	Error on the current centroid, normalized by plasma minor radius	EAST
q95	Safety factor at 95% flux surface	EAST
Br(n=1)/Bϕ	Non-rotating component of n=1 radial field, normalized to toroidal field	DIII-D
δ	Plasma triangularity	DIII-D
ζ	Plasma squareness	DIII-D
Te peaking	Peaking factor of electron temperature profile in the core	DIII-D
ne peaking	Peaking factor of electron density profile in the core	DIII-D
Prad Core peaking	Peaking factor of radiation profile in the core	DIII-D
Prad X-div peaking	Peaking factor of radiation profile in the divertor	DIII-D

One of the main advantages of using an ensemble of tree-based learners is the possibility to access the relevant information contained in the numerous decision paths that are developed during the training process [PAL2014, REA2019]. This characteristics guarantees the extraction of average interpretability metrics called feature contributions. The final disruptivity that is obtained as an average frequentist probability from the trained Random Forest can be deconstructed into two terms, the model bias and the contributions: $disruptivity = \frac{1}{M}\sum_{m=1}^{M} bias_m + \sum_{m=1}^{M}\left(\frac{1}{M}\sum_{n=1}^{N} contrib_m(n)\right)$, where M is the total number of trees in the forest and N is the total number of input features. Both the bias and the contributions are intrinsically related to the training dataset composition: the former reflects the actual discrepancy between the initial populations of the two classes (disruptive vs non-disruptive samples), while the latter are obtained by evaluating the variation in the sample population in each class, after each statistical test in the construction of the individual decision paths during training. For more information and examples please refer to [REA2019].

3. DIII-D DPRF IMPLEMENTATION AND EXPERIMENTS

Accurate predictions of impending disruptions are essential to a robust DPAM strategy, in particular data-driven or physics-based approaches (e.g., threshold on mode locking amplitude) can be used to trigger as last resort the mitigation system. Nevertheless, if the aim is to optimize plasma performances and actively avoid the occurrence of instabilities in the plasma, a continuous monitoring of relevant plasma features is needed and this also leads to the development of scenario-dependent event probabilities. If the predictive algorithms foresees the adoption of supervised machine learning techniques, the most important resource consists in the availability of

labelled pre-disruptive events. This entails quite a cumbersome effort, since the metadata needed to qualify the disruption causes lacks broadly in experimentally datasets and is usually obtained by employing an extraordinary amount of time from a domain expert to validate all these cases. To this aim, efforts exist to adopt physics-based or semi-supervised approaches to accelerate the automated identification of disruption causes [SAB2018, MON2021].

DPRF training for the DIII-D application relied on a set of 310 unstable events manually identified and validated, where the transition time into an unstable operational space was recorded for the individual disruptive discharges. Additionally, DPRF is a binary classifier: therefore individual time samples from each plasma discharge are labelled as either non-disruptive or unstable. To complement the ~ 300 disruptive discharges, 983 non-disruptive ones were also randomly selected and included in the training set. Figure 1 shows for the unstable samples the event type counts at the recorded transition time: as expected, MHD-driven disruptions (e.g., rotating modes that can eventually lead to locking) are the most frequent root causes of instability on DIII-D. Empirically it was also found that for such supervised machine learning-based algorithm the composition of the training set can influence the sensitivity of the model towards certain scenarios. Therefore, a particular effort was dedicated to diversify such scenarios and include impurity accumulation events (IMP), UFOs and MARFEs to differentiate the disruption causes. The analysed DIII-D data was obtained from the experimental campaigns spanning 2015 and 2016: during the latter a ring of tungsten tiles was installed in the lower divertor area for material migration studies [BYK2017]. In addition, the prevalence of MHD-driven disruption can be explained by the inclusion of 2015 dedicated experiments on rotating tearing mode control and magnetic islands formation.

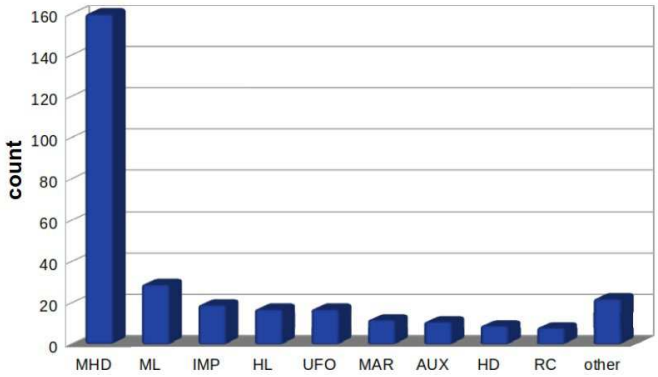

Fig. 1 Event type at transition time. Tags for different events are taken from [DEV2011]. As expected, MHD events leading eventually to mode locking (ML) on DIII-D dominate the disruptive chain across the more than 300 disruptions analyzed.

The combination of a more refined training set that includes different disruption causes with input features tailored to capture these particular events, such as the profile peaking factors, aims at increasing the model's sensitivity to less frequent disruption dynamics on DIII-D. To test these assumptions, actual experiments were performed on DIII-D, where DPRF was tested in ITER Baseline Scenario (IBS) plasmas [DOY2010]: closed-loop performances were verified by successfully integrating the disruptivity alarm in the Off Normal Fault Response [EID2018] algorithm for emergency response. In different discharges, a pre-programmed threshold on the disruptivity signal triggered (1) the Massive Gas Injection (MGI) in response to dangerous tearing mode activity, (2) the Electron Cyclotron Heating (ECH) system aimed at q=3/2 surface when such precursor emerged, and (3) an early shutdown via fast ramp down also due to the presence of a tearing mode (see Fig. 2). Additionally, the peaking factors were verified to represent a good metric for impurity accumulation events, when probing via Argon injection: in discharge 180808 at 2 seconds into the flattop, one of the gas valves at the bottom of the chamber opens to inject Argon. The disruptivity (Fig 3, first panel) responds on impurity accumulation time scales, as also reflected in the middle panel by the core radiation (Prad PF) and density (ne PF) peaking factors. The bottom panel in Fig. 3 shows the feature contributions calculated in real-time: the core radiation peaking factor (Prad PF) results as the most important driver of the disruptivity increase. An important detail about these discharges is that the Prad X-div peaking feature was unavailable, as the experimental session had to be run in upper single null plasmas to reach conditions that would match IBS plasmas [TUR2018] and favour H-mode transition during the flattop phase.

Further experiments are scheduled during the 2021 experimental campaign to demonstrate the importance of such metric in impurity-driven disruptions.

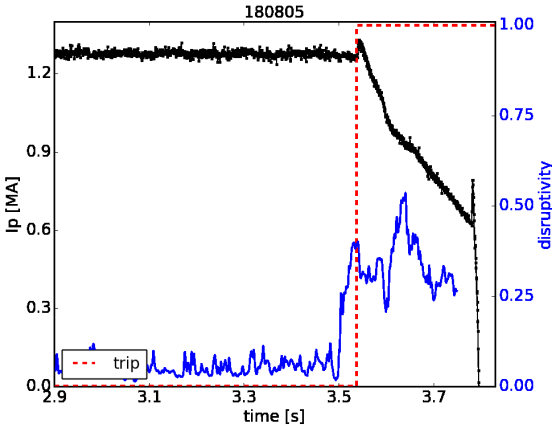

Fig. 2 In DIII-D discharge 180805 DPRF triggers an early shutdown by means of fast ramp-down in response to the presence of a tearing mode. Plasma disrupts anyway but at halved plasma current value, thus strongly reducing the stored energy that gets dissipated. The red dashed lines indicates the trigger time (t ~ 3.6 s) at the preset disruptivity threshold used by the PCS.

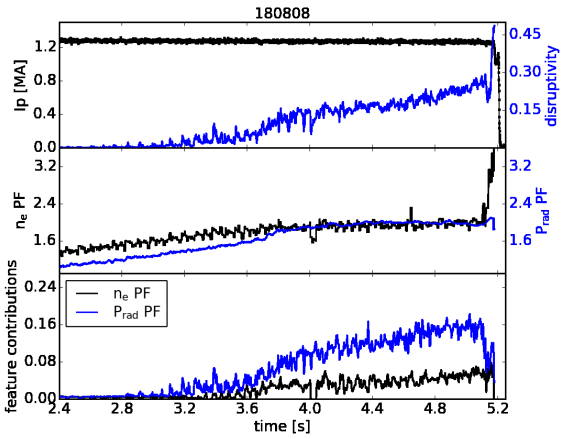

Fig. 3 ITER Baseline discharge in DIII-D, at 2s gas valve opens to inject Argon. The disruptivity responds on impurity accumulation time scales, as reflected in the middle panel by radiation (P_{rad} PF) and density (n_e PF) peaking factors. The bottom panel shows the feature contributions calculated in real-time: P_{rad} PF is the most important driver of the disruptivity increase, less relevant features not shown here.

3.1 Improvements to DPRF for real-time implementation on DIII-D

A further difference between DIII-D and EAST DPRF consists in the strict real-time implementation: EAST DPRF is trained and optimized offline using off-the-shelf routines from scikit-learn [PED2012], then the trained algorithm is translated to C through an open source library[2] and the contributions are precomputed in the final C version following what documented in https://github.com/andosa/treeinterpreter. The C version is then compiled and linked in the PCS as an external library and enabled via a dedicated category connected to the 'Alarm' system.

On DIII-D, the real-time implementation runs through a new HDF5 interface. Since the collection of decision trees can often translate into a huge external function (sizes of hundreds of MB and more), this strongly affects the PCS compiling time. Therefore, we remapped the DPRF trained structure from a python object to an HDF5 file format. By doing so, there is no need of a priori linking it into the PCS and the model data can be loaded in (even different model data for different phases) at runtime, during experiments. This is a newly developed and general HDF5 interface that can be adapted for other data-driven algorithms, outside DPRF. The final advantage is the speed-up of PCS rebuilding/compiling time and the flexibility for retraining the algorithms between run days or even discharges. Similar solutions are foreseen for EAST's implementation as well, thanks to the synergy of the two control systems: this will be part of the next steps in the international collaboration.

3.2 DPRF and the DIII-D proximity control architecture

The successful integration of DPRF in DIII-D PCS can be considered part of a broader approach to qualify advanced disruption prevention strategies to address ITER's critical needs. The "Disruption Free Protocol" [J. Barr et al IAEA-FEC 2020] consists in a suite of novel prevention techniques that are being developed and tested on DIII-D to ensure ITER's success. A new block (the "Proximity Controller") in the DIII-D PCS architecture was designed and is being validated to regulate the plasma proximity to instabilities. The Proximity Controller was first tested to prevent VDEs in controlled experiments and more recently DPRF was also connected to it, as it provides a general plasma stability metric through the real-time disruptivity computation, associated to the crucial information on disruption causes given by the feature contributions. The Proximity Controller takes in the disruptivity (f_{danger}), the particular feature contribution ($f_{i,contrib}$), and by combining it with the trend of the actual input plasma signal, the requested change in control targets is calculated through a proportional-integral-derivative (PID) controller. Eventually, these changes are sent to the appropriate actuator control algorithm. When DPRF is used for continuous plasma monitoring in combination with the Proximity

[2] sklearn-porter, https://github.com/nok/sklearn-porter, Morawiec, unpublished.

Controller (Fig. 4), plasma performance can be regulated ahead of time before any asynchronous response pre-programmed via the ONFR algorithm. Following the example sketch in Fig. 4, by monitoring the elongation contribution to the disruptivity, the Proximity Controller will calculate the requested target change as: $\Delta\kappa = PID\left[f_{danger} * f_{\kappa,contrib} * sign\left(\frac{d\kappa}{dt}\right)\frac{\Delta\kappa_{target}}{\Delta f_{\kappa,contrib}}\right]$ and act on DIII-D shaping actuators to keep the plasma away from reaching a dangerous stability limit.

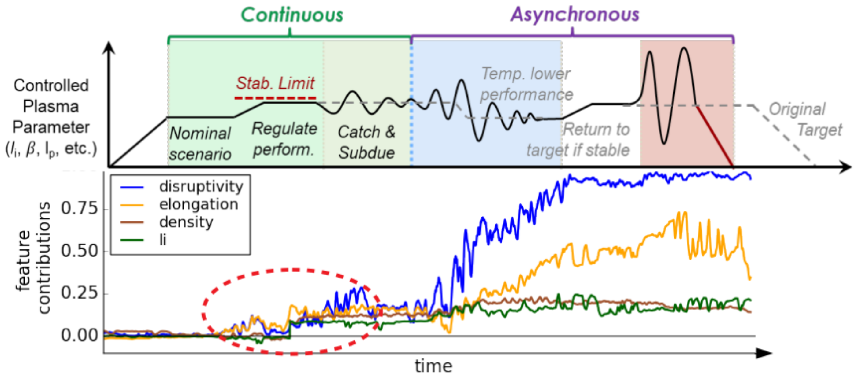

Fig. 4 Cartoon showing different control regimes, continuous vs asynchronous, to monitor controllable plasma parameters. The bottom panel shows an example situation where the disruptivity starts deviating before a clear loss of performance, and thanks to the information provided by the feature contributions, a requested change to the target plasma features can be calculated via the proximity control architecture.

DPRF integration with the Proximity Controller has already been tested in piggyback experiments during 2020 and dedicated control test sessions are scheduled in 2021 to fine tune the PID controller and demonstrate such active disruption prevention strategy.

4. EAST DPRF EXPERIMENTS

Among the different precursors of the final loss of control in tokamak plasmas [DEV2011], density-limit phenomena mark the reaching of stability limits or poor confinement modes that can narrow the operational space. ITER is expected to operate at a Greenwald density fraction of about 0.85 [MUK2003], making it prone to density limit events. To properly identify this disruptive dynamics, one should investigate the chain of events in all collected high density disruptions on EAST to find patterns similar to an uncontrolled increasing plasma density combined with a rise in impurity radiation from the edge, edge cooling and the steepening of the current profile in the proximity of the q = 2 surface, together with destabilization of n = 1 MHD modes (principally those with poloidal mode number m = 2). This is a rather complex effort, making the identification of density-limit cases quite difficult to be carried out at scale. In this application of DPRF to EAST, we have considered high density disruptions that include density-limit cases, but might also contain spurious dynamics such as uncontrolled density ramps, density feedback failures, etc… As part of future work, we will improve the training dataset by implementing automated ways [MON2021] of detecting the transition to an unstable scenario, and therefore refine our classification sensitivity. Nevertheless, this is the first time that a real-time data-driven

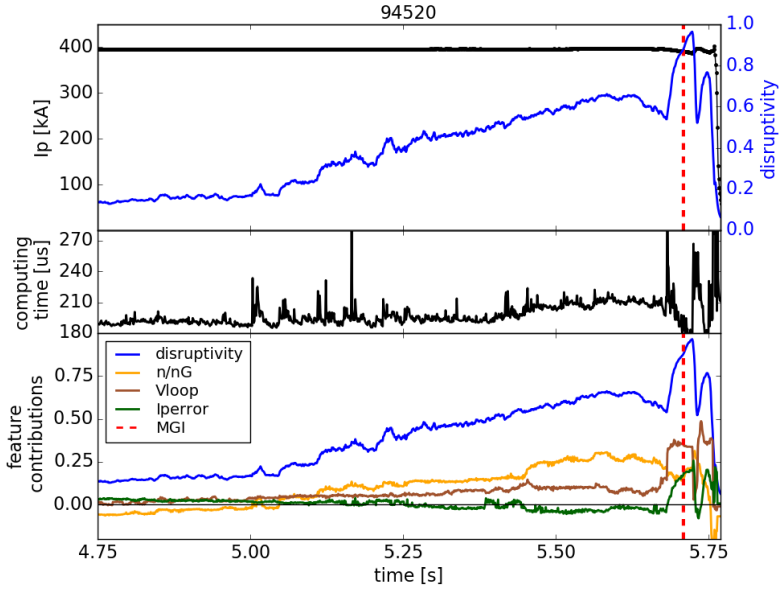

Fig. 5 DPRF predictions of impending disruption and its main drivers on EAST. The top panel shows the last second before the disruption, when the disruptivity (in blue) rises above 80%. The middle panel reports the PCS computing time of both predictions and its interpretations (200-220 μs). The bottom panel displays in color the disruptivity, and the most important drivers, i.e. the Greenwald density fraction, the loop voltage, and the fractional error on the plasma current. Less relevant contributions are not reported here. The red dashed line shows the MGI trigger time at 5.71 s.

predictive model for disruption warning is operational on the EAST superconducting tokamak.

EAST DPRF was trained with data from 2015 to 2018 experimental campaigns, and in particular 483 disruptive discharges were selected where the line-averaged density value was found to reach or go beyond 80% of the Greenwald density limit during the flattop phase of the plasma current. An analogous number of non-disruptive discharges was randomly included from a different variety of scenarios – all samples were taken from the flattop phase of the plasma current. DPRF was installed as part of the available alarms that can trigger EAST's mitigation system: to gain confidence in the algorithm, DPRF has run in background in EAST PCS during the 2020 experimental campaign and only a few dedicated sessions were used to test the closed-loop performances. In Fig. 5 we show one example of a discharge from January 2020, where real-time predictions and interpretations were provided in around 200 microseconds computing time. EAST discharge 94520 shows symptoms of a density-limit disruption, due to the high contribution from the Greenwald density fraction, but eventually disrupts due to a radiative collapse caused by impurities present in the plasma, as also indicated by the dominance of the loop voltage contribution at the end. The MGI is triggered with no gas injection at 5.71 s, 60 ms before the current quench.

Generally, EAST DPRF has shown good performances in post-mortem analysis, achieving a successful alarm rate of 92% and a false alarm rate of 9.9% with a disruptivity threshold set at 80%. The accumulated fraction of detected disruptions shown in Figure 6 as a function of the warning time provides more information about DPRF performances: a plateau of 92% of successful alarms is reached before incurring in the late detection area, where $t_{warning}$ < 30 ms, i.e. the required time for MGI to be effective. The more extended, submitted paper by the authors [HUW2021] presents more details about EAST DPRF analysis and results.

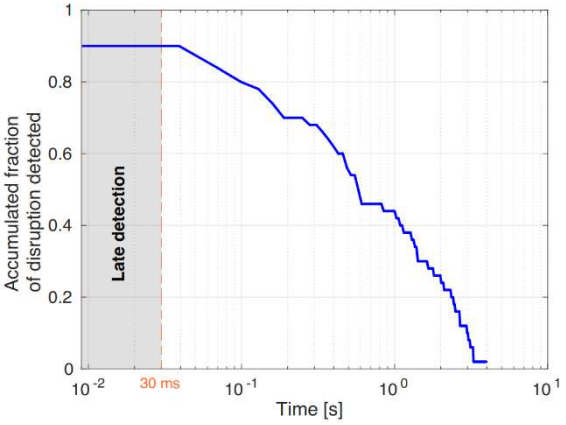

Fig. 6 Figure reproduced from [HUW2021]. Accumulated fraction of detected disruptions on EAST by DPRF, as a function of the warning time (in log scale). The shadowed gray area identifies the late detections, where $t_{warning}$ < 30ms, for which the MGI would not have any mitigation effect. EAST DPRF reaches a performance plateau between 40 and 50 ms before the disruption, guaranteeing more than 90% of correct classifications, while keeping the false alarm rate at values less than 10%.

The developed algorithm still presents some residual flaws as it shows a relatively high false alarm rate. Furthermore, more work will be required to refine the training set and identify on a shot-by-shot basis the transition time into an unstable operational space[3], and additional efforts will be needed to implement peaking factors from kinetic profiles similarly to what we have done on DIII-D. These steps are crucial to validate a data-driven algorithm in a device-independent framework aimed at properly diagnosing disruption precursors.

5. CONCLUSIONS AND OUTLOOK

The DPRF algorithm provides explainable predictions in real-time thanks to the feature contribution analysis that can be used to identify the main drivers of the disruptivity. DPRF is now part of the real-time algorithms suite in DIII-D and EAST Plasma Control Systems, although the two implementations differ slightly: on DIII-D, a more refined training dataset combined with the availability in real-time of profile peaking factors allows the model to be sensitive to a variety of pre-disruptive precursors, such as impurity accumulation events or MHD-driven disruptions. On EAST, a "less mature" version of DPRF was trained only on high density disruptions with a limited set of input features mainly coming from the P-EFIT equilibrium reconstruction code. Peaking factors for the electron temperature, density and radiation profiles are already available offline on EAST, and we plan to use the radiation peaking factors in the upcoming 2021 experimental campaign, as more AXUV channels will be digitized in real-time and made available to the PCS.

DPRF was tested in closed-loop mitigation experiments on both DIII-D and EAST tokamaks, showing average computing times on both devices around 200 μs. On DIII-D, DPRF is part of the "Disruption Free Protocol" and on track to be tested in 2021 dedicated experiments in conjunction with the Proximity Controller to continuously monitor the plasma state and optimize plasma performances. On EAST, DPRF was actively used in a limited

[3]For EAST DPRF application, the transition time was kept fixed at 1.7 s before the disruption, and found as an additional model parameter through an offline K-fold optimization following what described in [MON2019].

number of discharges to demonstrate its efficacy as MGI trigger. To gain confidence in the tool, DPRF was used in background during routine plasma operations to gather performance statistics and in post-mortem analysis it has shown a successful alarm rate of 92% and a false alarm rate of 9.9%. More details about EAST DPRF analysis and results can be found in another submitted paper by the authors [HUW2021].

Current results represent a step forward in the development of data-driven predictive algorithms for scenario optimization and disruption avoidance, first in existing devices (DIII-D and EAST) and then for ITER and the next generation of tokamaks. This work establishes the importance of developing tools capable of identifying and informing in real-time the PCS on the dangerous plasma parameters deviations to the disruptive space in order to enable the proper actuators' response. The final aim is to obtain scenario-dependent event probabilities that can extend the plasma survival by promptly navigating away from instability boundaries, thus optimizing plasma performance.

ACKNOWLEDGEMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, and Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award DE-FC02-99ER54512, DE-SC0014264, DE-SC0010720, DE-SC0010492, and DE-FC02-04ER54698. Additionally, this work is supported by the National MCF Energy R&D Program of China, Grant No. 2018YFE0302100. Part of the data analysis was performed using the OMFIT integrated modelling framework [MEN2015]. **DISCLAIMER**: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

REFERENCES

[BRE2001] BREIMAN, L. 2001 Machine Learning 45 5–32.
[BYK2017] BYKOV, I. et al. 2017 Phys. Scr. T170 014034.
[CPP2020] "Advance multiple methods for disruption prediction, avoidance, and mitigation to inform FPP design", page 71, A Community Plan for Fusion Energy and Discovery Plasma Sciences, Report of the 2019–2020 American Physical Society Division of Plasma Physics Community Planning Process. https://drive.google.com/open?id=1w0TKL_Jn0tKUBgUc8RC1s5fIOViH5pRK
[DEV2011] DE VRIES, P.C. et al. 2011 Nucl. Fusion 51 053018.
[DOY2010] DOYLE E.J., et al 2010 Nucl. Fusion 50 075005.
[EID2018] EIDIETIS, N.W., et al. 2018 Nucl. Fusion 58 056023.
[FER1998] FERRON, J., et al. 1998 Nucl. Fusion 38 1055–66.
[HEN2007] HENDER, T, et al. 2007 Nucl. Fusion 47 S128–S202.
[HUA2016] HUANG, Y., et al. 2016 Fusion Eng. and Des. 112 1019–1024.
[HUA2020] HUANG, Y., et al. 2020 Nuclear Fusion.
[HUW2021] HU, W., et al 2021 Nucl. Fusion in review.
[MEN2015] MENEGHINI, O.M. et al. 2015 Nucl. Fusion 55 083008.
[MON2019] MONTES, K.J., et al. 2019 Nucl. Fusion 59 096015.
[MON2021] MONTES, K.J., et al. 2021 Nucl. Fusion 61 026022.
[MUK2003] MUKHOVATOV, V. et al. 2003 Plasma Phys. Contr. Fusion 45 A235.
[PAL2014] PALCZEWSKA, A. et al. 2014 in Integration of Reusable Systems (Advances in Intelligent Systems and Computing vol 263) ed T. Bouabana-Tebibel and S.H. Rubin (Cham: Springer)
[PAU2018] PAU, A., et al. 2018 IEEE TPS 46 7 2691.
[PED2012] PEDREGOSA, F. et al 2012 J. Mach. Learn. Res. 12 2825–30.
[REA2019] REA, C., et al. 2019 Nucl. Fusion 59 096016.
[REA2020] REA, C., et al. 2020 Fus. Sci. Technology 76 912–924.
[SAB2018] SABBAGH, S.A. et al., "Disruption event characterization and forecasting in tokamaks", 27th IAEA-FEC, Conf. Proc. EX/P6-26 (2018).
[TUR2018] TURCO, F., et al. 2018 Nucl. Fusion 58 106043.