Adaptive Learning for Disruption Prediction in Non-Stationary Conditions

EUROfusion Consortium, JET, Culham Science Centre, Abingdon, OX14 3DB, UK

by A.Murari[1], M.Lungaroni[2], M.Gelfusa[2], E.Peluso[2], J.Vega[3] and JET Contributors*

1) Consorzio RFX (CNR, ENEA, INFN, Universita' di Padova, Acciaierie Venete SpA), Corso Stati Uniti 4, 35127 Padova, Italy.

2) Department of Industrial Engineering, University of Rome "Tor Vergata", via del Politecnico 1, Roma, Italy

3) Laboratorio Nacional de Fusión, CIEMAT. Av. Complutense 40. 28040 Madrid. Spain

* See the author list of "X. Litaudon et al 2017 Nucl. Fusion 57 102001

Abstract

In the last decade, machine learning tools have proved to be very powerful disruption predictors in Tokamaks. On the other hand, the vast majority of the techniques deployed assume that the input data are independent and are sampled from exactly the same probability distribution for the training set, the test set and the final real time deployment. This hypothesis is certainly not verified in practice, since the experimental programmes evolve quite rapidly, resulting typically in ageing of the predictors and consequent suboptimal performance. This paper describes various adaptive training strategies that have been tested to maintain the performance of disruption predictors in non-stationary conditions. The proposed approaches have been implemented using new ensembles of classifiers, explicitly developed for the present application. The improvements in performance are unquestionable and the final predictors meet the needs of the next generation of experimental machines, such as ITER. Given the difficulties encountered so far in translating predictors from one device to another, the proposed adaptive methods *from scratch* can therefore be considered very good candidates for the next generation of devices, particularly at the very beginning of their operation.

Keywords: Disruptions, Machine Learning Predictors, Adaptive training, De-learning, Obsolescence, Ensembles of Classifiers

Corresponding author: michele.lungaroni@uniroma2.it

1 Predicting disruptions in non-stationary conditions: present and future devices

For about two decades, machine learning tools have been applied to the task of predicting disruptions in Tokamaks [1]. The main motivations behind these efforts have been the potential damage of disruptions to the devices and the lack of physical models even vaguely suitable for prediction. The results obtained so far are quite impressive but not entirely satisfactory. For example, most of the predictors tend to age quickly, providing very soon suboptimal performance. Moreover, the transfer of predictors from a device to another has proved to be very problematic [2]. This is worrying since the gap between the present machines and ITER or DEMO is quite large [3].

In almost all applications of machine learning to disruption prediction, the problem is formulated as a classification task. For every time slice, the predictors have to decide whether the plasma belongs to the class of the safe discharges or of the disruptive ones. Except for a few cases [4-6], all the classifiers deployed on Tokamaks assume that the data are independent and identically distributed (*i.i.d.*). This hypothesis means that the inputs are expected to be sampled independently from the same probability distribution function for both the training set, the test set and the final real life application. Physically, the *i.i.d.* assumption implies that the properties and the results of the predictors are guaranteed only in steady state conditions, not only during discharges but also between them. These assumptions are certainly violated in practice. There are indeed strong correlation effects not only between subsequent time slices but also between discharges (see Section 3 and 4).

The fact that the *i.i.d.* conditions are not met typically results in predictors, which age very quickly and provide suboptimal performance. Indeed, as occurred very clearly on JET in the last years, the experimental programmes nowadays evolve rapidly, creating often new situations different from the ones used to train the predictors. These changes in the plasma operation have more severe repercussions with the ITER Like Wall, due to the difficulty of predicting the effects of heavy impurities. Quite delicate is the control of tungsten influxes, since W can easily cause radiative collapses particularly in the ramp down phase of the plasma current. It should also be mentioned that the robustness of disruption predictors performance, which is already a severe challenge on JET, will be even more important on ITER, particularly at the beginning of operation, when not enough cases will be available for traditional batch training and the experiments will have to explore a large operational space in limited time. Indeed it should be remembered that, even if a lot of examples are available, the translation of predictors from one device to another has always proved to be a very difficult task so far; therefore predictors trained with data of smaller devices could not meet the specifications of larger machines. For all these reasons, adaptive predictors, which require a minimum number of examples to be trained and can follow the evolution of the experimental programme, would be very useful.

In this paper, various strategies for adaptive training of disruption predictors are presented and their performance assessed with a large database of JET discharges with the ITER Like Wall (ILW). The approaches developed are all compatible with real time applications; they also implement a *from scratch* training, which means that a minimum number of examples, in principle one safe and one disruptive shot, is enough for the predictors to start working with acceptable performance. The predictors have been optimised for mitigation and their application to avoidance will be addressed in a future publication.

The paper is structured as follows. Next section is devoted to the introduction of the main machine learning tools used to obtain the results reported in the paper: Classification and Regression Trees (CART) and their ensembles [7,8]. The details about the main strategies of adaptive learning, to be deployed between shots, are provided in Section 3 The approach implemented to take into account the memory effects during discharges is described in Section 4. The large database of discharges analysed is described in Section 5. The results in terms of success rates and false alarms are reported in Sections 6, 7, 8 and 9. The conclusions and lines of future investigations are the subject of the last section of the paper.

2 Classification trees and ensembles

As already mentioned, disruption forecasting is considered a classification problem; for each time slice, the predictors have to decide whether the plasma conditions are such that the plasma is going to disrupt or not. The task to solve can therefore be formalised as follows. Given a training sample of n observations, the class variable is indicated by Y and can in general take a finite set of discrete values 1, 2, ... , k. In our application, the predictors have to discern between safe and disruptive discharges and therefore the classification is binary; however, the tools are fully general and can be adapted easily to multiclass problems, such as the classification of disruption types. The set of p features used as predictor variables are indicated by $X_1,..., X_p$. The goal of the training consists of identifying a model, which can predict the class Y from new X values. As will become apparent in the rest of the paper, one of the main advantages of the adaptive approaches, proposed in the following, is the minimisation of the number of classified examples required for training the predictors.

The adaptive learning strategies proposed in this work are very general. The machine learning tools deployed have therefore to be sufficiently powerful and flexible to allow exploiting the potential of the adaptive procedures to the full. To this end, the methods, adopted to implement the various strategies of adaptive learning described in the next section, belong to the class of rule-based machine learning (RBML). RBML are machine learning methods that derive "rules" for solving a problem directly from the data available.

In more detail, the basic technologies implemented for the studies reported in this paper are Classification And Regression Trees (CART). CART extract knowledge directly from the available databases and formulate it as trees, which are their final output [10]. Their approach to classification consists of recursively partitioning the data space and then applying to each partition the best rule to fit the data. The algorithm is iterated until convergence to a solution of acceptable performance, quantified in terms of misclassification costs. The final classifier is therefore a series of rules that have the same representational power as propositional logic.

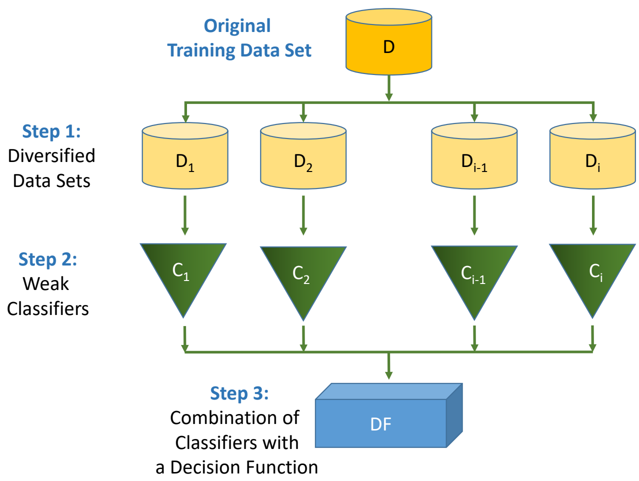

Figure 1. The typical architecture of ensemble classifiers.

RBML classifiers of the CART family present various advantages, such as flexibility and interpretability. On the other hand, they are affected by a quite severe weakness; excessive sensitivity to the details of the training set. Indeed, the CART models have the unfortunate tendency not to be very stable, in the sense that small variations in the examples of the training set can affect excessively the final trees and therefore the final classification. A typical approach to alleviate this difficulty is based on the so called "weak learners". These are not extremely successful classifiers that, on the other hand, are computationally not too heavy to train and deploy. Many of these weak classifiers are therefore pooled in so called *ensembles* and their predictions are then combined to obtain more stable and performing results (see Figure 1). The essential element in this combination of classifiers is to achieve enough diversity in the classifiers of the ensembles. Various alternatives are possible and two of the most popular are Bagging and Random Forests. They reduce the high-variance of CART trees by training many individual classifiers with slightly different sets of examples. A CART model is derived for each of the subsets and, for each new example to classify, the predictions of the various models are aggregated with a specific decision function, which implements one of many possible alternatives. Overfitting the data by individual trees becomes a much less serious issue for the members of the ensembles, because the resulting high variance is remedied by the aggregation implemented by the decision function. On the other hand, the individual trees are so flexible that they present typically a low bias, a property, which is then inherited by the final predictor.

In the case of Bagging, variability is introduced by randomly subsampling the original database with replacement. Random Forests (RF), or random decision forests, introduce variability not only by sampling the set of examples but also by implementing a random choice of the variables in the database. An original complement to these strategies is based on the observation that the measurements in Tokamaks are affected by many sources of significant noise, which are additive and independent. From a statistical point of view, this situation satisfies the validity conditions of the central limit theorem. It can therefore be assumed that the statistical distribution of the noise is Gaussian of zero mean. This fact suggests building different sets of data, for training the predictors in the ensembles, by summing random noise to the measurements. The natural choice for the statistics of the noise is a Gaussian distribution with variance equal to the error bar of the measurements. This method will be indicated with the name Noised-based Ensemble in the following sections. To improve the variability of the training sets, it has proved advantageous to apply the noise before building the Bagging and Random Forests ensembles, as detailed in Section 6. It should be noted that systematic errors in the measurements cannot be handled in the same way as random ones and should be addressed at the level of the individual diagnostic. On the other hand, if they remain consistent, systematic errors may have a very low impact on the performance of predictors, even if course can strongly impact the interpretation of the physics.

3 Adaptive training strategies between discharges

Adaptive methods for learning in non-stationary environments basically adopt one of two main strategies: active or passive training. In active training, the predictors actively try to determine the shift in the conditions and react accordingly. Passive approaches periodically update the models independently from the detection of changes in the system. Both passive and active approaches provide up-to-date models and can be very effective. There are no principled reasons to prefer a priori one method or the other. Moreover, in theory, combining the two approaches could be advantageous. Unfortunately, passive strategies for training are not viable in the present application at least for two reasons. First, the resulting training set would be too imbalanced, since in normal operating conditions the disruptive discharges are a small fraction of the safe ones. Moreover, the computational time scales quite badly with the number of examples in the training set and therefore there is also the need not to widen the training set unnecessarily. In the following therefore only active learning methods are considered and investigated.

In the context of updating the predictors between shots, the adaptive strategies proposed in the paper can be conceptually divided in two main categories. The first one consists of techniques, which react to individual errors and changes in the performance of the predictors. These approaches mainly consist of a suitable updates of the training set, as described in detail in subsection 3.1. The second class comprises strategies, which consider the entire history of the predictor performance up to the present time and act on the decision function and the training set, as discussed in subsections 3.2 and 3.3.

3.1 Basic adaptive learning: retraining after errors

This first part of the adaptive training is meant to react to errors and in particular to missed or tardy alarms. In this perspective, the last trained predictor, starting with the one obtained after the first training "*from scratch*", is deployed on the following discharges until the first missed or tardy alarm. When the predictor misses a disruption or gives a tardy alarm, the shot not properly classified is included in the training set. In this way a new model is derived, which is deployed to analyse the following discharges until the next error, which provides an example for a new retraining. Therefore, as far as the disruptions are concerned, the models are refined every time they commit an error, by missing an event or providing a late alarm.

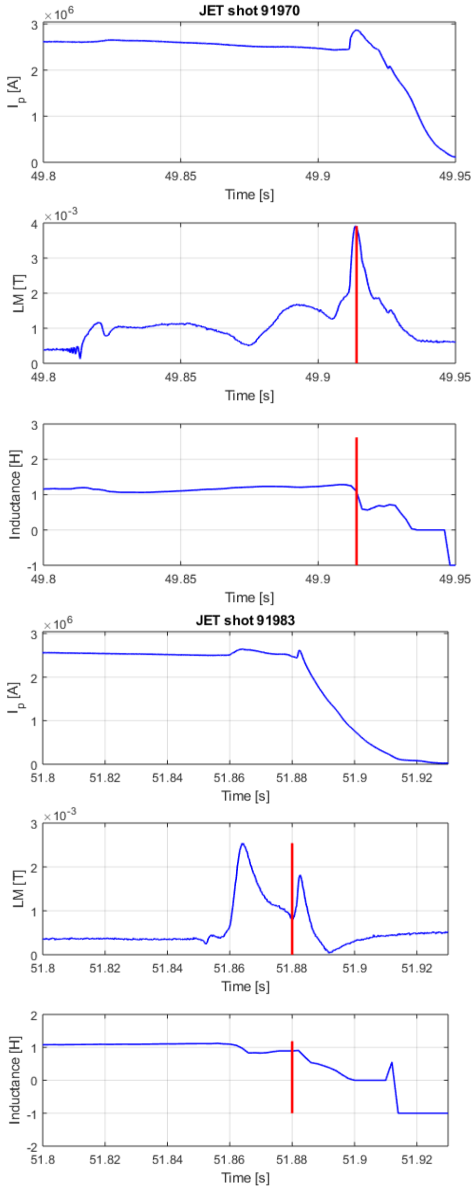

Figure 2.Examples of the locked mode time evolution. Top; disruption preceded by increasing mode locked amplitude. Bottom: a discharge in which the disruption occurs during the decaying phase of the locked mode signal. A predictor based only on the amplitude of the LM would be certainly suboptimal.

The case of the false positives, the false alarms, is more delicate. In closed loop operation, indeed, it is not necessarily the case that false alarms can be always recognised, after the discharge has been terminated prematurely or, in any case, interfered with following the false alarm. Of course off line analysis investigations can indicate that the discharge was shut down without any good reason, but this cannot be guaranteed to take place systematically. On the other hand, retraining based only on disruptive examples typically results in an unacceptable level of false positives. Such an issue becomes particularly evident during long campaigns and/or when the scenarios evolve rapidly, exploring completely new regions of the operational space. This has become a quite severe problem in the last years, particularly since the installation of the ILW. On JET, the experimental programme tends to vary quite quickly and, at the same time, the metallic wall renders the configurations more sensitive to small changes in the operational parameters. The predictors need therefore to adapt rapidly to the variations of the experimental programme. On the other hand, as already mentioned, passive training, after each safe discharge, is not a viable option, at least because it would lead to a training set too unbalanced. Therefore a more targeted criterion has to be devised to include new safe cases. The approach tested is based on the observation that one should include in the training sets the safe shot that precede a disruption. Typically, subsequent discharges are indeed quite similar since the exploration of the operational space is performed with great caution. Therefore an efficient way to update safe shots in the training set consists of adding the safe shot preceding any alarm triggered by the predictor. In this way, the models can follow the evolution of the safe operational space reducing the number of false alarms. The details of the algorithms to update the training sets as just discussed are provided in Section 6. In case of more abrupt changes in the sequence of experiments, of course additional automatic checks on the plasma configuration (such as current, field additional heating etc.) could be performed. Even manual updates of the training set could be envisaged but none of these additional precautions have proved necessary for the campaigns investigated in the present work.

3.2 Updating the decision function

The adaptive training strategies described in the previous subsection are based on the idea of updating the training set once a classification error has occurred or an alarm has been triggered. Alternative methods can try to take advantage of the entire history of the predictor performance. To this end, 6 different decision functions have been implemented, to decide whether a certain time slice is disruptive or safe, on the basis of the outputs of the individual weak predictors described in Section 2. All these predictors are then run in parallel even if only one can be in feedback loop at any point in time. Between shots, the track records of the various decision functions since the beginning of the campaign, or from any suitable starting point, are compared. The most performing prediction function, according to the chosen metric (see Section 7), is the one to be selected and included in feedback loop for the next discharge. The process is then repeated after each new discharge. In the case of the offline analysis presented in this paper, the preferred predictor at any point in time is the one used to calculate the global performance, since it is the one which would have been in feedback loop. Section 7 provides the essential details about the decision functions implemented and the results obtained.

3.3 De-learning: discarding old examples

The adaptive strategies, introduced in subsection 3.1 result in a continuous increase in the training data set. The switching from one decision function to another, as described in subsection 3.2, does not reduce the training set either. In addition to posing computational issues in the long term, this situation can lead to non-optimal results, because old and obsolete examples can remain in the training set. A possible improvement would consist of discarding too old examples. The main issue is the identification of the shot number from which the examples become obsolete. To this end, suitable performance indicators have been defined to determine automatically how to update the training set without retaining too old examples. The details about this point are provided in Section 9.

4 Handling memory effects during discharges: trajectory learning

As mentioned in Section 1, the plasma status at a certain point in time is not completely independent from the previous history. Therefore, the features sampled at various time slices are strongly correlated and in particular depend on the past evolution of the discharge. This is a clear violation of the i.i.d. hypothesis. In the case of disruption prediction, this situation renders predictors based on the simple absolute value of a measured quantity, such as the locked mode, clearly suboptimal. A relevant example is reported in Figure 2, where it is shown how disruptions can occur during phases when the locked mode amplitude is either increasing or decreasing. Unfortunately, during the discharges there is not enough time to implement complex active learning strategies, which are on the contrary deployed between shots (see Section 3). The time scales of plasma disruptive instabilities are in the ms range and indeed the cycle time of JET real time system is 2 ms. The only viable alternative is therefore "trajectory learning" [11]; the approach consists of training the predictors with series of subsequent time slices so that they can learn these correlation effects occurring during the discharge. Time series of various lengths have been tested for the case of the analysed data base as described in detail in Section 8. The proposed solutions extend the approach already pioneered in [12]. In any case, it should be emphasized that all the techniques presented in this paper allow implementing an approach "*from scratch*", which means that the developed tools can in principle start predicting with one single example of each class, a disruptive and a non-disruptive discharge. The first model can therefore be trained with the first useful discharges at the beginning of a campaign.

5 JET Databases with the ILW and definition of the statistical indicators

To test the various adaptive training strategies, described in the previous sections, a large database of JET discharges with the ILW has been considered. The set of discharges covers the interval of campaigns C28-C32. This database is particularly interesting, for the purpose of assessing the potential of adaptive learning strategies, because C28 is the first campaign of JET with the new ITER Like Wall. The discharges included in the DB are therefore well suited to simulate the learning process of predictors at the beginning of the life of a device. As far as the definition of disruption is concerned, the traditional JET threshold has been retained, i.e. discharges with a current decay faster than 5MA/s are considered disruptive. The disruption time is taken as the beginning of the current quench. Moreover, again for coherence with the past, only discharges, whose plasma current exceeds 750 kA, have been included. On the contrary, intentional and mitigated disruptions have not been considered. Overall the database comprises 2428 discharges, 430 disruptive and 1998 non disruptive shots. Plots showing the operational space covered by the database are reported in Figure 3. For all the discharges of the DB, all the signals have been resampled at 1kH frequency and the predictions have been performed every ms.

In terms of the classification criteria, alarms, which are launched 5 ms or less from the

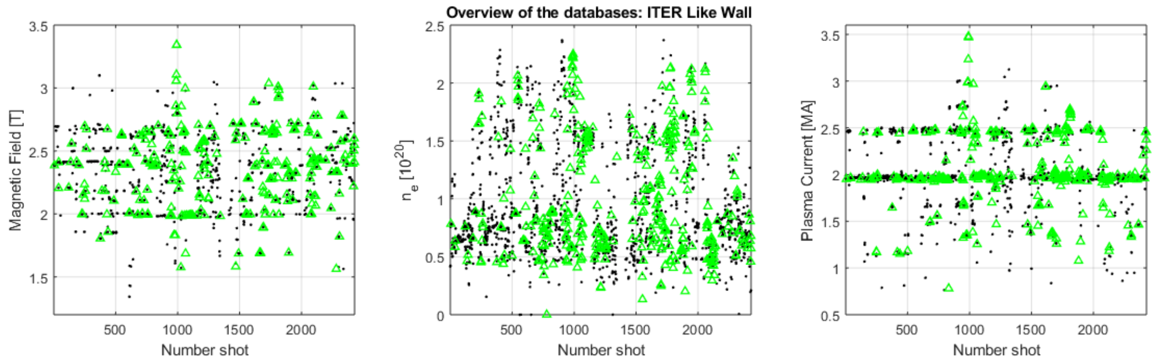

Figure 3.Overview of the databases for the ILW- A characteristic point for each shot in the database has been reported. The green triangles identify the disruptive shots.

beginning of the current quench, are considered tardy, since 5 ms is approximately the minimum time required on JET to undertake mitigation action. Alarms triggered more than 3 s before the beginning of the current quench are considered early; this is a conservative choice because many disruptions on JET are due to impurity accumulation, which can be detected even earlier than 3 s before the beginning of the current quench. The false positive plus the early alarms are summed to obtain the total rate of false alarms in the tables presented in the rest of the paper.

6 Results obtained by updating the training set in case of errors at specific time slices

JET database with the ILW, described in Section 5, has been analysed first by applying the basic adaptive training introduced in Subsection 3.1. The first model has been obtained with 5 safe discharges and the first disruption. For the safe shots the training times are four averages over 10 ms randomly chosen during the flat top. For the disruptions, three averages of 5 ms each, in the last 15 ms before the beginning of the current quench, have been provided as inputs to the classifiers. When updating the training set, for the missed and tardy alarms the same points are selected (three averages of 5 ms each, in the last 15 ms before the beginning of the current quench). To follow the evolution of the safe discharges, as discussed in Section

3, after each alarm, four points averaged over 10 ms are chosen on the previous safe discharge in the 40 ms before the alarm time of the current discharge.

With regard to the architecture of the ensembles, eleven noise realisations have proven sufficient to obtain good results. The bagging and Random Forests classifiers, applied to these eleven databases with noise, consist of 40 trees each. Therefore, the ensemble classifiers, described in the rest of the paper, are the combination of 440 individual trees. This

Table I Performance of the single CART tree and the ensembles by updating the training set in case of errors at specific time slices.

Method	Success Rate	Missed	Early	Tardy	False	Mean [ms]	Std [ms]
Single CART	93.24 (400/429)	0.47 (2/429)	3.50 (15/429)	2.80 (12/429)	14.39 (289/2008)	385	397
BAG	89.04 (382/429)	0.70 (3/429)	6.53 (28/429)	3.73 (16/429)	11.13 (225/2021)	369	407
RF	95.10 (408/429)	0.93 (4/429)	0.93 (4/429)	3.26 (14/429)	8.37 (167/1996)	371	401
Noise+RF	94.87 (407/429)	1.17 (5/429)	0.47 (2/429)	0.47 (2/429)	7.17 (143/1995)	364	397
Noise+BAG	95.10 (408/429)	0.70 (3/429)	0.70 (3/429)	3.73 (16/429)	7.97 (159/1996)	374	400

is the level of complexity of all the classifiers used to obtain the results reported in the following. In terms of aggregation, to derive a unique decision from the outputs of the various predictors in the ensembles, a simple threshold, in the percentage of the weak learners detecting a disruption, has been adopted at this stage (DF1 functions described in the next section).

As a preliminary step, the single CART tree and the ensemble classifiers have been used to explore the database and to identify the most relevant global variables for mitigation. With this goal, they have been trained with a series of global quantities as inputs, with the aim of determining the ones providing the best results. Only individual time slices have been provided to the predictors in the traditional way (no trajectory learning at this stage). Also only signals routinely available in real time have been considered in the perspective of future deployment in JET real time network. The quantities analysed as candidate inputs have been: the vacuum toroidal magnetic field B, the plasma current I_p, the amplitude of the locked mode LM, the amplitude of the internal inductance *li*, the safety factor at 95% of the plasma radius q_{95} and the edge density *n*. Confirming previous analyses and common knowledge widely reported in the literature, the locked mode and internal inductance signals have proved to be the most informative [13-20]. These two quantities not only allow reaching good performance on JET but adding additional global quantities typically does not improve the outputs. Therefore, this preliminary analysis confirms various studies reported in the past showing the importance of these two quantities in predicting disruptions not only on JET but also on other devices. Overall the performance obtained by applying the ensembles to individual time slices of the locked mode and internal inductance signals are reported in Table I. To interpret the numerical results of this and the following tables, it should be mentioned that the single disruption and the 5 safe shots used to train the first predictor are not included in the statistics.

7 Results obtained by updating the decision function

The performance of the noise based ensembles, reported in Table I, are reasonable and quite competitive with traditional predictors, particularly if one takes into account the adaptive nature of the approach adopted. On the other hand, there are still margins for improvements.
First, it has been tested whether other decision functions can fare better than the simple threshold in the percentage of the weak learners detecting a disruption. To this end 6 different decision functions have been tried, comprising all the most widely formulas typically used to aggregate the results of ensembles. The mathematical details are provided in Table II.
As expected, there is no individual decision function which always outperforms all the others. The best alternative therefore consists of switching to the most performing decision function dynamically during the evolution of the campaigns. In this perspective, the main difficulty is always the definition of suitable indicators to reliably implement the switching between the decisions functions. Theoretical methods to guide in this task do not exist and empirical solutions have to be found on a case-by-case basis. In our application, a quite systematic investigation of possible alternatives has shown that the following two indicators are very effective in determining when to switch between the individual classifiers:

$$IND1 = VAR1 - \frac{VAR2}{\%Alarm\ Threshold} \qquad (1)$$

$$IND2 = \frac{VAR1}{VAR2} - \frac{30 \cdot VAR2}{\%Alarm\ Threshold} \qquad (2)$$

VAR1 is a counter, which is increased by one unit every time a predictor in the ensemble does not trigger any alarm for an entire safe shot. In its turn, VAR2 is increased by one unit every time a predictor in the ensemble misses a disruption or triggers a tardy alarm. So basically the two indicators are increased by right predictions and decreased by wrong ones. It is worth mentioning that the two versions of the numerical indicators are quite robust, in the sense that reasonably small variations in the constants do not affect significantly the final performance, at least for the database analysed in this paper. On the other hand, the

Table II The basic features of the families of decision functions applied to the JET database with the ILW. The percentage of disruptivity mentioned in the table is to be understood as the percentage of trees detecting a disruptive behaviour.

Decision Function	Basic mathematical description	Colour
DF1	Percentage threshold: the classification decision is given when exceeding a threshold in the percentage of the 440 total models detecting a disruption. The votes that the time slice is disruptive are transformed into a percentage and if it exceeds the 80% threshold then an alarm is triggered.	Blue
DF2	Tree majority voting: the classification decision is given by the majority voting of the 40 trees, one for each noise realisation. For each tree the percentage of disruptivity is calculated and if the percentage is below 30% or above 70% the tree is retained in the final majority voting.	Tan
DF3	Tree percentage threshold: the classification decision is given by the percentage over the 40 trees, one for each noise realisation. If this percentage exceeds the 80% threshold then an alarm is triggered. For each tree the percentage of disruptivity is calculated and if the percentage is below 30% or above 70% the tree is retained in the final majority voting.	Gold
DF4	Ensemble majority voting: the classification decision is given by the majority vote of the 11 ensembles. For each ensemble, all the votes of the related trees (40) are considered (majority voting).	Dark Olive Green
DF5	Ensemble majority voting with filtering: the classification decision is given again by the majority vote of the 11 ensembles as for DF4 but only the ensembles, whose percentage of disruptivity is below 30% or above 70%, are retained in the calculation of the majority voting.	Turquoise
DF6	Ensemble majority voting with filtering: the classification decision is given again by the majority vote of the 11 ensembles as for DF4 but only the ensembles, whose percentage of disruptivity is below 30% or above 70%, are retained in the calculation of an average; if the disruptivity is above the 80%threshold then an alarm is triggered .	Purple

deployment of the two indicators IND1 and IND2 on the existing database has not provided very significant improvements in performance, which is only marginally better than the values reported in Table I. The reason, for the very limited impact of switching between different decision functions, resides in the data provided as inputs to the predictors, which constitute a too limited description of the plasma. Selecting dynamically the best decision function provides much better results if applied after trajectory learning, as described in the next section.

8 Results obtained by taking into account the history of the individual discharges

Given the outputs of Section 6, the approach of trajectory learning has been applied only to the locked mode and internal inductance signals. A series of systematic tests has been performed to identify the most appropriate lengths of the signals to be used as inputs for the

Table III. Performance obtained by implementing the trajectory learning to the noise based ensembles with Bagging and Random Forests.

Decision Function	Success Rate	Missed	Early	Tardy	False	Mean [ms]	Std [ms]
1 RF	95.10 (408/429)	2.56 (11/429)	0.23 (1/429)	2.10 (9/429)	4.91 (98/1994)	335	355
1 BAG	94.41 (405/429)	3.73 (16/429)	0.23 (1/429)	1.63 (7/429)	4.01 (80/1994)	339	355
2 RF	97.44 (418/429)	1.86 (8/429)	0.23 (1/429)	0.47 (2/429)	5.57 (111/1994)	343	354
2 BAG	97.44 (418/429)	1.40 (6/429)	0.00 (0/429)	1.17 (5/429)	3.81 (76/1993)	338	357
3 RF	95.34 (409/429)	2.80 (12/429)	0.00 (0/429)	1.86 (8/429)	4.21 (84/1993)	338	364
3 BAG	96.97 (416/429)	1.86 (8/429)	0.23 (1/429)	0.93 (4/429)	4.46 (89/1994)	319	335
4 RF	96.97 (416/429)	1.17 (5/429)	0.23 (1/429)	1.63 (7/429)	6.37 (127/1994)	349	358
4 BAG	97.90 (420/429)	0.93 (4/429)	0.23 (1/429)	0.93 (4/429)	6.32 (126/1994)	342	360
5 RF	97.67 (419/429)	0.70 (3/429)	0.47 (2/429)	1.17 (5/429)	7.07 (141/1995)	377	402
5 BAG	97.67 (419/429)	1.63 (7/429)	0.23 (1/429)	0.47 (2/429)	5.52 (110/1994)	348	353
6 RF	96.97 (416/429)	1.40 (6/429)	0.23 (1/429)	1.40 (6/429)	4.76 (95/1994)	335	351
6 BAG	96.27 (413/429)	1.86 (8/429)	0.23 (1/429)	1.63 (7/429)	4.61 (92/1994)	336	352

ensembles to best learn the history effects during the discharges. The best results have been obtained using, in addition to the locked mode and internal inductance amplitudes at the beginning of the current quench, two additional samples for each of the two signals: the locked mode amplitude at 40 and 50 ms before the beginning of the current quench and the internal inductance amplitude at 20 and 40 ms before the beginning of the current quench.

In the course of this analysis, the issue of the assertion time has been investigated in detail. Given the fluctuations of the various plasma parameters, triggering an alarm after the first time slice, when a disruptive behaviour is detected, could not necessarily be the best choice. In the interest of reliability, it might be better to wait for the ensembles to identify an abnormal situation for two or three subsequent time intervals. In the present application, particularly after the implementation of the trajectory learning, the best performances have been obtained by triggering an alarm immediately after the predictors detect a disruptive behaviour. With this choice of the signals as inputs, and implementing also the switching of

RANDOM FOREST		**RANDOM FOREST**	
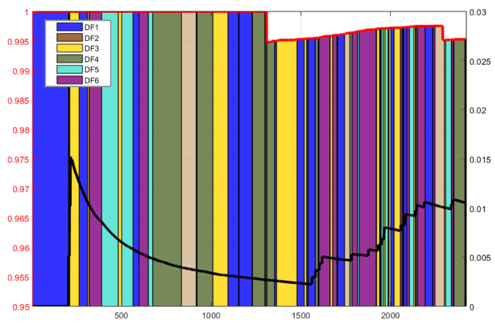		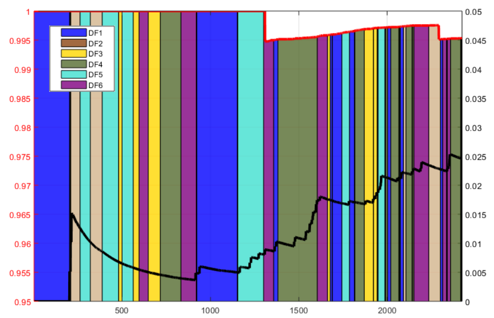	
Success Rate	99.53	**Success Rate**	99.53
False	1.05	**False**	2.45
BAGGING		**BAGGING**	
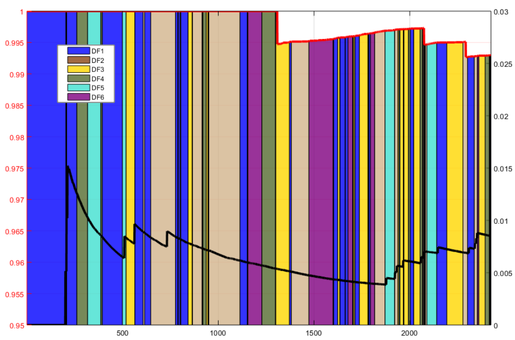		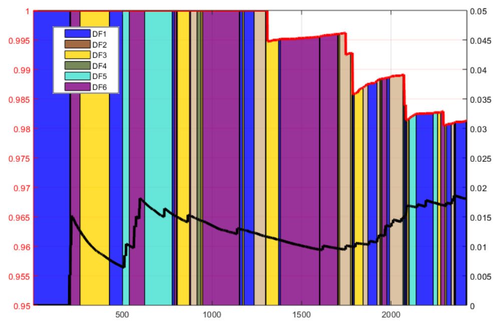	
Success Rate	99.3	**Success Rate**	98.37
False	0.85	**False**	1.80

Figure 4. Performance obtained with implementing de-learning by retaining only the last 10 discharges in the training set and zeroing the indicators IND1 and IND2 after each switch of the decision function. Top: results obtained with IND1. Bottom: results obtained with IND2. The red line and the left vertical axis report the success rate; the black line and the right vertical axis report the number of false alarms.

the decision functions as described in the previous section, the results are significantly better, particularly in terms of false alarms, as can be seen from the statistics of Table III. From the results reported, it can be concluded that trajectory learning improves noticeably the performance. On the other hand, there is clear trade-off between success rate and false alarms; increasing performance in one of these two fundamental indicators produces a significant degradation in the other. This fact is not remedied by including information about the signals trajectories. The observation that the performances tend to decrease toward the end of the C32 campaign, after being significantly better earlier on, suggests that the oldest examples could be obsolete and could be profitably discarded, as investigated in the next section.

9 De-learning to manage obsolescence

Discarding old examples is an important aspect of any form of learning in non-stationary conditions. JET programme is now evolving very quickly and therefore data of past discharges can become quickly misleading and reduce the performance of the predictors. Therefore adaptive strategies, which do not include effective forms of de-learning, can give very rapidly suboptimal performance. On the other hand, it is quite difficult to find effective criteria to discard old examples. General theories are not available and this is another typical aspect of developing predictors which has to be addressed empirically.

Among the various alternatives tested, the ones providing the best results, with reasonable computational efforts, are the two described in the following. First, at each retraining only the last 10 discharges in the training set are retained (with obvious clerical checks to make sure that at least a disruptive and a safe shot are included in the set). The second form of de-learning implemented is the simple zeroing of the indicators IND1 and IND2 after each switching of the decision function. This is of course a form of de-learning; old examples are discarded and do not enter into the calculation of the various decision functions performance.

The obtained results are summarised in the plots of Figure 4. The improvement is very clear and the overall performance reach unprecedented levels in terms of both success rates and false alarms for the case of Bagging combined with IND2. The errors, in terms of both success rate and false alarms, are in the range of a few cases per thousand. Particularly the case of the bagging using IND2, which is always the most performing alternative, achieves success rate in excess of 99 % and false alarms less than 1 % (see bottom left plot of Figure 4). For this predictor, only two disruptions are detected late and one early (out of 429) and the false alarms are only 15 discharges (out of 1993). Such levels of accuracy are particularly relevant if one considers that the method is adaptive and therefore the predictors have to learn as the campaigns evolve [13-20]. Moreover, in the statistics of the errors are included various cases, for which it is either very doubtful that the predictors really made a mistake or that could have been avoided by considering in advance the objectives of the experimental sessions. In particular of the 15 false alarms, at least 10 occur in the ramp down phase of the plasma current and are triggered by minor disruptions. A couple of representative cases are reported in Figure 5. Of course with more experimental time, strategies could be devised to avoid such situations by designing better ways to terminate the discharges. Once optimised the exit rom the H mode and the ramp down of the current, if necessary specific predictors could be developed for the detection of actual disruptions in the final phases of the discharges, which can reasonably be expected to have even better performance.

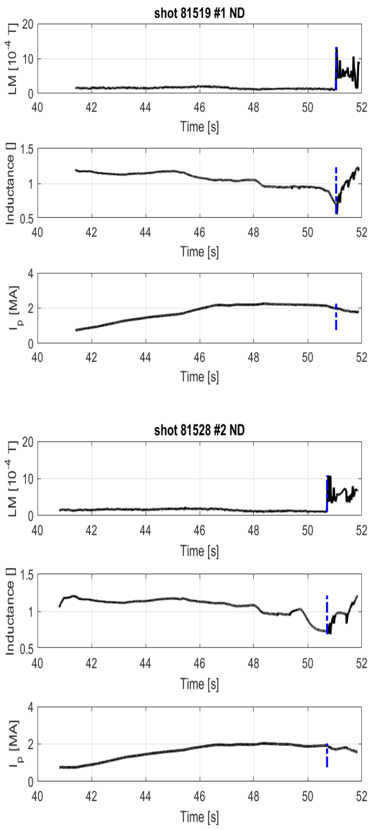

Figure 5. Examples of false alarms during the current decay phase of the discharge. The vertical dashed lines identify the beginning of the current quench.

10 Conclusions and further developments

The adaptive techniques developed in this work are the most sophisticated ever applied to Tokamak signals. They have been implemented with a quite powerful set of ensemble classifiers, which have proved more than adequate to the task of predicting disruptions for mitigation using a large database of JET discharges with the ILW. All the steps of the developed adaptive training are compatible with real life implementation, either during or between discharges. The proposed approach is also from scratch and therefore the various tools can be deployed after the first disruption. On the other hand, it should be stressed that implementing adaptive strategies from scratch does not mean that no prior information is required. The proposed methodology needs indeed to set some parameters, in particular the length of the training data sets for de-learning and the amount of time points to be retained for trajectory learning. These quantities have been derived by a first analysis of the last campaigns of JET with the carbon wall.

The obtained performances are unprecedented both in terms of success rate and false alarms, being in the range of a few per thousand. Even more importantly, the performance are always well in excess of ITER requirements. During the entire range of campaigns investigated, indeed, the success rates of the best predictors developed always hover above 99% and the false alarms rates remain well below 2%. Therefore, in the light of these results, the approach of adaptive learning is expected to be very important at the beginning of operation of the next generation of devices, especially at the beginning of their operation. ITER in particular will have to increase the current to 15 MA in a relatively short time. Moreover, to prepare DT scenarios, various isotopic compositions will have to be investigated. Previous experience in JET shows that different fuel mixtures can alter significantly various aspects of the discharges, which can affect their disruptivity. Good examples are the ELMs, the emitted radiation and detachment. It is not at all evident that all these effects can be taken into account from the beginning without any need for retraining of the predictors. Therefore some form of adaptive predictors, not necessarily exactly deploying the same technologies investigated in this paper, will probably be important for ITER. Another important issue to consider is avoidance. Also in this perspective, adaptive strategies could prove to be essential to cope with the great variety of potentially disruptive situations, particularly at high fusion power when also the alpha particles could strongly influence the stability of the configurations. For avoidance, of course, additional signals, and in particular profiles, will have to be included in the list of features to be provided as inputs to the predictors; indeed quantities such as the radiation profiles are considered essential to obtain early warnings of problems in integrated scenarios [21,22].

Aknowledgements

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work was also partially supported by the Spanish Ministry of Economy and Competitiveness under the Projects Nos. ENE2015-64914-C3-1-R and ENE2015-64914-C3-2-R.

References

[1] A.H. Boozer. "Theory of tokamak disruptions". Physics of Plasmas.19, 058101. 2012.

[2] C.G. Windsor et al. *A cross-tokamak neural network disruption predictor for the JET and ASDEX Upgrade tokamaks*. Nucl. Fusion 45 337. 2005.

[3] R.Wenninger et al "*Power Handling and Plasma Protection Aspects that affect the Design of the DEMO Divertor and First Wall*" submitted for publication in Proceedings of 26th IAEA Fusion Energy Conference

[4] J. Vega, A. Murari, S. Dormido-Canto, R. Moreno, A. Pereira, A. Acero and JET-EFDA Contributors. "Adaptive high learning rate probabilistic disruption predictors from scratch for the next generation of tokamaks". Nuclear Fusion. 54 (2014) 123001 (17pp).

[5] S. Dormido-Canto, J. Vega, J. M. Ramírez, A. Murari, R. Moreno, J. M. López, A. Pereira and JET-EFDA Contributors. "*Development of an efficient real-time disruption predictor from scratch on JET and implications for ITER*". Nuclear Fusion. 53 (2013) 113001 (8pp).

[6] G. Pautasso et al. *An adaptive real-time disruption predictor for ASDEX Upgrade.* Nucl. Fusion 50 075004. 2010.

[7] Murari et al "*Adaptive predictors based on probabilistic SVM for real time disruption mitigation on JET*" Nuclear Fusion, Volume 58, Number 5, March 2018, https://doi.org/10.1088/1741-4326/aaaf9c

[8] Rokach, Lior; Maimon, O. (2008). "*Data mining with decision trees: theory and applications.*" World Scientific Pub Co Inc. *ISBN 978-9812771711*.

[9] L. Breiman, J. H. Friedman, R. A. Olshen, C. J. Stone "*Classification and Regression Trees*" Chapman&Hall, London 1998.

[10] Sollich, P. and Krogh, A., *Learning with ensembles: How overfitting can be useful*, Advances in Neural Information Processing Systems, volume 8, pp. 190-196, 1996

[11] Sayed-Mouchaweh, M. "*Learning from Data Streams in Dynamic Environments*" Springer Briefs in Applied Sciences and Technology 2016

[12] J.Vega et al "*A linear equation based on signal increments to predict disruptive behaviours and the time to disruption in JET*" to be submitted to Nuclear Fusion

[13] A.Murari et al, Volume 49, Number 5 April 2009 doi.org/10.1088/0029-5515/49/5/055028

[14] A.Murari et al Nuclear Fusion, Volume 48, Number 3 February 2008 doi.org/10.1088/0029-5515/48/3/035010

[15] G.Rattà et al Nuclear Fusion, Volume 50, Number 2 January 2010 doi.org/10.1088/0029-5515/50/2/025005

[16] Y.Zhang et al Nuclear Fusion, Volume 51, Number 6 May 2011 doi.org/10.1088/0029-5515/51/6/063039

[17] J. Vega, R. Moreno, A. Pereira, S. Dormido-Canto, A. Murari and JET Contributors. "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.

[18] .Vega, A. Murari, S. Dormido-Canto, R. Moreno, A. Pereira, G. A. Rattá and JET Contributors. "Disruption Precursor Detection: Combining the Time and Frequency Domains". Proc. of the 26th Symposium on Fusion Engineering (SOFE 2015). May 31st-June 4th, 2015. Austin (TX), USA

[19] G. Pautasso et al. Plasma shut-down with fast impurity puff on ASDEX Upgrade. Nuclear Fusion, Volume 47, Number 8. 2007.

[20] S.P. Gerhardt et al. Nuclear Fusion, Volume 53, Number 6 17 May 2013 • 2013 IAEA, Vienna

[21] Ongena, J. et al "*Towards the realization on JET of an integrated H-mode scenario for ITER*" (2004) Nuclear Fusion, 44 (1), pp. 124-133. DOI: 10.1088/0029-5515/44/1/015

[22] Puiatti, M.E. et al, "*Radiation pattern and impurity transport in argon seeded ELMy H-mode discharges in JET*" (2002) Plasma Physics and Controlled Fusion, 44 (9), pp. 1863-1878. DOI: 10.1088/0741-3335/44/9/305