

A new phase-oriented disruption prediction strategy for mitigation, prevention and avoidance in JET

G.A. Rattá¹, J. Vega¹, A. Murari², D. Gadariya¹ and JET Contributors*

¹Laboratorio Nacional de Fusión. CIEMAT, Madrid, Spain.

²Consorzio RFX (CNR, ENEA, INFN, Università di Padova)

* See the author list of E. Joffrin et al. accepted for publication in Nuclear Fusion Special issue 2019,

<https://doi.org/10.1088/1741-4326/ab2276>

Abstract

The ideal operational scenario for the future Tokamak reactor is disruption free. However, so far all the experimental evidence indicates that disruptions are unavoidable and can occur with alarming frequency, when approaching reactor conditions (low q_{95} , high radiated fraction, divertor detachment, etc.). In this article, a unified strategy for disruption avoidance, prevention and mitigation is proposed and validated on JET data. The approach is based on three phase-oriented predictors to detect the main instabilities, leading to the undesired and sudden end of the discharge. The first model detects dangerous profiles as an early indication of a critical situation. The second one is designed to identify MARFEs (Multifaceted Asymmetric Radiation From the Edge) and other abnormal radiative events. The third model is devoted to mitigation and triggers alarms around few tens of ms before the beginning of the current quench. The models have been trained and tested with a database of almost one thousand JET discharges of recent campaigns with the ITER Like Wall. The overall performances are very close to 100 % of successful detections with a few percent of false alarms. In addition to the first systematic use of imaging cameras, the most relevant aspect of this work is related to the distribution of the alarms of the three predictors, which do not overlap and are sequential. Consequently, the three predictors are meant to work in parallel over running discharges and, depending on which one triggers the alarm, the cause can be determined and the approximate remaining time to intervene can be estimated, potentially allowing the optimisation of the remedial actions.

Keywords: Disruptions, avoidance, prevention, mitigation, profile indicators, MARFEs, Genetic Algorithms, SVM.

1- Introduction

In the perspective of the next generation of tokamaks, such as ITER and DEMO, it is urgent to deal with the problem of disruptions. The sudden and unplanned termination of the plasma can indeed induce huge heat loads on the plasma facing components and high electromechanical forces on the structures of the devices. In the future commercial reactors, it is estimated that disruptions will have to be completely avoided, since even a single one could compromise their integrity [1].

On the contrary, in present Tokamaks, disruptions are not only unavoidable but can occur at an alarming frequency. The baseline at low safety factor (around $q_{95}=3$), the reference scenario in ITER, is particularly vulnerable. On JET, in some high current ($I_p \geq 2.5\text{MA}$), low q_{95} campaigns (as the one of 2016), the disruptivity rate reached 60%, even for a low radiated fraction. Nevertheless, those rates were posteriorly improved. For example, in the 2019-2020 campaign, it was reduced to a 21% for $I_p > 2\text{MA}$ disruptions. It should also be remembered that the next generation of devices will have to operate with radiation above 90% of the input plus alpha particle power, and with fully detached plasmas [2].

At low q_{95} in metallic first wall machines, excessive radiation in the core, combined with hollow electron temperature profiles, is believed to be one of the most common causes of disruptions [3]. Of course, many other different instabilities can lead to the plasma collapse, ranging from local density limit to Neoclassical Tearing Modes (NTMs), Edge-Localize Modes (ELMs), sawteeth, etc [3].

Tackling disruptions is therefore recognised as a complex task, which the community has tried to address in the past, focusing on three strategies: avoidance, prevention and mitigation. Avoidance means remaining in the safe operational region or returning to it sufficiently quickly not to alter the behaviour of the discharge. The next step is prevention, which comprises of a series of measure to terminate a discharge safely, once the plasma has shifted into a situation, which barring correcting actions would lead to a disruption. Mitigation consists of the set of measures, such as massive gas injection or shattered pellets, to alleviate the consequences of disruptions, once they have become unavoidable. Therefore, from the point of view of the control systems, avoidance, prevention and mitigation correspond to three different situations, characterised by completely different measures to be undertaken. They will be called disruption phases in

the rest of the paper. Again with regard to nomenclature and definitions, in this work the beginning of the current quench is considered the time of the disruption.

The main approach to avoidance, of course, would consist of running the plasma in the part of the operational space not prone to disrupt. This strategy is constantly evolving and the evidence indicates that, in JET with the carbon wall, a significant reduction in disruptivity was obtained far from the operational boundaries, leading to the conclusion that decreasing the disruptivity rate is also linked to a better operational experience [3]. In this context, the development of theoretical models and plasma simulators are promising lines of research. However, this is a challenging task at the point that theoretical models and plasma simulators have not reached enough reliability to be tested in real-time on JET [4].

Therefore, another approach to achieving sufficient avoidance consists of identifying early precursors of the disruptions, in order to intervene and steer the pulse back to a safe region of the operational space. In this specific case, the minimum time to develop avoidance actions will depend on the type of the incoming disruption and its characteristic evolution time (i.e. from the moment that the first early precursor can be detected until the disruption occurs). Therefore, the requirements for avoidance are: 1. to detect precursors hundreds of millisecond before the sudden end of the discharge and 2. to identify the type of disruption (or its root cause) in order to perform the most adequate avoidance action. Currently at JET, significant efforts have been devoted to pursuing these challenging objectives. The idea is to recognise a cascade of a real-time events with the system PETRA (Plasma Event and TRigger for Avoidance) to improve the current real-time protection modules. PETRA is meant to comprise several disruption detection systems based on both physics and data-driven predictors. On the other hand, in the last years, there have been only a reduced number of published works devoted to document the effectiveness of avoidance actions in JET. One of these was based on Generative Topographic Maps (GTM), a methodology already applied for disruptions more than a decade before [5]. This first paper (2010) exploited the main potential of GTM, the dimensionality reduction, to understand in 2 dimensions the way 7 signals (7 dimensions) evolve as disruptions approach. In the second one [6], the same technique is used to find a model with large warning times, aiming at performing avoidance actions, with the purpose of monitoring the disruptions dynamics and its physics mechanisms. The results of that work are promising and it would be very interesting to see the performance of that model over a wide database.

Prevention is also based on the early detection of precursors, so that the control system can intervene and shut down the plasma safely. The approaches to prevention most commonly pursued so far have been extensions of the one adopted for mitigation. In practice, the results have not been very positive in the sense that predictors can provide sufficient anticipation time but with high standard deviations. In other words, when an alarm is raised, there is no way to know the time to the disruption and, therefore, although prevention strategies are fired, most of the times, mitigation actions are finally applied. To improve the drawback of not being able to complete prevention measures due to the lack of estimations of the time to the disruption, research on predicting the time to the disruption is more than welcome. A recent approach about this [7] presented promising results but it is only a starting point.

The research involving mitigation has been intense, particularly in the last 15 years. The strategy, in this case, is based on puffing gas or injecting pellets, once the disruption is imminent and unavoidable, in order to alleviate the consequences of the thermal and current quenches. On JET, the predictors for mitigation, attempt to detect precursors tens of ms before the disruptions (the time required by the actuators plus the flight time of pellets or gas). Historically, the most common models used for mitigation in JET have been based on setting thresholds in the locked mode signals and occasionally also on others such as plasma energy and radiation measurements. Using these simple methods, the prediction results are modest but with a clear physics understanding of the reason why each alarm is triggered. Nevertheless, the plasma evolution towards a disruption is more complex. Thus, the consideration of non-linear relationships among a wider set of plasma parameters can improve the prediction rates. To model these complex relationships accurately, machine-learning predictors have been explored and developed in the last decades, including approaches from scratch [8], real time implementations [9] and adaptive learning [10]. Among these systems, the APODIS prototype [11] gained critical relevance since its detection rates were the highest ever in JET at the moment of its publication. Therefore, a revised version of the system was installed in the JET real-time network with a success rate of 98.36% and only a 0.92% of false alarms [12]. The robustness of APODIS was verified in practice over the years; its prediction rates remained high, without any retraining of the model, even after considerable structural modifications of the device such as the installation of the new ITER Like Wall (ILW). Another two prediction models worth mentioning are based in centroid methods [7] and outlier detections [13][14]. This last model (SPAD predictor) achieved high success rates

in JET by identifying abnormal changes in the locked mode signal. These systems have been the only three, based on machine learning techniques, successfully applied online [9] on JET so far. Their success is due, at least in part, to the fact that they were tested with a wide independent dataset covering real-time operational conditions.

In any case, even the best predictors deployed on JET have not been very useful for either avoidance or prevention. This is mainly due to excessive spread of their warning times, which can range wildly from a few ms to more than 1 second. In the present article, several limitations of the state of the art predictors (particularly the ones for avoidance) are addressed. The goal is creating not only a model but, instead, a whole system to deal with disruptions in general from avoidance to prevention and mitigation. For that, three predictors able to work independently but also simultaneously over running discharges have been developed. Each one of these models targets a specific phase of the discharge and the related disruption-types, in order to get not only an alarm once a precursor is identified, but also the reason why it has been triggered together with an estimation of the time to the disruption, as it has been recently explored in other research with promising results [7].

The first of the three models, introduced in Section 4.1 of this article, is aimed at detecting dangerously situations in the core, mainly hollow electron temperature profiles. For that, it uses data from the High Resolution Thomson Scattering (HRTS) [15]. To gain extra accuracy, also plasma density profiles, SXR and bolometric measurements are used.

The second predictor, detailed in Section 4.2, is tuned to identify MARFEs (Multifaceted Asymmetric Radiation From the Edge) [16] and other radiative events. In this case, the approach differs significantly from the most common one, based on the use of bolometric measurements to detect spikes or unusual behaviours in the radiation. Instead, the data is extracted and processed only from videos of the wide angle operational visible camera. The selection of several Regions Of Interest (ROIs), and their processing to generate time-traces susceptible of identifying MARFEs, has proved crucial to obtain accurate results.

Finally, the third system is focused on mitigation and it is designed and optimized to trigger alarms around 40 ms before the beginning of the current quench. This time is enough to develop mitigation actions with both massive gas injection and pellets.

The three models are meant to work in parallel during the discharges and when one of them activates an alarm, the control system is provided with specific information to optimise the remedial action. The integrated system can indicate which predictor is the

one triggering the alarm, with the associated information about the discharge phase, the disruption type and the bounds on the remaining time to the beginning of the current quench.

With regard to the structure of the paper, next section is an overview of the machine learning tools, on which the predictors are based. The analysed database is described in detail in Section 3, while the signals and indicators provided as inputs to the predictors are covered in Section 4. Section 5 is devoted to the integration of the various predictors into a single system. The results and performances are detailed in Section 6 before the summary and discussion provided in the last section of the paper.

2. The main Artificial Intelligence techniques implemented by the predictors

2.1. Introduction to supervised learning

The field of data mining or data exploration has flourished since the beginning of the century thanks also to the progress in computational and data storage technologies. Recently techniques inspired by nature, such as Genetic Algorithms (GA), have started providing excellent results, especially in cases where it is infeasible to reach reliable outcomes by conventional algorithms or in reasonable computational times.

Machine Learning (ML) (that includes techniques such as Support Vector Machines (SVM) [17] and Artificial Neural Networks [18]), aims at creating mathematical models based on data samples. The samples used to develop the models are known as “training samples”. They could be, for instance, a feature vector that includes the value of several parameters (as the plasma current, the plasma internal inductance and the “hollowness” of a Te profile) at a given time.

“Supervised” ML consists of labelling each training sample (i.e. the feature vector with the parameters). The tag could be, for example, “disruptive” or “NON disruptive”. The fundament of ML for classification consists of calculating a rule in order to differentiate the samples according to their label (mathematically speaking, it computes a separating hyper-plane able to split non-disruptive and disruptive feature vectors in the feature space).

Once the model is trained, it can be validated using new feature vectors. The goal of the training is to reach a general rule to correctly predict the label of the new data. In case of insufficient performance, the classifier would need to be optimized (with methods like

the Genetic Algorithms), by tuning several parameters or revising the training samples. This revision is quite relevant and sometimes even crucial to obtain good results: in case incorrect examples (e.g. disruptive samples labelled as “non-disruptive”) are provided, the ML system will develop an inaccurate classifier.

The validation dataset may be a subgroup of the “training dataset”, a fraction of it. The latter dataset is used exclusively to create the models and it must be completely different than the final “testing” dataset. Once the model has been trained (and validated) the “testing” dataset is employed for the last evaluation. The final and valid results, the only relevant ones to be trusted, are the ones obtained in this final stage.

In this work the applied ML classifier is SVM and the optimization method GAs, both introduced in the following subsections.

2.2. Support Vector Machines

As in many previous studies by some of the authors, the supervised ML technique applied in this case is SVM [17]. The main reasons for choosing this method are its good performance, fast computational times and proven generalization capabilities. In addition, it is relevant to mention that SVM provides always the same final result in case it is trained with the same data. Other techniques, as most types of Neural Networks, may provide different results with the same training data and very different outcomes with slightly modified training datasets, which can deeply affect the extrapolation of these models to other devices or regimes of operation.

The solution provided by SVM is the linear hyper-plane that splits the classes maximizing the distance between the boundary samples. However, real problems (especially the ones that require complex techniques as AI) may have a non-linear nature. To solve them with the linear SVM margin maximization principle, Kernels functions can be applied. These functions map the input space into a higher dimensional feature space. Transforming the space in such a way allows computing a linear solution able to maximize the quality of the classification. It is important to keep in mind that the solution is linear in the higher dimensional space but it would generate a non-linear model in the input space. Kernel functions can take different forms. The one used in this work has been one of the most versatile, the Radial Basis Function (RBF):

$$K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2) \quad (1)$$

It is important to notice that the value of the parameter γ (as well as the slack variable c) must be defined in advance. This c variable is fundamental to avoid overfitting. A higher value of c implies a harder penalization to those objects in misplaced with respect to the computed hyper-plane.

For this work the free licensed software LIBSVM has been used and has been adapted to run under MATLAB.

2.3. Genetic algorithms

GAs are a set of computer methods, based on *artificial selection* [19], mainly used for finding good solutions to complex optimization problems. *Artificial selection* differs from *natural selection*. In natural selection, living beings are optimized for generic survival and reproduction. In artificial selection, the objectives are more specific and predefined. A widespread example is the case of wolves. They were selected by humans based on their docility, fidelity, obedience and even cuteness and in only few thousand years (a very short time in evolutionary terms) some of them went through drastic changes becoming puddles, toy spaniels and all the wide range of sub-breeds.

GAs are based on the principle of artificial selection. Given a problem, a population of possible models or individuals (each one with their own characteristics/genes) is created. The goodness of each individual is determined with the help of a predefined metric called Fitness Function (FF). That means that the objective to be optimized can be predefined and the solutions are going to evolve towards that target.

The application of the technique to disruption prediction can be summarized in the following computational steps:

STEP 1- generation of the population of individuals.

In this study 50 *individuals* per generation were chosen as a good trade-off between results quality and computational times. Every different *individual* is a vector (see Fig. 1.a) that contains all the required instructions to create a disruption predictor: the combination of signals to be used and the values of each Kernel parameter and the slack variables in SVM.

Notice that the *individual* contains rudimentary instructions: 'ones' and 'zeros' are assigned to each position in the vector. Each position is linked to a parameter (e.g. the Hollow Factor, the Bolometry Factor or the SVM slack variable C). In the codification,

the 'ones' indicate that their linked parameters must be included in the predictors' development whereas the 'zeros' mean they must be not. Only in the first iteration these ones and zeros are assigned randomly. In the posterior iterations/generations the GA heuristics take place in the creation of the new generations' characteristics.

STEP 2- Training of the predictors.

The individuals are not predictors but instructions to create them. Each individual, then, determine the set of parameters to train and create an SVM predictor.

In this step, the 50 individuals' instructions (ones and zeros indicating which parameters must be included and the different values for the SVM variables) are followed to train 50 SVM predictors. The training database is used.

STEP 3- Evaluation.

The 50 trained (in the previous STEP) predictors are evaluated with the validation dataset. 10 points are assigned each time a predictor triggers an alarm before the disruption. 5 points are given to those predictors that, correctly, do not activate an alarm in a non-disruptive shot. Notice that this criterion (usually called Fitness Function) can be adjusted *ad hoc* to optimize and shape the predictors according to specific requirements. As a result, each predictor has a number (total summation of the points) that quantify their fitness, i.e. their performance for solving the problem.

STEP 4- Selection of parents.

The idea is to assign a higher chance to be selected as parent to those individuals with a higher fitness score, in order to mix their good characteristics. For this purpose, the roulette method has been applied: i) The Fitness Function scores are sorted and normalized between [0; 1]; ii) a random value between 0 and 1 is chosen (roulette); iii) all individuals whose normalized fitness is over the randomly selected value are added to the bag of selected parents; iv) the process is repeated from step i) until the number of parents is equal to the population number (50 in this case).

STEP 5- Cross-over and creation of children.

To create children as a combination of parents' characteristics, the 2 points cross-over operation has been applied (see Fig. 1.b). It consists of randomly selecting 2 points

in 2 paired parent's vectors. Then, the sections determined by these points are interchanged between those pair of parents to create offspring. The operation is repeated for all the pair of parents to create a new population of 50 children. This population replaces the previous one.

Notice that the concept (combining the ‘genes’ of fitted parents to create new promising combinations) is carried out with an efficient simplicity.

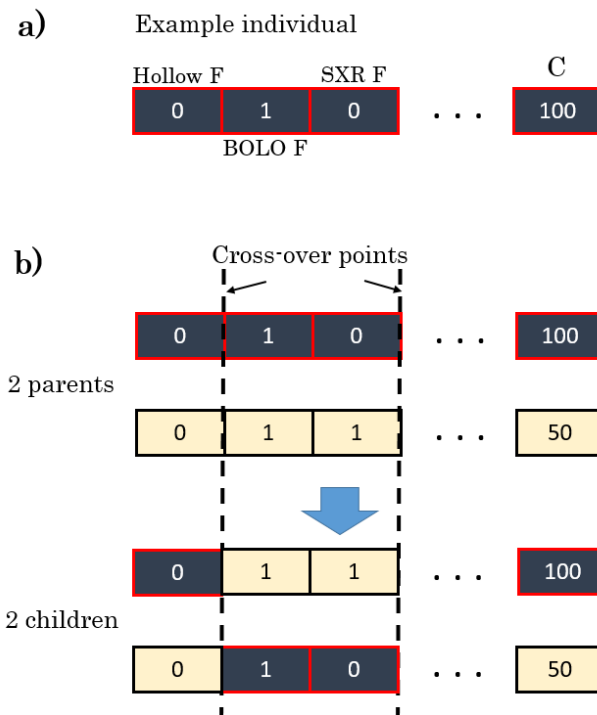


Figure 1. a) Example of an individual codification. b) Cross-over operation to create children.

STEP 6- Iterate until the ending condition is satisfied.

Unless an ending condition is satisfied (reach 100 iterations), iterate from STEP 2. The 100 iterations were selected as a trade-off between good results in affordable computational times (12 hours per run in an Intel(R) Core (TM) i7-8700 CPU @3.2GHz with 16 GB of RAM memory).

A more detailed explanation of the GA routines used in the present work can be found in [\[20\]\[21\]\[22\]](#).

3. Database

Most of the problems in the development of predictive models (not only the ones addressing disruptions) derive from the way they are derived and tested. Models built using small databases may lead to unreliable conclusions. Therefore, in this work a large and very recent dataset has been gathered and used in order to obtain reliable results.

The *initial* database consisted of 1227 discharges (900 non-disruptive and 327 disruptive) corresponding to the experimental campaigns C38 and C38b (June 2019 till March 2020, from #94152 to #96745). All the non-intentional disruptive shots in the period were included. Since there is no selection of the shots, there is a wide diversity in the range of the plasma current of the shots (maximum values from 1 MA to 3,6 MA). It is relevant to mention that all discharges with $I_p > 2$ MA have the Disruption Mitigation Valves (DMVs) included as a protection system. So, in most of the disruptive cases, the plasma was terminated by these DMVs, generally triggered by the locked mode signal.

Regarding the data collected for each one of the 1227 pulses, the following signals have been gathered in the database:

- a) videos of the operational camera the MARFE detector;
- b) data from HRTS (both Te and ne profiles), Soft X Rays and bolometry for the Profile predictor;
- c) 1. Locked mode amplitude; 2. Total input power; 3. bolometry (the same time-trace used in the Profile predictor); 4. Plasma vertical centroid position and 5. Plasma Current for the mitigation model.

The dimension of the database is around 50 Gigabytes.

Since the objective is to use the three predictors scanning for precursors in parallel over each discharge, a necessary condition to consider any given shot is the availability of all the signals required by the 3 models. Unfortunately, in some discharges, a signal, a profile or a video is not available. In that situation, the shot has to be discarded from the database. The most common missing quantities are the Te and ne profiles (not available in 240 shots). For this reason, together with the fact that a few other signals and videos are missing, a total of 974 shots remains in the database (263 of them disruptive and 711 non-disruptive).

All the time-traces used for the three predictors (including the ones derived from videos or profiles) require to be synchronized, which means to have a common time base.

They have been resampled to 500 samples/s (the same as JET real time network), following a similar interpolation methodology as the one detailed in [23].

One of the most common and accepted criteria to develop predictors is to use 2/3 of the database for training and validation and to save 1/3 of the data for the final testing. This is the criterion adopted in this article. For the training/validation database 176 (randomly selected) disruptive and 470 (randomly selected) non-disruptive discharges have been used. The final test is performed using the remaining 87 disruptive and 241 non-disruptive shots.

For every time-trace in the database, Time Derivatives (TD) have also been computed [21]. TDs represent the local increment of the amplitude values of a signal X (ΔX) over a predefined temporal difference (i.e. $\Delta X/\Delta t$). Two different Δt have been used: 2 and 50 ms. These TDs may contribute significantly to improve the results because they get rid of possible off-sets in the signals (their calculation is performed over closely separated samples, so any inherited DC component is avoided). Also, since the past samples are involved in their computation, they supply valuable information about what has happened before.

For the sake of clarity and as an example of the notation for the TD of time-traces, to express the TD of the Hollowness Te Factor for $\Delta t=50$ ms the notation $\text{Hollowness Te Factor}_{\text{TD}50}$ will be used. The TD of the rest of temporal evolution data will be expressed in a similar way, with the help of analogue subscripts.

Regarding the training, validation and testing of the model, the following procedure has been implemented:

For training and validation of each predictor, the training/validation database has been used. In each case, the process is similar to the ones performed and detailed in previous publications [21][22]. To train each predictor, SVM with a Gaussian kernel is used. The training process uses a combination of SVM and GAs to find out not only the typical ridge penalty term of SVM and the kernel parameter but also a proper combination of time traces and time derivatives. To this end, a different set of parameters (detailed in the next Section) is provided to a combined system that joins SVM and GAs.

SVM generated models are non-linear and they include a high dimensional matrix (see [22]). The parameters (i.e. the combination time-traces, time derivatives, the slack variable C and the parameter γ required by SVM) are chosen and refined with GAs.

4. The disruption phase-oriented indicators

Each disruption type oriented predictor requires specific indicators to properly trigger the alarms. These indicators are sets of variables obtained as result of a combination of plasma measurements. They can be selected from JET database or processed in order to target the phenomenon, which disruption predictor is focused on.

In the next subsections the computation of these indicators, for all the predictors, are detailed.

4.1. Profile indicators for avoidance

Hollow Te profiles can be an early indicator of an incoming disruption and therefore they were studied in the past [6] as possible precursors. Since in JET not all hollow Te profiles lead to a disruption, it is necessary to quantify the degree of “hollowness”. Moreover, the disruptivity of a discharge depends not only on the hollowness of the Te profile but is influenced by other factors, which have to be properly quantified.

Regarding the kinetic profiles, the High Resolution Thomson Scattering (HRTS) has been selected. One of the main reasons for using the HRTS is the availability of its data in most of the discharges. The Electron Cyclotron Emission (ECE) Michelson Interferometer, which has been successfully used to detect hollow temperature profiles in real time during the current ramp phase in JET [24], could have been a reasonable alternative. Unfortunately, it is frequently affected by cut-off at specific radial locations, which happens for various combinations of density and magnetic field. This can be a serious drawback, especially in the perspective a possible future deployment online.

The HRTS measures the electron Temperature (Te) and the electron density (ne) with a rate of 20 light pulses per second (20 Hz), providing 63 data points per profile; its spatial resolutions for the core region and the pedestal are 1.6 and 1 cm, respectively. The sampling rate is clearly not high enough to identify fast events for mitigation, but since in the present application it is meant to contribute to early detection for avoidance, the reduced time resolution is a tolerable limitation.

The metric to quantify the Te Hollowness Factor has been determined as explained in the following. To this end, a reduced database of 20 Hollow Te discharges and 20 randomly selected non-disruptive discharges (DB_{HOLLOW}) has been built. This database has been thoroughly tested in order to deduce the following two key parameters:

1. The subset of relevant channels (out of the total of 63 measured by the HRTS).
2. The Expression (E) that defines the relationship among these channels to represent the Hollowness Te Factor.

A scan, exploring a combination of the 1st to the 40th more internal channels of the HRTS, has been performed. The final adopted selection includes the channels from number 1 (corresponding to 2,98 m) to channel 29 (3,44 m) and the expression implemented for E, supported by the GA optimization, is:

$$E = \text{channel 29} - \text{channel 1} \quad (2)$$

Notice that, in this simple formula, a positive value means that the measurement closer to the plasma core has a lower temperature than the outer one. Then, the higher E the hollower the profile.

The same formula (E) with the opposite sign has been implemented to calculate the peaked ne Factor:

$$E = - \text{channel 29} + \text{channel 1}$$

The change in the sign is chosen to obtain high values for peaked profiles (and therefore to have a more intuitive visualization of the phenomenon).

These two indicators are a good starting point, proving the potential of the HRTS to quantify the character of the profiles. However, the initial tests (using only these variables) resulted in a high number of false alarms (over the 20%) and some extra parameters are evidently needed, in order to improve the accuracy of the predictions.

To this end, the Soft X Rays have been considered. The rationale for this choice is that the SXR emission is influenced by the W content, therefore potentially providing information that can contribute to detecting whether the plasma is drifting towards a critical situation [25].

In the present application, the vertical SXR- camera, with 250 μm Be-filter, has been used [26]. To create a ‘‘SXR Factor’’, emissions coming from the plasma core (C) and the plasma edges (I representing the Inner part of the vessel and O the Outer side) have been synthesized in the following formula:

$$\text{SXR Factor} = \frac{2C}{I_{1,2} + O_{1,2}} \quad (3)$$

Where C represents the plasma core line of sight and I₁, I₂, O₁ and O₂ correspond to the line integrals depicted in Fig. 2. Over bars indicate averages.




Figure 2. The SXR Factor is computed with the central, Inner and Outer line integrals of the vertical camera.

Finally, a total radiation time-trace has been calculated, using the metal foil bolometer arrays [27], namely 4 chords of the 24 belonging to the horizontal camera (see Fig. 3). Again, the idea is to pick differences between the core and the outer parts, in this case the Core (C1 and C2), Upper (U) and Lower (L) regions of the plasma. The BOLO Factor is calculated as:

$$\text{BOLO Factor} = (C1 + C2) - (L + U) \quad (4)$$




Figure 3. The “Bolometry Factor” is computed as (C1 + C2) - (U+L)

4.2. MARFE indicator for prevention

In tokamaks MARFEs [16][28] appear as a toroidal radiation belts that can move from the divertor to the upper part of the inner wall (or *vice versa*). Since they emit in the visible region of the spectrum, the visible cameras can be used to detect them, to quantify their intensity and to investigate their influence on the plasma stability. This is a different and rather unexplored solution to predict disruptions: indeed the most common approach is the direct analysis of bolometry signals.

JET has several operational cameras installed looking at the vacuum vessel. In this work we have used the one with a wide view of octant 8 (suitable for the visualization of radiation events as MARFEs). The idea is to identify Regions Of Interest (ROI), i.e. specific zones where the instability appears in the videos and constitutes a pre-disruptive pattern.

Each frame of the videos corresponds to a time slice of 50ms (this is the sampling rate of this camera). A ROI is a group of pixels, a small spot in the image. Since the mean intensity (brightness) of each ROI can be calculated, a time-trace can be computed. The 7 ROIs, which were initially selected for a first analysis, are shown in Fig. 4, while an example of these time-traces is reported in the right hand plot of Fig. 5.

Notice that these ROIs are distributed along the inner wall in the locations where (after a thorough visual inspection over tens of MARFEs) the instability appears most clearly.




Figure 4. Seven ROIs have been tested to detect MARFE instabilities. Their location along the inner wall from the divertor to the top of the vacuum vessel can capture MARFEs vertical movements and abnormal radiation events.

To determine the optimal size of the ROIs, a scan in pixel size has been performed. Different $n \times n$ dimensional ROIs have been tested (with $n=1,2,\dots,10$). A visual interface has been used to assess the influence of the ROI dimensions on the generated time-trace (see Fig. 5). At the right of the Figure, the time evolution of the mean brightness of the three ROIs (black, blue and red little squares in the left Figure) is depicted. Fast variations in these signals would be the evidence of a sudden brightness and therefore a possible indication of a radiative phenomenon.

As the result of this procedure of careful visual inspection of 50 MARFEs, it has been verified that big ROIs ($n>2$) generate time evolution signals without abrupt changes, because the averaged area is too large (in each ROI the average brightness is computed). Therefore, a sudden variation in the intensity of some pixels within the ROI is attenuated by the surrounding ones, masking valuable information.

On the other hand, ROIs with $n=1$ (just 1 pixel) are too instable. Even a small movement of the vessel or change in the background light would have a profound impact on the resulting time-trace.

The best trade-off consists of ROIs of dimension $n=2$, i.e. square ROIs of 2×2 pixels.

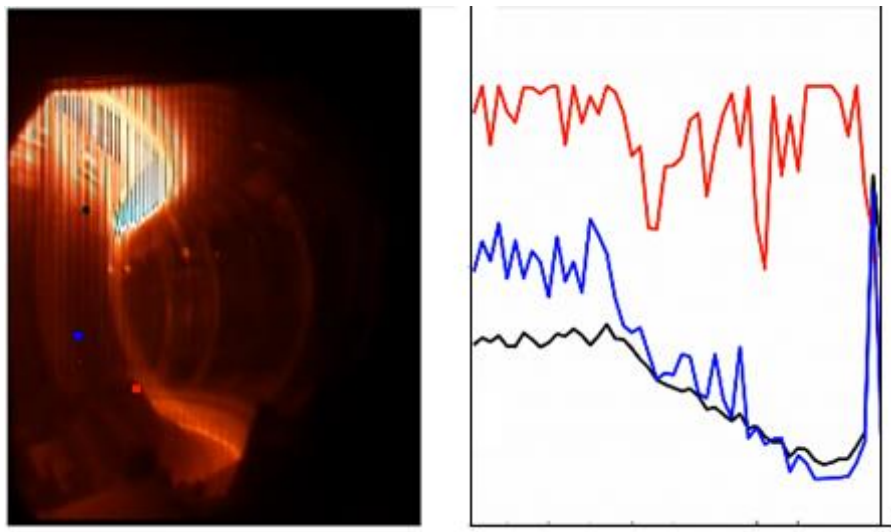


Figure 5. In the left figure, the ROIs sizes have been visually augmented in order to provide a better view of their location. The best ROI size consists of 2×2 pixels. In the plot on the right, three example time-traces corresponding to the three example ROIs (red, blue and black small

squares in the left figure) are shown. For each frame, the mean brightness values of these ROIs are computed to create the time-traces. Sudden variations in these processed signals may indicate a MARFE.

A comment regarding the time-traces computation is required. The underlying idea is to detect temporal changes in the luminosity of the ROIs. To this end, again Time Derivatives (TDs) of the signals have been used in order to capture these differences.

As a final remark, it is convenient to clarify that this method, even if it is MARFE-oriented, may capture other sudden radiation instabilities able to drive the pulse to a disruption.

4.3. Mitigation indicator

The strategy pursued in this work is to prioritize avoidance over mitigation. Consequently, the mitigation model has been designed to react to strong indicators of imminent disruptions with short warning times, allowing the avoidance predictors to operate and act first.

The mitigation predictor is developed following the same general procedure of the avoidance predictors (SVM optimized with GAs). In this case, the task is more straightforward, thanks to the availability of several previous works that successfully tackled mitigation. The crucial task of developing accurate indicators to be used as inputs of the model is not necessary since the main parameters to predict disruptions with short warning times have been identified in previous studies, as detailed for example in [22].

On the basis of past experience, therefore, the following signals have been included (as possible indicators) in the list of inputs for the mitigation predictor (the final set is selected by the GA optimization): 1. Locked mode amplitude; 2. Total input power; 3. Bolometric Factor (the same time-trace used in the Profile predictor); 4. Plasma vertical centroid position and 5. Plasma Current. Again, the time derivatives of these signals have also been considered for the training and validation of the models.

5. The predictors

The models of the three predictors have been obtained by deploying the GA routines to optimise them with the time-series of the various signals and of the devised indicators as inputs.

In the case of the Profile predictor for avoidance, the four computed profile Factors (1. Hollow Te, 2. Hollow ne, 3. SXR and 4. BOLO) and their time derivatives with $\Delta t : 2; 50$ ms are used to train and validate an SVM predictor. The process is optimized by the use of GAs (in the same fashion that is explained in Section 3 and detailed in the references therein,) leading to a model that include the following 8 time-traces:

1. Factor BOL_2 ; 2. Factor SXR; 3. FactorSXR₂; 4. FactorSXR₅₀; 5. Hollow Te;
6. Hollow ne; 7. Hollow ne₂ and 8. Hollow ne₅₀.

These 8 parameters are the only ones included in the Profile avoidance predictor.

To train the predictor, times slices containing these 8 parameters (also called “feature vectors”), at the time a pronounced Hollow profile is detected in 20 clear cases, have been used as pre-disruptive examples. As non-disruptive examples, 200 feature vectors from non-disruptive discharges and from Hollow profiles not leading the discharge to the disruption, have been chosen. Conceptually, what the system is shaped to do is only to trigger an alarm in case it detects a signature of a clear hollow Te and peaked ne and SXR profiles that will force the plasma to disrupt.

An analogue procedure has been carried out for the MARFE model. In this case, according to the GAs optimization, the only relevant ROIs selected by the GA are ROI 1 and ROI 2, both of them in the upper part of the inner wall, discarding the others as too unstable in terms of changes in the pixels’ luminosity (see Fig. 4). To train the predictor/hyper-plane able to separate pre-disruptive and non-disruptive time slices (each one containing the values of the two selected ROIs), it was crucial to identify clear and significant examples of MARFEs discharges that ended in a disruption. After a thorough trial and error procedure it has been evidenced that the best results can be attained with a reduced set of 12 very clear examples of MARFEs that lead the discharges towards the disruption (they were used as pre-disruptive examples). Non-disruptive examples are abundant and 200 of them have been extracted from non-disruptive shots and from pulses with MARFEs that do not drive the plasma towards a disruption.

Finally, the system, designed for mitigating the disruptions not identified by the avoidance and prevention predictors, has also been trained and validated using SVM in combination with GA. According to the GA, the optimal set of signals for the mitigation model are:

1. Plasma Current;
2. Plasma Current₂;
3. Factor BOL₂;
4. Total input Power;
5. 6. Total Input Power₂;
7. Locked Mode₁;
8. Plasma Vertical Position;
9. Plasma Vertical Position₂ and
10. Plasma Vertical Position₅₀;

In this case, the disruptive examples correspond to 2 ms before the disruption (aiming at triggering the alarm with very short anticipation) for 50 shots and 200 non-disruptive examples have been randomly chosen from non-disruptive shots.

Only shots of the training dataset have been used to create the three predictors. A completely independent set of shots, never used before, has been reserved for the testing database, the one used to evaluate the predictors and to get the results detailed in the next section.

6. Results

6.1. Performance of the avoidance and prevention models

The independent test dataset has been used to obtain the results reported in this section. To emulate real-time operational conditions, all three predictors have analysed each discharge from the beginning of the plasma current plateau until the current falls below 750 kA. In the case that any of them triggers an alarm, the analysis of that shot is stopped and the alarm is recorded.

It is important to note that the procedure suffers the risk of triggering a wide number of false alarms, because each discharge is monitored by three models, acting in parallel to find precursors (also, of course, in non-disruptive shots). For that reason, special care has been devoted to the training of the models, severely penalising false alarms in the GAs optimization.

Before describing the overall results, it is interesting to analyse the response of the avoidance models (the one for mitigation does not provides significant new findings compared to previous models as [22]). One representative example of a correct alarm triggered by the Profile avoidance predictor is shown in Fig. 6. Notice that the Te Factor increases (meaning that the temperature in the core is decaying) and also the ne Factor peaks (indicating a peaked density profile). This shot disrupts 920 ms after the alarm. The

yellow trace shows the outcome of this predictor. Taking into account that the predictor is a SVM model (basically, a high dimensional hyper-plane) this outcome is just the distance of the evaluated time slice to the hyper-plane. Negative values represent non disruptive behaviours. Positive ones mean that the sample under analysis is disruptive and an alarm has to be triggered. In the Fig. 6. example, as it is expected, the model's output increases as the Te profile grows hollow (red trace arising) and the ne factor becomes more peaked (black line).

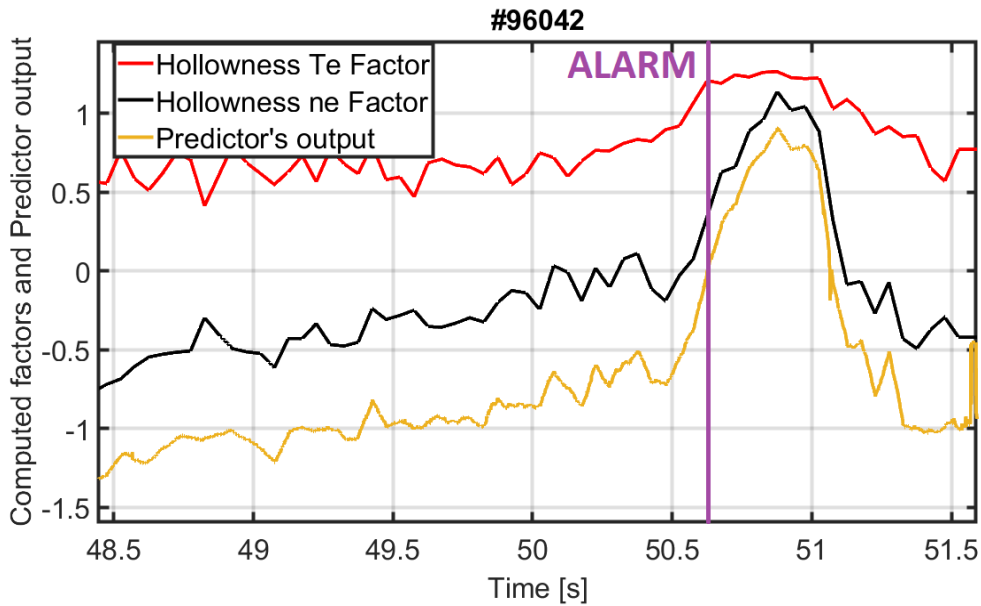


Figure 6. Correct alarm of the Profile predictor triggered 920 ms before the disruption. In this case, the Te profile becomes hollow due to impurities accumulation.

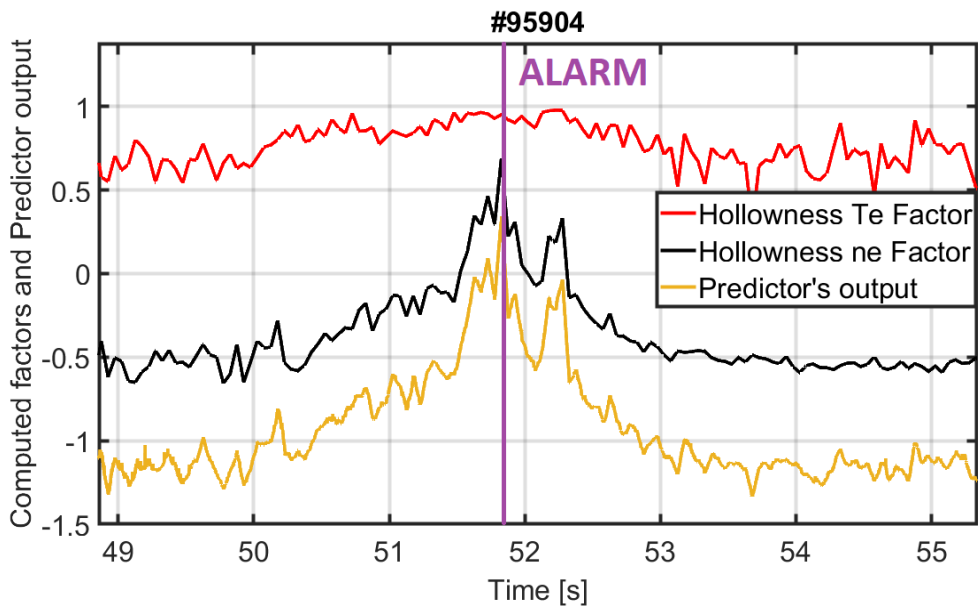


Figure 7. False alarm triggered at 51.709 s. This shot belongs to experiments about the isotope effect on H-mode detachment and density limit.

However, not all the predictions of this model are correct. In #95904 (Fig. 7) a discharge belonging to a density limit with a false alarm has been represented. This is one of the few false alarms triggered by the profile predictor in a very atypical evolution of the discharge, showing a pattern similar to the one of the previous Fig. 6 (disruptive). At the time of the alarm, the Te Factor is high (near to 1) and simultaneously the density is quite peaked, with a peaked ne Factor near 0,5. This is the most common situation in the 34 correct alarms correctly triggered by this avoidance system.

One typical correct alarm of the MARFE predictor for prevention is reported in Fig. 8 for shot #96201. Once again, the vertical line marks the time of the triggered alarm, more than 300 ms before the disruption. The pattern that identifies the precursor shows a significant difference in the pixels' intensities between the 2 ROIs (Zones 1 and 2 respectively; see Fig. 4 in Section 4.2). Note that after the firing of the alarm, the MARFE still keeps moving and the upper ROIs (Zone 1) show an oscillating behaviour.




Figure 8. Correct alarm activated by the MARFE predictor. The vertical line indicates the time when the model triggers the alarm (316 ms before the disruption). The differences in the pixels' intensities of the ROIs reach a maximum at the time of the alarm.

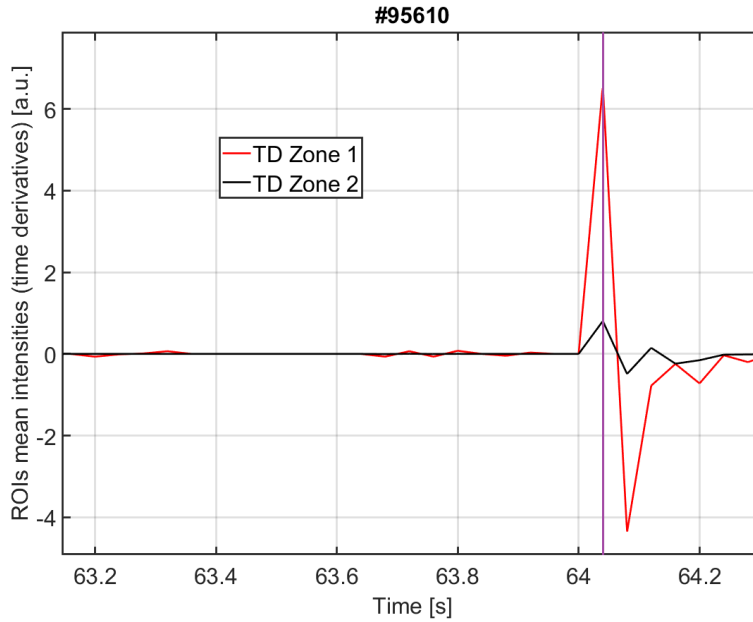


Figure 9. A false alarm is triggered due a sudden variation in the pixels' intensities caused by the injection of gas using the Disruption Mitigation Valves (DMV). In the experiment, the valves are not used to mitigate disruptions but for mitigation of the loads on the divertor tiles. The pattern is similar to a strong MARFE. In this case, this alarm might not be considered as false, taking into account the nature of the conducted experiment.

Something similar happens in the false alarm displayed in Fig. 9. There, discharge #95610, belonging to an experiment that involved the use of the Disruption Mitigation Valves (but not for disruption mitigation), shows the same pattern at 64,029 seconds and a (false) alarm is fired. This alarm has been classified as a false one, even if, due to the nature of the experiment and the gas injection, it is at least arguable to consider it as a mistake.

6.2 Overall results including the predictor for mitigation

The overall results are depicted in Fig. 10. This is a commonly used plot that shows the accumulated fraction of predicted disruptions (in percentage) versus the time to the beginning of the current quench (reported as time zero). The figure is quite informative, because it provides not only the total prediction rates of the models but also a general overview of the alarm anticipation times.

Fig. 10 shows that the vast majority of the first alarms with large warning times are triggered by the Profile model (green line). As the shots approach the disruption, in case the Profile predictor does not activate an alarm, the one more likely to intervene is the MARFE model. Finally, the mitigation model captures the most imminent indications of a disruption to start last resort measures, when the other models have failed to detect them.

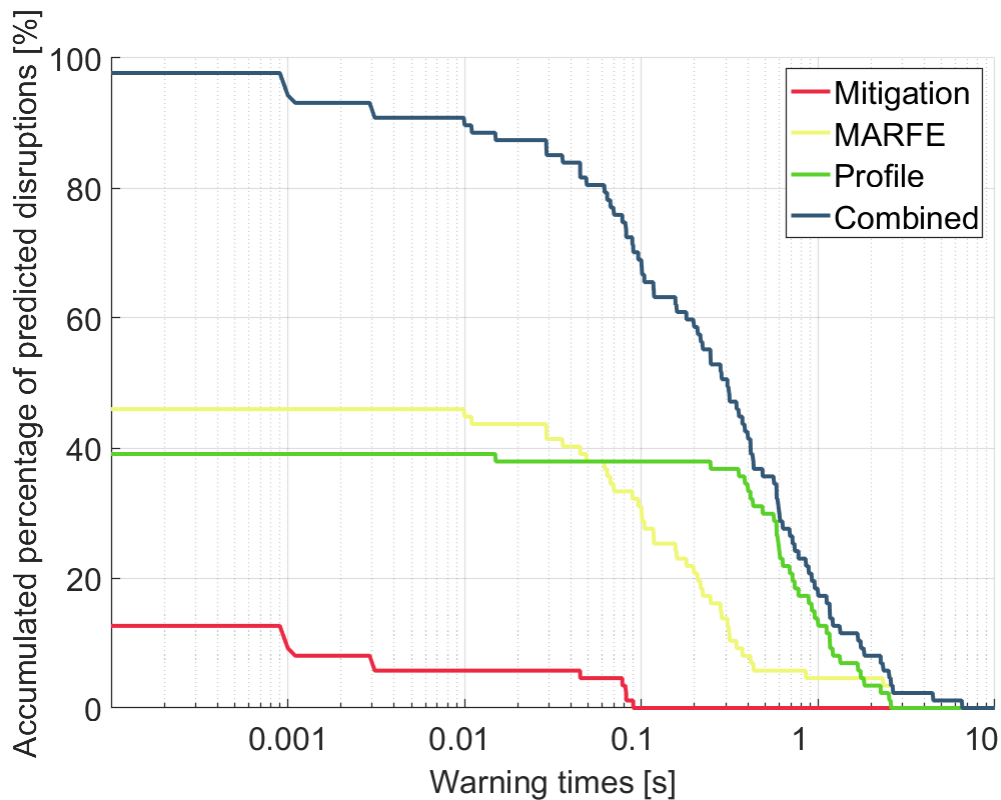


Figure 10. In this plot, only the first alarm has been plotted. The vast majority of the alarms with large warning times are triggered by the “Profile” (green line) predictor. The combined strategy reaches the 97,7% of predicted disruptions.

Another desirable outcome is the “flatness” of the green curve (for warning times between 0 and 400 ms). It means that almost no alarm is triggered by this predictor with less than ~400 ms of anticipation. This practically guarantees a minimum time (400 ms) to intervene. In addition, the yellow curve, reporting the warning time of the radiation/MARFE predictor for prevention, tends also to flatten out after 50 ms before the beginning of the current quench, showing a behaviour complementary to the other two classifiers. A deeper analysis of the alarms statistics is the subject of next section.

6.3. Statistical analysis

For a better understanding of the overall system response, a statistical analysis of the outcomes has been performed. Fig. 11 shows the distribution of warning times for the Profile avoidance predictor in the interval, in which it triggers the vast majority of the

alarms (between 400 and 2500 ms before the disruption). This distribution can be fitted to an exponential model of the form:

$$f(w) = f_0 \cdot \exp(-w/T) \quad (5)$$

where $f(w)$ is the fraction of detected disruptions with a warning time w , f_0 is the fraction of detected disruptions with positive warning times and T is the average warning time of the predictions. The resulting model parameters for this Profile predictor fit are: $f_0 = 58,79 \pm 0,14$ and $T = 675 \pm 13$ ms, where the estimations have been performed with 95% of confidence interval and the R-square factor of the fit is 0.987.

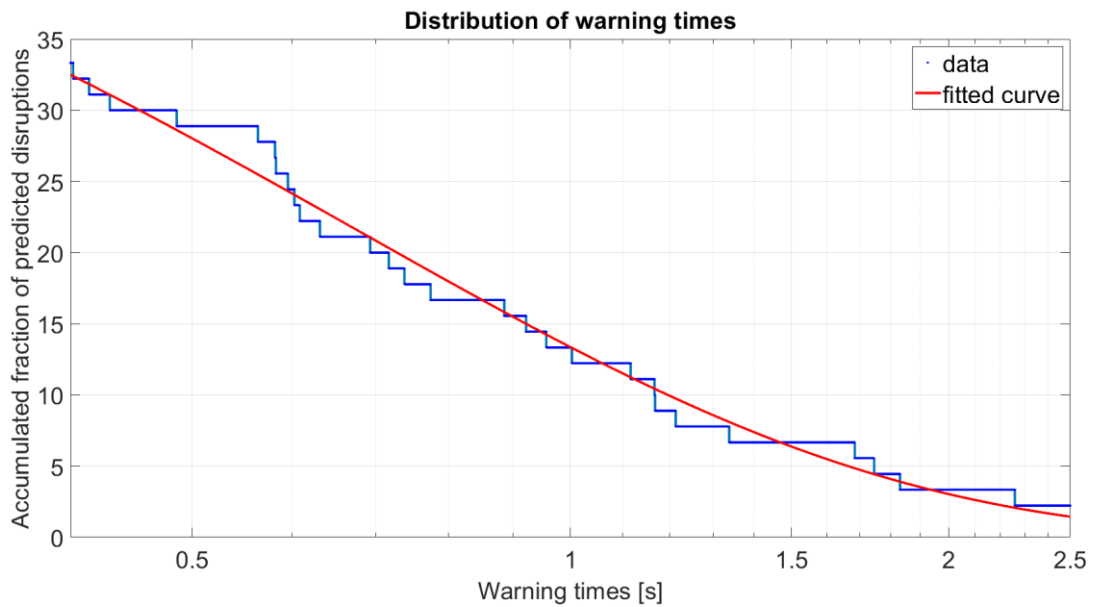


Figure 11. The distribution of warning times for the Profile predictor for avoidance follows an exponential model with an average warning time of 675 ms.




Figure 12. Exponential model for the MARFE predictor for prevention.

The MARFE predictor has also been fitted with an exponential model (Fig. 12), again in the interval, in which it triggers most of the alarms (between 30 and 430 ms before the disruption); its parameters are $f_0 = 45.85 \pm 0.09$ and $T = 230 \pm 15$ ms, where the estimations have been performed with 95% of confidence interval and the R-square factor of the fit is 0.99.

A simple and clear representation of the statistics derived from the exponential models is shown in Figure 13. There, the mean warning times of the predictors (\pm their standard deviations) are represented in different colours. Notice that the Profile model barely overlaps with the MARFE's. This in another prove of what stated before: by design, the three predictors are able to act in cascade, prioritizing the intervention of the avoidance and prevention models. In case these two miss an alarm, the mitigation model intervenes to trigger measures of last resort.

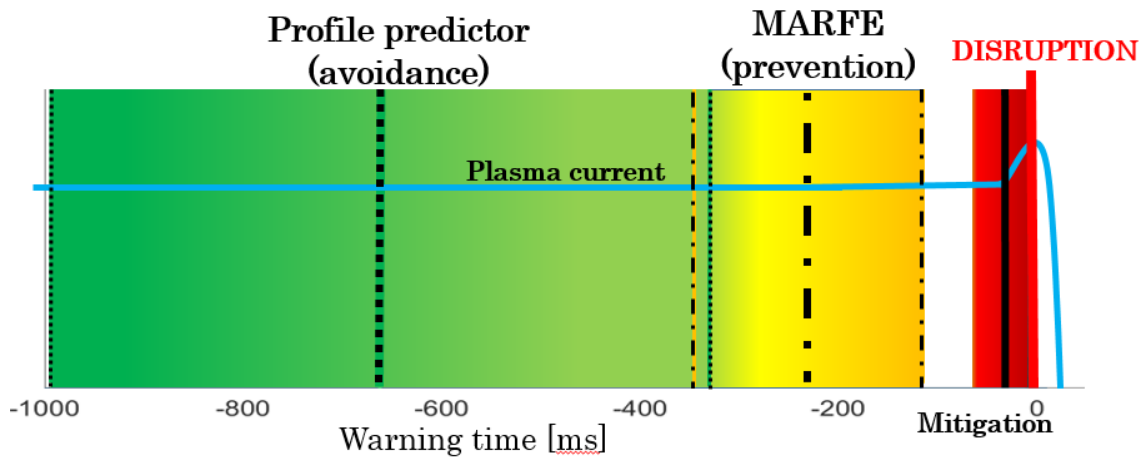


Figure 13. Mean times and standard deviations of the three predictors, according to the statistical analysis. It's remarkable that the anticipation times barely overlap. Disruptions not anticipated by the avoidance and prevention model can be detected by the mitigation predictor.

Another important summary statistics refers the distribution of the alarms, reported in Table 1. There, it is possible to see how the Profile predictor detects 34 disruptions (representing 40% of the total alarms triggered), with a mean warning time of 675ms. In the case of the MARFE predictor, it is the first one detecting precursors in 40 cases (47% of the total triggered) with a mean warning time of 230 ms.

Predictor	Profile	MARFE	Mitigation
Mean	0,675	0,23	0,028
# of alarms	34	40	11
Percentage of the total triggered	40%	47%	13%
Overall results			
	False alarms 4,7%*	Missed (tardy) alarms 2,3%	Overall success rate 97,7%
	*75% of them before/after fast stop		

Revised stats	1,175%	2,3%	97.7%
---------------	--------	------	-------

Table 1. This table summarizes the results for the test database, with 87 disruptive discharges and 85 correct alarms (and 2 tardy alarms, triggered ~2ms after the disruption). The false alarms, after a revision, could be considered a 1,175% since the models triggered an alarm as consequence of a fast stop (after it) or predicting an anomaly that made the JET protection system intervene (before a fast stop).

The obtained results support the interpretation of the predictors as being optimised for different phases of the discharges, with respect to the beginning of the current quench.

The profile predictor would be the most adequate for avoidance, since it triggers alarms earlier and guarantees that there are always at least 400 ms to the beginning of the current quench after its warning. The classifier based on the visible camera, intervenes just after the profile predictor, with practically no overlap and could be considered a good candidate for prevention, since it provides an intermediate anticipation time. The last predictor springs into action only when the disruption is imminent and seems therefore to be the most suited to mitigation. Again, it is relevant to mention that the trained system not only fires alarms in sequence but provides also a first classification of the disruption types from an operational standpoint.

7. Summary, discussion and future developments

Disruptions are a priority research subject in sight of next step devices such as ITER and DEMO. Up to now, very few works can be found in the literature able to deal with their early prediction for avoidance and/or prevention. There still are too many loose ends, especially due to the complexity of detecting very subtle precursors without triggering a large number of false alarms. Even more, it is not only relevant to have a trigger, but also it is necessary to have in real time at least a hint of the remaining time available to intervene once an alarm is fired. Moreover, identifying the main operational cause of the incoming disruption is also very important to optimise the countermeasures.

These issues are efficiently addressed in this work with a prediction strategy instead of with a single predictor. Three models, each one of them targeting a specific disruption phase, have been trained. The first one, aimed at avoidance, uses HRTS measures to detect Hollow Te and peaked ne profiles that end in a disruption. The high success rate

of this classifier is achieved also by the inclusion of a SXR and bolometry data, to help reducing false alarms in shots with hollow Te and peaked ne profiles but not ending in disruption.

The second predictor, aimed at prevention, has been built with the innovative approach of using only videos from an operational camera, to follow the brightness of ROIs, with the objective of capturing radiation anomalies that can drive the pulse to a premature end. After the training, an unexpected result has emerged: the model only requires two small ROIs (ROI 1 and ROI 2), located in the upper part of the inner wall, to detect MARFEs and other radiative precursors of disruptions.

The third and last predictor is designed to capture those alarms not triggered by the avoidance and prevention models. The point is to use this mitigation predictor as a last resort trigger to inject gas or pellets.

The total database analysed in this work contains 977 discharges. The models are created with a database of 646 pulses (176 disruptive, 470 non-disruptive) and tested with a completely independent set of 328 discharges (87 disruptive and 241 non-disruptive). The results have therefore a quite substantial statistical basis.

The overall rates are a 97,7% of predicted disruptions. This means that whole system only misses 2 of the 87 disruptions in the test database. In these 2 cases, the mitigation system triggers the alarm only ~2ms after the beginning of the disruption, which is not considered a major issue since the current quench lasts hundreds of ms, leaving plenty of time to undertake mitigation actions.

The false alarms are less than the 5%. Even more, a deeper analysis reveals that the 50% of them are activated after JET control system intervened with a fast stop. In these cases, the cause of the alarm is the detection of these abrupt control actions. The other 50% of the false alarms are activated as consequence of detecting severe pathologies in ill plasmas that finally were terminated by the control system without disruptions.

In many case, the disruptions in the JET database are caused by the DMVs activation triggered by the JET control system. This means that the timing of the disruption is determined by the active mitigation trigger and not by the time scales of a natural disruption. But in any case, the avoidance and prevention predictors detect that something is wrong with the plasma behaviour before any intervention by the control system. Specifically, the avoidance predictor and the prevention model trigger 89% and 75% of their alarms before the first fast stop respectively. Moreover, both the prevention and

avoidance models trigger a very low rate of false alarms. A legitimate interpretation of this evidence is that: a) the predictors learn quite well the natural evolution of the plasmas (very low level of false alarms when JET control system does not take any action) b) the number of unnecessary plasma terminations caused by JET control systems is quite low (otherwise the avoidance predictor would tend to trigger an alarm after a stop much more often).

This indicate that: 1) every time they detect a pre-disruptive signature, a severe anomaly will posteriorly appear and the discharge will be driven towards a forced landing; 2) it strongly suggests that the detected pre-disruptive signatures are taking place in shots that are whether going to disrupt or ill enough to force the posterior intervention of the control system.

The most important aspect of this work is related to the distribution of the alarms and their anticipation times. The Profile predictor activates the first alarm in 34 instances (40% of the total triggered) with a mean anticipation of 675 ms. The MARFE/radiation predictor is the first one firing the alarm in 40 cases (47% of the total triggered alarms), with a mean warning time of 230 ms. Finally, only 10 alarms (13% of the total triggered) are not detected by previous avoidance models and they are activated by the mitigation system with a mean time of 28 ms. The overlap of the alarm times of the three predictors is minimal; the information provided is therefore quite detailed and would allow a control system to optimise the remedial countermeasures much better than the present generation of tools.

It worth mentioning that the conclusions reached in the present work are in good agreement with a very recent publication that also links, among other parameters, the T_e and n_e profiles to disruptions [30]. The cited article describes how hollow T_e profiles in the core and edge cooling increase the probability of destabilizing 2/1 tearing modes. The main time scales are also basically the same, with the hollowing of the temperature profiles in the core being followed by the current much later than the cases of disruptions due to edge cooling. On the other hand, the present work is more orientated toward real time prediction, whereas [30] is more focussed on explaining the physical mechanisms leading to the disruptions.

Instead, the work detailed here has been developed and test considering its possible on-line application. It may represent a significant breakthrough in the field. For the first time not only an alarm is fired but also, depending on the model that triggered the alarm,

the phase of the discharge and the mean warning time are known. Moreover, the first systematic use of videos has shown the great potential of this class of diagnostics for contributing to addressing the issue of disruptions. On JET, unfortunately, the available actuators probably would not be versatile enough to take full advantage of such a system of predictors but the obtained performance are very promising and could constitute also a good basis for the design of diagnostics and actuators on the next generation of machines.

On the other hand, it should be recognised that the pursued approach to training requires many data, which cannot be necessary available at the beginning of operation of new devices. The methodology proposed in this work can therefore be considered complementary to the open world training approaches recently developed. From scratch predictors [8], adaptive learning [7] and transfer learning [4] are techniques that have proved to be very effective. They can start operating with a minimum of examples, can follow the evolution of the operational programme and even effectively be transferred from one device to another [7] [31]. However, they have always been plagued by a very large spread in the warning times, never managing to provide an estimate of the time remaining before the beginning of the current quench. The two approaches could be therefore profitably combined. At the beginning of operation of a new device, the first tools deployed could be those based on open world learning. Then, once enough examples have been gathered, the ones based on GA optimization could take over. It should be mentioned that the techniques of adaptive learning could also help in the selection of the most suitable examples to be used for GA based methods.

Again regarding future improvements, the most interesting next step could be to train other different predictors targeting specific causes or types of disruptions to include them in the main integrated system [31][32]. This needs to be accompanied by a thorough comparison with the triggers issued by the JET protection system. Also, it would be convenient to replicate the methodology in other tokamaks in order to evaluate the transferability of the strategy in the perspective of next step devices.

Acknowledgments

This work was partially funded by the Spanish Ministry of Economy and Competitiveness under the Project PID2019-108377RB-C31.

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014–2018 and 2019–2020 under Grant agreement No. 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

- [1] Lehnen, Michael, et al. "Disruptions in ITER and strategies for their control and mitigation." *Journal of Nuclear materials* 463 (2015): 39-48.
- [2] Chen, Francis. *An indispensable truth: how fusion power can save the planet*. Springer Science & Business Media, 2011.
- [3] De Vries, P. et al. "Statistical analysis of disruptions in JET". *Nuclear Fusion*. 49. 055011. 10.1088/0029-5515/49/5/055011. 2009.
- [4] Murari, A. et al. *Nucl. Fusion* 60 056003 (18pp). "On the transfer of adaptive predictors between different devices for both mitigation and prevention of disruptions". 2020.
- [5] Rattá, G. A., Vega, J. A., Murari, A., & Vagliasindi, G. "Inspection of disruptive behaviours at JET using Generative Topographic mapping". In *From Physics To Control Through An Emergent View* (pp. 315-320). 2010.
- [6] Pau, A. et al. *Nucl. Fusion* 59 106017. 2019.
- [7] Vega, J., et al. "A linear equation based on signal increments to predict disruptive behaviours and the time to disruption on JET." *Nuclear Fusion* 60.2 (2019): 026001.
- [8] S. Dormido-Canto, et al. *Nuclear Fusion*, volume 53 (2013) 113001 (8pp).
- [9] Esquembri, S. et al. "Real-Time Implementation in JET of the SPAD Disruption Predictor Using MARTE". In *IEEE TRANSACTIONS ON NUCLEAR SCIENCE*, volume 65, 2018.
- [10] Vega, J. et al. "Adaptive high learning rate probabilistic disruption predictors from scratch for the next generation of tokamaks". *Nuclear Fusion*. 54 (2014) 123001 (17pp).
- [11] Rattá, G. A. et al. "An advanced disruption predictor for JET tested in a simulated real-time environment." *Nuclear Fusion*, 50(2), 025005. 2010.
- [12] Vega, J. et al. "Results of the JET real-time disruption predictor in the ITER-like wall campaigns." *Fusion Engineering and Design* 88.6-8: 1228-1231. 2013.
- [13] Vega, J. et al. "Advanced disruption predictor based on the locked mode signal: application to JET". 1st EPS Conference on Plasma Diagnostics. April 14-17, 2015. Book of abstracts. Frascati, Italy.
- [14] Vega, J. et al. "Disruption Precursor Detection: Combining the Time and Frequency Domains". SOFE program. 26th Symposium on Fusion Engineering (SOFE 2015). May 31st-June 4th, 2015. Austin (TX), USA.
- [15] Pasqualotto, R. et al. "High resolution thomson scattering for Joint European Torus (JET). Review of Scientific Instruments, 75(10), 3891-3893. 2004.
- [16] Lipschultz, B. J. *Nucl. Mater.* 145 15. 1987.
- [17] Cortes, C. and Vapnik, V. "Support vector networks." *Mach. Learn.*, 20, 273–293. 1995.
- [18] Hertz, J. et al. *Introduction to the theory of neural computation*. Addison-Wesley. 1991.

- [19] de Jong, K.A. et al. Parallel Problem Solving from Nature. LNCS, vol. 496, pp. 38-47. Springer, Heidelberg. 1991.
- [20] Rattá, G. A. et al. "Improved feature selection based on genetic algorithms for real time disruption prediction on JET." *fusion Engineering and Design* 87.9 : 1670-1678. 2012.
- [21] Rattá, G. A., Vega, J., Murari, A., & JET Contributors. A multidimensional linear model for disruption prediction in JET. *Fusion Engineering and Design*, 146, 2393-2396. 2019.
- [22] Rattá, G. A., et al. "Global optimization driven by genetic algorithms for disruption predictors based on APODIS architecture." *Fusion engineering and design* 112: 1014-1018. 2016.
- [23] Lopez, J.M. et al., Integration and Validation of a Disruption Predictor Simulator in JET, *Fusion Science and Technology* 63.1 (2013): 26-33.
- [24] Fontana, M., et al. "Real-time applications of Electron Cyclotron Emission interferometry for disruption avoidance during the plasma current ramp-up phase at JET." *Fusion Engineering and Design* 161 (2020): 111934.
- [25] Czarnecka, A., et al. Analysis of metallic impurity content by means of VUV and SXR diagnostics in the presence of ICRF induced hot-spot on the JET-ILW poloidal limiter. *Review of Scientific Instruments* x, 2018.
- [26] Sertoli, Marco, et al. Determination of 2D poloidal maps of the intrinsic W density for transport studies in JET-ILW. *Review of Scientific Instruments*, vol. 89, no 11, p. 113501. 2018.
- [27] Huber, A. et al. Upgraded bolometer system on JET for improved radiation measurements. *Fusion Engineering and Design*, 82(5-14), 1327-1334. 2007.
- [28] Chanin, A. V. On the poloidal localization and stability of multi-faceted asymmetric radiation from the edge (MARFE). *Physics of Plasmas*, 11(4), 1484-1492. 2004.
- [29] A.Murari et al *Nucl. Fusion* 59 (2019) "Adaptive learning for disruption prediction in non-stationary conditions" 086037 (11pp) <https://doi.org/10.1088/1741-4326/ab1ecc>
- [30] Pucella, Gianluca, et al. "Onset of tearing modes in plasma termination on JET: the role of temperature hollowing and edge cooling." *Nuclear Fusion* (2021).
- [31] Rattá, G. A., J. Vega, and A. Murari. "Viability assessment of a cross-tokamak AUG-JET disruption predictor." *Fusion Science and Technology* 74.1-2 (2018): 13-22.
- [32] A.Murari, et al "Stacking of Predictors for the Automatic Classification of Disruption Types in Support to the Control Logic" 2021 *Nucl. Fusion* 61 036027 DOI <https://doi.org/10.1088/1741-4326/abc9f3>