EFDA–JET–PR(08)35

A. Murari, J. Vega, G.A. Rattà, G. Vagliasindi, M.F. Johnsons, S.H. Hong
and JET EFDA contributors

Unbiased and Non-Supervised Learning Methods for Disruption Prediction at JET.

Unbiased and Non-Supervised Learning Methods for Disruption Prediction at JET.

A. Murari, J. Vega, G.A. Rattà, G. Vagliasindi, M.F. Johnsons, S.H. Hong
and JET EFDA contributors*

JET-EFDA, Culham Science Centre, OX14 3DB, Abingdon, UK

[1]Consorzio RFX-Associazione EURATOM ENEA per la Fusione, I-35127 Padova, Italy.
[2]Asociación EURATOM-CIEMAT para Fusión, CIEMAT, Madrid, Spain
[3]Dipartimento di Ingegneria Elettrica Elettronica e dei Sistemi-Università degli Studi di Catania, 95125 Catania, Italy
[4]EURATOM-UKAEA Fusion Association, Culham Science Centre, OX14 3DB, Abingdon, OXON, UK
[5]CEA, IRFM, F-13108 Saint-Paul-lez-Durance, France.
** See annex of F. Romanelli et al, "Overview of JET Results ",*
(Proc. 22nd IAEA Fusion Energy Conference, Geneva, Switzerland (2008)).

Preprint of Paper to be submitted for publication in
Nuclear Fusion

ABSTRACT.
The importance of predicting the occurrence of disruptions is going to increase significantly in the next generation of Tokamak devices. The expected energy content of ITER plasmas, for example, is such that disruptions could have a significant detrimental impact on various parts of the device, ranging from erosion of plasma facing components to the structural damage. Early detection of disruptions is therefore needed with evermore increasing urgency. In this paper, the results of a series of methods to predict disruptions at JET are reported. The main objective of the investigation consists of trying to determine how early before a disruption it is possible to perform acceptable predictions on the basis of the raw data, keeping to a minimum the number of "ad hoc" hypothesis. Therefore the chosen learning techniques have the common characteristic of requiring a minimum number of assumptions. Classification and Regression Trees (CART) is a supervised but on the other hand a completely unbiased and non-linear method, since it simply constructs the best classification tree by working directly on the input data. A series of unsupervised techniques, mainly K-means and hierarchical, have also been tested, to investigate to what extent they can autonomously distinguish between disruptive and no-disruptive groups of discharges. All these independent methods indicate that, in general, prediction with a success rate above 80% can be achieved not earlier than 180 ms before the disruption. The agreement between various completely independent methods increased the confidence in the results, which are also confirmed by a visual inspection of the data performed with pseudo Grand Tour algorithms.

1. INTRODUCTION

The Tokamak configuration remains the most serious candidate for a magnetic confinement fusion reactor, thanks to its higher performance in terms of plasma parameters, compared to the most credited alternatives. On the other hand, Tokamak plasmas are particularly vulnerable to sudden losses of confinement, called disruptions [1], which produce a violent and unforeseen termination of the discharge. Disruptions are potentially very harmful events. First of all, in a very short time the energy content of the plasma is deposited on the first wall, causing very high thermal loads. Secondly, the fast termination of the plasma current induces eddy currents on the surrounding metallic structures, which can give rise to high induced forces. The risk involved in disruptions is already quite significant in present day large devices like JET and it is going to increase significantly in the next generation of machines like ITER, which will work at much higher plasma currents and thermal energy. Therefore the need to develop reliable algorithms capable of predicting sufficiently in advance the occurrence of a disruption, in order to have time to undertake remedial action, are becoming increasingly more stringent. On the other hand, the multiplicity of causes leading to a disruption and the nonlinear nature of the phenomenon have so far prevented the formulation of a consistent theory, capable of providing reliable prediction. As a consequence, in the last year various "soft computing" methods have been explored, ranging from Artificial Neural Networks (ANNs) to Fuzzy Logic (FL). The first attempts [2,3,4] with ANNs in the middle of the nineties showed a good rate of success only very near, typically a few milliseconds, to the occurrence of the disruption. Later predictors, based also on FL, seem to provide a good predictive capability significantly earlier even some hundreds of ms before the disruption [5,6,7,8,9,10]. Significant work ahs also been done at JET sing supervised methods [11,12] and trying to develop a multi-machine classifiers [13].

On the other hand the studies performed so far have not yet addressed two major issues. First of all, no systematic investigation of the information contained in the signals has been performed to determine, in general, how early in the discharge the incoming disruptions manifest themselves with enough clarity for a predictor to work acceptably. Second, in order to improve the performance of the predictors in terms of success rate, increasingly more specific training has been performed and therefore it is not clear how the reported success rates will scale to future devices.

In this paper, the application of various "Learning Methods" (LM) methods to the problem of disruption prediction is described. The main aim of the study consists of trying to obtain objective indications about how early in JET discharges, disruptions manifest themselves clearly in the signals. To increase the applicability of the approach, the most general and least biased methods have been investigated. First of all Classification and Regression Trees (CART) have been investigated to identify the signals carrying more information about incoming disruptions (see section 2). Since CART is a very general, unbiased and fully nonlinear approach, it has also been used to classify new discharges after having been trained on a validated database. Notwithstanding its generality, CART remains a supervised method and therefore requires the users to prepare a training set of discharges, already divided in disruptive and non disruptive, for the algorithm to learn. In order to double-check that no bias had been introduced in the preparation of the training set, a series of unsupervised methods have been also applied to the JET database of disruptions. In particular K-means and hierarchical techniques have been considered to determine to what extent and how early JET discharges could be classified into disruptive and safe by unsupervised clustering algorithms (see section 3). These two completely independent techniques, CART and unsupervised clustering, applied to the same or different JET databases provide very accurate predictions up to about 180 ms before the disruption. Around this time, their performance significantly degrade, indicating that the information content in the signals is reduced and that disruptions do not leave a clear footprint of their future occurrence any more. These results have been finally confirmed by simply visualising the multidimensional signals in a lower-dimensional representation space (2D) using a Grand Tour algorithm (see section 4). Possible directions to extend the present work are briefly outlined in section 5.

2. AN UN-BIASED NON-LINEAR APPROCH: CLASSIFICATION AND REGRESSION TREES

2.1 THE METHOD

CART is an algorithm explicitly conceived to construct trees for classification (if the output variable is discrete) or regression (if the output variable is continuous) [14]. To explain the basic properties of the approach, it is better to focus the discussion on the case of our interest, the classification of discharges into two classes, disruptive or non disruptive. The CART algorithm traverses the entire database and, for each input variable, tries to find the value that best divides the database into the two desired groups. In a certain sense, the algorithm seeks to maximize the purity of the two subclasses, called child nodes, by splitting the database into two subgroups each containing discharges belonging only to one class. In general no variable has such an explanatory power to allow a perfect division of the entire database into two completely separated child nodes. Therefore at each stage the algorithm chooses the variable that provides the best division of the database in two homogenous groups and then the process is repeated for the child nodes until a complete separation is achieved (or all the variables have been processed). By construction, the resulting tree has the more important variables toward the root and the ones with less explanatory power toward the leaves.

The CART approach is a supervised technique and therefore a training set containing properly classified examples is required (in our case a set of discharges with a well diagnosed time of disruption). Once properly trained the resulting tree can be used to classify further examples. It is sufficient to take a new set of input signals and traverse the tree to see in which leaf node (disruptive or safe) the new example is classified. To optimise this aspect of prediction, the final operation of pruning the tree is particularly important. This phase consists of eliminating the final nodes which increase the complexity of the tree, without bringing sufficient improvement in the classification. This pruning operation not only provides way to find a trade-off between complexity and accuracy but also influences the generalisation properties of the final tree.

It must be emphasised that the method is very general. It is totally unbiased, since no manipulation of the raw data, not even normalisation, is needed. The algorithm is also fully nonlinear since it exhaustively traverses the entire database, maximizing the purity of the child nodes simply on the basis of the information content of the signals. It is therefore particularly suitable to the prediction of disruptions, which are very variegated and nonlinear phenomena.

2.2 DATABASE AND FEATURE EXTRACTION

The database used for the study, whose results are presented in this section, is extracted from the JET disruption database and is a subset of the one introduced in [15, 16]. It is composed of signals from 440 pulses, 220 of which end with a disruption and are, therefore, referred to as “disruptive”, whereas the remaining 220 are pulses which terminated normally, and are therefore referred to as “safe”. The dataset for each disruptive pulse consists of several signals (see later) made of 21 points each (one sample every 20 ms), in the time interval [t_D-440, t_D-40], where t_D is the time the disruption takes place. For safe pulses, i.e. pulses where a disruption does not take place, the data is composed again of 21 samples, sampled at 20 ms, taken 7 seconds after the X-point formation. This choice has been taken with consideration that, on average, this is the time of disruption for the set of disruptive discharges selected. Finally, the data from the database has been divided into two groups, a training and a testing set, both composed of 220 pulses, half disruptive and half safe. Both training and test sets include a selection of discharges belonging uniformly to all the campaigns in the database.

The signals which have been included in the database are listed in Table I. They are all available in real time from JET diagnostics except the Net Power (P_{net}) which is the arithmetic difference between the total input power (P_{inp}) and the total radiated power.

This choice of these signals as the most relevant has been based on the use of CART for feature extraction. This aspect of feature extraction is a major issue in the analysis of big databases and in our applications consists of deciding which, among the thousands of JET acquired signals, are the most relevant to study the problem disruption prediction. As a preliminary step, various experts have been asked to identify the most relevant quantity for this end. They have converged on a list of 13 signals, which is the one used in previous papers like [15]. To further reduce the number of variables, they have been used as input to the CART algorithm. By construction, the CART software locates the most relevant signals towards the root of the tree and the ones with less explanatory power towards the leaf nodes. This allows the algorithm to be used to classify the input signals in terms of their relative importance and therefore provides an unbiased and nonlinear method to select the measurements more relevant to predict the phenomenon. It is worth mentioning that the assessment of the variable relevance includes also the inputs that do not explicitly appear in the final optimal tree. This is appropriate since it can happen that variables are masked, which means that they can be a surrogate of other variables, never occurring in this way as a primary splitter although they are the best choice after the chosen variables in the selected tree. CART allows this problem to be overcome by the so-called "variable ranking method", which consists of summing across all nodes in the tree the improvement scores that the variable induces when it acts as a primary or surrogate splitter. In this way, variables that never appear in the tree being always a surrogate of another input are also considered in the final classification.

The relevance values so produced allow ranking the different input signals from high to low importance. In this way, CART can be used for feature selection, being able to identify the most important variables to describe the output. The results of this process of feature extraction using CART is indeed the selection of signals reported in table I.

During the feature selection process aimed at identifying the most relevant signals, it has emerged from the available database that this classification depends significantly on the interval considered before the disruption, as reported in table II (a possible way to overcome this general problem that the relative importance of the variables depends on the time interval chosen is discussed later in this section).

In order to assess the robustness of the described signal selection process, Gaussian noise was added to the input signals to simulate different level of signal to noise (S/N) ratio. The CART algorithm has been applied again and for reasonable values of the S/N ratio, above 20, certainly realistic for JET measurements. The results of this sensibility study are reported in table III. The classification of the variables is quite robust since the most significant variables remain the same. The small variations in the ranking of some signals is a characteristic intrinsic to the database, in the sense that these variables have a very similar explanatory power and therefore small relative variations induced by the noise can affect their relative rating.

2.3 THE CLASSIFICATION RESULTS

In addition to assessing the explanatory power of the input signals, the CART method can be used for the purpose it was originally conceived, to perform unbiased and non-linear classifications of discharges in the two groups of disruptive and safe. To accomplish this, first of all a complete training of the CART for different time intervals was performed using the training datasets. The time intervals chosen are the same reported in Table II and used to analyze the relative importance of the variables at different times. Then the classification of the test dataset was performed with all these CARTs, trained in their specific subintervals. These results are reported in fig. 1. From a close look at this figure, it is possible to observe that the three networks performing best in proximity of the disruption, i.e. the [t_D-200, t_D-160], [t_D-140, t_D-100] and [t_D-80, t_D-40], have a steep slope around t_D-180 ms, which indicates that the footprint of the disruption is unclear for times prior to t_D-180 ms. On the other hand, training data further than 200 ms from the discharge tends to result in poor performance over the entire dataset confirming that, probably, so far away from the disruption there is not enough information in the signals for an effective prediction.

The robustness of the results reported in fig.1 have been double-checked by adding Gaussian noise to the signals and determining the success rate of the various trees with these new input signals. For reasonable levels of the S/N ratio, above 20, the noise does not have a very significant impact on the performance of the classifiers and the trends of fig.1 are reproduced, as shown in fig.2. For S/N ratios below 30 the success rate of the trees is strongly reduced particularly close to the disruptions. The global performance of the classifiers become flat around 60% over the entire time interval, showing that the additional information present in the signals is masked by such a level of noise. This of course emphasizes one more time the need of good measurements for the prediction of complex phenomena like disruptions.

To summarise, the CART method indicates quite clearly that in general disruptions manifest themselves clearly only with a maximum notice of about 180 ms. On the other hand expressing the results of fig.1 in terms of false and missed alarms, as illustrated in fig.3, suggests a possible strategy to train the predictor depending on the objectives of a certain session or campaign. For example, on a device with effective tools to terminate discharges quickly (in the order of 100 ms), it could be a good strategy to train the CART tree with data not earlier than 200 ms before the disruption. For instance, training with data from the interval [t_D-200: t_D-160] would grant a very good rate of success close to the intervention time disruptions at a prize of a high level of missed alarms earlier. On the other hand, in a device with a need to detect the imminence of disruptions as soon as possible, it could be a more effective approach to train the tree with data taken earlier than 200 ms before the disruption (for example in the interval t_D-260: t_D-220) to have a lower percentage of missed alarms earlier in the discharge, at a price of a higher level of false alarms.

3. UNSUPERVISED APPROACHES: K-MEANS AND HIERARCHICAL METHODS

3.1 THE METHOD

As mentioned in the introduction, CART is a powerful unbiased and nonlinear method but it remains as a supervised approach. It has therefore been considered important to see if its results could be confirmed by a general, unsupervised clustering technique. The tested ones have been partition-based (K-means) and hierarchical algorithms [17] but since the former has provided better results the discussion will be particularised for it in this paper.

The K-means approach has been developed to identify groups or clusters of datapoints in a multidimensional space. The main objective of the technique consists of partitioning the data into separate clusters of similar points, according to an appropriate definition of distance. Intuitively, therefore, a cluster can be interpreted as comprising a series of datapoints near to each other compared to the distance of points outside the cluster. The K-means algorithm has been devised to partition the original dataset in K clusters of similar points. To achieve this, the algorithm proceeds iteratively calculating firstly for each new point the most appropriate cluster, on the basis of an mathematical distance (Euclidean, City Block, Mahalanobis or Correlation), and then recalculating the cluster barycentre given the new points. The iteration is halted when the points are grouped in such a way as to minimize an overall global function, which represents the distance between the points in the database and the barycentre of the clusters.

One of the most significant issues in the implementation of unsupervised clustering methods is the adequate determination of the number of clusters in which the dataset should be partitioned. In our application to disruption identification, it would look natural to use only two clusters, one for the disruptive and one for the non-disruptive discharges. In reality, the problem is more complex due not only to the nature of the disruptions but to the diversity of the non-disruptive shots.

As will be discussed in more detail in section 5, disruptions can be classified into six main categories, on the basis of the cause leading to the loss of control of the discharge, and for that reason the plasma behaviour in each case is necessarily dissimilar. Besides, in each session during a typical JET campaign, there can be tens of “safe” experiments, every one of them of diverse nature but with common evolution without major unexpected instabilities. This “a priori” information is extremely useful to establish the cluster number as 7, under the assumption that six clusters would naturally classify the different types of disruption and the last one would capture the “safe” shots. That assumption has been verified after a careful study of the distributions of the discharges with the K-Means. In figure 4 it can be seen that in two intervals ([t_D-80ms, t_D-40ms] and [t_D-200ms, t_D-160ms], t_D being the time when the disruption occurs), the safe shots are mainly grouped in one cluster while the others in the remaining ones.

3.2 THE DATABASE

In order to simplify the comparison with the results of the CART method, the same database described in section 2.2 has been selected for the unsupervised analysis. For each discharge the signals adopted to perform the classification are the nine already identified by CART as the most relevant. Not all of them have the same time base and therefore they have been re-sampled when necessary to obtain a time resolution of 1 ms. This choice has proven to give enough samples to perform the analysis over time intervals of 40 ms, which is a good trade-off between time resolution and statistical basis for these clustering methods. In order to apply the K-mean algorithms in a convenient way, the 9 signals have been appended, forming a feature vector. The final database consists therefore of a sequence of these feature vectors, each one containing the samples pertaining to the time slice chosen for the particular analysis.

3.3 THE FEATURE EXTRACTION

In each discharge, the signals under study represent the evolution in time of different plasma parameters. Some of the information about disruption precursors is quite hidden in them and therefore,

as will be shown, careful signal processing is required to extract the most relevant features for the classification task at hand. In the course of this study, several data processing algorithms have been applied to each signal and the success rates in the final classification have been compared. The best results have been obtained with the sequence of operations described in the following.
First of all, it is necessary to normalize the data due to the high amplitude differences in the involved signals. A standard normalization formula has been implemented:

$$\text{Normalized signal} = \frac{\text{Raw\_signal\_Min}}{\text{Max - Min}}$$

Where Min and Max are the minimal and maximal values computed in all the dataset for each signal.

Secondly, the FFT (Fast Fourier Transform) has been calculated, discarding the first coefficient, in each time slice of the signals to extract valuable frequency information. Thirdly, the standard deviation of the normalized Fourier coefficients has been computed to measure the spread of the values. This sequence of mathematical operations has proved to be extremely useful in the construction of the training set. The CART analysis of the previous section, on the other hand, indicates clearly that different quantities can have a completely dissimilar explanatory value when it comes to the problem of disruption prediction. To profit from this basic information, each standard deviation has been multiplied by the corresponding CART coefficient that represents the relative importance of the signal for disruption prediction. In this way, the various quantities are weighted by a coefficient representing their relative explanatory power for the problem at hand.

At this point, each shot, initially characterized by a “global” vector of values corresponding to time slices of 40 ms of the 9 raw signals, each one re-sampled at 1 sample/ms (leading to a total amount per pulse of 360 samples) has been manipulated into a new feature vector of 9 values (the concatenation of the 9 values obtained by the previously described steps). This extreme dimensionality reduction not only achieves better classification results but also provides a simpler input to the K-Means algorithm, reducing the computational time.

3.4 TRAINING AND RESULTS

One of the most important issues for these unsupervised clustering methods is making sure that the database is properly chosen. Not only the training and test groups of discharges must be selected randomly to avoid any biasing but also each obtained result should be repeated and averaged to avoid misleading results due to unforeseen spurious correlations or statistical fluctuations. To obtain the most robust results given the available database, n-fold cross validation method [18] has been applied. It consists in splitting randomly the available database in n groups (each one with the same amount of haphazardly selected disruptive and “safe” pulses). One of the groups is kept for test and the other n-1 for training.

For every 40 ms interval, the training dataset is given as input to the K-Means algorithm (with seven clusters). The algorithm groups the data according to the Squared Euclidean distance (because of its better performance) and constructs the classification system, which consists basically of the barycentres of the clusters. Once the positions of the cluster barycentres have been determined, the test dataset is used to evaluate the success and error rates of the obtained classifier. The distance between every new object in the test group and the barycentre of each cluster of the classification system is measured. The object will be classified into the cluster with the minimal distance. The rate of success is then expressed in the usual percentage terms. The overall error is divided in 2 groups: when a “safe” shot is classified as disruptive, it is counted as a False Alarm, and in the opposite case, the error is counted as Missed Alarm.

This procedure is repeated n times, leaving in each occasion a different group for test purposes. Hence, n error rates are calculated independently and to obtain the final percentages the average of the n values has been performed.

The results are shown in figure 5. It can be noticed that the CART weighting improves considerably the performance of the classifier. The Missed and False Alarm rates correspond to the trace that includes the CART weighting in the feature extraction process.

The global results are in good agreement with the CART classifier, since they confirm that, earlier than 180 ms before the disruption, the success rate significantly decreases, proving that the information content in the signals is much lower further away from the sudden termination of the discharge.

4. CONFIRMATION OF THE RESULTS WITH A PSEUDO GRAN TOUR VISUALIZATION

4.1 DATA TOURS

Human beings are used to detecting patterns in their every day life. We easily recognize faces, places or follow object’s trajectories. Also, without any significant effort, we can detect differences or relationships between groups (people, cars, animals, buildings, landscapes etc). Since images have very high information content and they are the main means for human beings to explore the external world, it would be beneficial to confirm the results of the previous sections by visualizing the database. More ambitiously, this visualization could help identify possible structures or tendencies in the signals that could lead to statistical hypotheses. However, when the dimensions of what is being seen are significantly higher than the normal three of physical space, our comprehension of visual information decreases drastically. In the specific case of disruption precursors, the dimensionality of the raw data base is huge. To overcome this problem, first of all the analysis can be limited to the most relevant features already identified with the CART method in section two. Secondly the database can be projected into a subspace of lower dimensions to make the information visually more intuitive. These projections are indeed performed by a series of methods which fall in the general category of “data tours”, which provide the most interesting representations of the available data in a reduced feature space.

One of those methods, the so called pseudo Grand Tour, is based on the idea of showing the data from all possible viewpoints in a sequence of lower-dimension projections, so its evolution can be converted into a running movie of scatterplots, providing an overview of the high dimensional space in a sequence of 2D plots.

The high dimensionality of the problem at hand becomes manageable, since techniques to project high dimensional datasets on a 2-D surface, in such a way that the sequence of planes is dense, i.e.

it comes close to any given 2-D projection, exist. For our investigation we applied a Data Tour method, which projects the data on a 2-D subspace and shows the results as scatterplots [19]. The output of the algorithm is an animated sequence of scatterplots, which is representative of all projections of the original dataset. Visualizing a moving sequence has some additional benefits because it contains additional information, related to the movement of the data points. The pseudo Grand Tour implementation, adopted to obtain the results reported in this paper, is similar to the Torus Winding Method originally proposed by Asimov [20], which exploits the topological properties of a torus to identify the projecting planes. The pseudo Grand Tour was chosen because it presents some advantages like speed, uniformity of the tour, ease of recovering the projections and of the calculations, and also because it was already implemented in MATLAB [18], making possible it application by a friendly software

The visualizations of the same features used for the K-means clustering, described in section 3.3, are presented in figure 6. In Fig 6.a, the time interval [t_D-0.08, t_D-0.04] seconds before the disruption is shown. It is evident that most of the safe discharges are closely grouped (also detailed in a closer view). Fig.6.b reports other projections with the same characteristics for the interval [t_D-0.26, t_D-0.22 s] before the disruption. Two main groups of shots can still be recognized. Finally, in Fig.6.c, a distant time interval ([t_D -0.44, t_D -0.4] seconds before the disruption) is considered showing that that the discharges are mixed and no evidence of clustering is apparent. A systematic visual exploration of the database with the Grand Tour method therefore confirms that earlier than 200 ms there is no clear footprint of the disruption in the signals.

FUTURE DEVELOPMENTS

An important result of the clustering analysis seems to be the number of relevant groups of discharges. All tested clustering methods seem to indicate that the database lends itself to the division in seven clusters. This is a quite significant result because it is coherent with the typical classification of disruptions in terms of the trigger mechanisms. As reported in various papers and summarised in [15], the main factors triggering disruptions are believed to be the following:

a) Mode Lock (ML)
b) Density Limit (DL)
c) High Radiated Power (RP)
d) H/L mode transition (HL)
e) Internal Transport Barrier (IT)
f) Vertical Displacement Event (VDE)

A further investigation of the results of the automatic clustering methods, to see if they naturally tend to partition JET database into groups corresponding to these types of disruptive discharges would be a very interesting subject for further investigation.

Another interesting development would be to compare the data of machines of different size to see if the physiology manifests different time scales depending on the dimensions of the device. This of course would be extremely important in assessing the extrapolations to ITER

Switching to the subject of earlier disruption detection, it must be emphasized that the approaches presented in this paper are very general. They therefore provide a good basis for the extrapolation to bigger devices but on the other hand their results do not exclude the possibility of obtaining better performance by fine tuning supervised methods to the characteristics of specific machines. Some further refinements are indeed already under way for JET. Moreover, the signals analysed are the ones historically present in the JET database. In recent years new sophisticated diagnostics have become operational, some of which could potentially contain information about the imminence of a disruption earlier than the signals analysed so far. A good example is certainly the wide angle endoscope recently installed of JET for imaging of the main chamber [21]. In the case of density limit discharges, both the infrared and the visible cameras show a clear formation and evolution of a MARFE [22] instability as shown in figure 7. This instability is always present before this type of disruption and in some cases evidence of it, in the form of IR emission possibly due to hydrocarbons, has been detected even more than one second before the disruption. A systematic analysis of these results has not been completed yet but the potential of IR views for the prediction of at least the density limit disruptions should be investigated very seriously. Automatic image processing algorithms, to detect this footprint of MARFE instabilities, can be envisaged since the signature on the IR frames is very clear.

ACKNOWLEDGEMENTS
This work, supported by the European Communities under the contract of Association between EURATOM/ENEA Consorzio RFX, was carried out within the framework of the European Fusion Development Agreement. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

REFERENCES

[1]. Schuller F.C. 1995 "Disruption in tokamaks" *Plasma Phys.Control. Fusion* **37** A135–62

[2]. Hernandez J. V. *et al* 1996 "Neural network prediction of some classes of tokamak disruptions" *Nucl. Fusion* **36** 1009–17

[3]. Vannucci A. *et al* 1999 "Forecast of TEXT plasma disruptions using soft x-rays as input signal in a neural network" *Nucl. Fusion* **39** 255–62

[4]. Sengupta A. and Ranjan P. 2000 *Nucl. Fusion* **40** 1993–2008

[5]. Wroblewski D. *et al* 1997†"Tokamak disruption alarm based on a neural network model of the high-β limit" *Nucl. Fusion* **37** 725–40

[6]. Pautasso G. *et al* 2002 "On-line prediction and mitigation of disruptions in ASDEX Upgrade" *Nucl. Fusion* **42** 100–8

[7]. Milani F. 1998 "Disruption prediction in JET" *PhD Thesis*University of Aston, Birmingham, UK

[8]. Cannas B. *et al* 2004 "Disruption forecasting at JET using neural networks" *Nucl. Fusion* **44** 68–76

[9]. Morabito F.C. and Versaci M. 1999 "A fuzzy neural approach to plasma disruption prediction in tokamak reactors" *Int. Joint Conf. on Neural Network (Washington, DC)*

[10]. Versaci M. and Morabito F.C. 2003 "Fuzzy time series approach for disruption prediction in tokamak reactors" *IEEE Trans. Magn.* **39** 1503–6

[11]. Cannas B., A. Fanni, P. Sonato, M.K. Zedda and JET-EFDA contributors Nucl. Fusion **47** No 11 (November 2007) 1559-1569

[12]. Cannas B., Cau F., Fanni A., Sonato P., Zedda M.K. and JET-EFDA contributors Nucl. Fusion 46 No 7 (July 2006) 699-708

[13]. Windsor C.G., G. Pautasso, C. Tichmann, R.J. Buttery, T.C. Hender, JET EFDA Contributors and the ASDEX Upgrade Team Nucl. Fusion 45 No 5 (May 2005) 337-350

[14]. Breiman L., Friedman J.H., Olshen R.A. and Stone C.J. 1984 *Classification and Regression Trees* (Belmont, CA:Wadsworth Inc.1993, New York: Chapman and Hall)

[15]. Vagliasindi G., *et al* 2008 Fuzzy logic approach to disruption prediction at JET *IEEE Trans. Plasma Sci. vol. 36, no. 1, Feb. 2008, pp. 253-262.*

[16]. A. Murari G., Vagliasindi at al "Prototype of an adaptive disruption predictor for JET based on fuzzy logic and regression trees" *Nucl. Fusion* ***48*** *(2008)*

[17]. Bishop C.M. "Pattern Recognition and Machine Learning" *Springer publish. 2006*

[18]. Martinez W.L. and A.R.Martinez "Exploratory data Analysis with Matlab" *Chapman&Hall 2005*

[19]. Wegman E.J. and J. Shen. "Three-dimensional Andrews plots and the grand tour" *Proceedings of the 25th Symposium on the Interface, 284-288. 1993.*

[20]. Asimov D. "The grand tour: a tool for viewing multidimensional data" S*IAM Journal on Scientific and Statistical Computing V.1, pg. 128-143, Jan. 1985*

[21]. Gauthier E., et al "ITER-like wide-angle infrared thermography and visible observation diagnostic using reflective optics" *SOFT 2006*

[22]. Lipschultz B., "MARFE: an edge plasma phenomenon," *Nuclear Fusion*, **24**, n8, Aug. 1984, p 977-89

SIGNAL NAME	**UNIT**
Plasma current I_{pla}	[A]
Mode Lock Amplitude *Loca*	[T]
Plasma density *Dens*	[m^{-3}]
Total Input Power P_{inp}	[W]
Plasma Internal Inductance L_i	
Stored Diamag. Energy Derivative dW_{dia}/dt	[W]
Safety factor at 95% of minor radius q_{95}	
Poloidal beta β_p	
Net power P_{net}	[W]

Table I: List of the signals used as predictors for the classification trees.

$[t_D-440, t_D-400]$		$[t_D-380, t_D-340]$		$[t_D-320, t_D-280]$		$[t_D-260, t_D-220]$		$[t_D-200, t_D-160]$		$[t_D-140, t_D-100]$		$[t_D-80, t_D-40]$	
Variable	Rank	Variable	Rank	Variable	Rank	Variable	Rank	Variable	Rank	Variable	Rank	Variable	Rank
P_{net}	100	*Ipla*	100	P_{net}	100	dW_{dia}/dt	100	dW_{dia}/dt	100	dW_{dia}/dt	100	dW_{dia}/dt	100
Ipla	84,83	P_{net}	24,01	*Ipla*	94,96	P_{net}	51,00	*Dens*	30,63	β_p	23,87	*Ipla*	6,65
dW_{dia}/dt	76,10	dW_{dia}/dt	21,88	dW_{dia}/dt	43,71	β_p	47,75	β_p	27,25	l_i	9,72	*Loca*	4,55
β_p	31,28	P_{inp}	21,63	*Loca*	23,80	l_i	28,44	l_i	9,53	*Loca*	7,21	l_i	1,50
l_i	30,64	*Dens*	10,98	β_p	18,32	*Ipla*	19,36	P_{net}	9,19	*Dens*	4,50	*Dens*	1,34
Dens	14,02	l_i	9,61	*Dens*	16,39	*Dens*	13,36	*Ipla*	6,27	q_{95}	2,40	β_p	0,69
P_{inp}	7,49	*Loca*	8,64	q_{95}	12,97	*Loca*	10,62	*Loca*	4,27	*Ipla*	1,24	P_{net}	0,44
q_{95}	0,0	q_{95}	4,70	l_i	6,49	P_{inp}	5,58	q_{95}	3,80	P_{net}	0,00	P_{inp}	0,0
Loca	0,0	β_p	2,19	P_{inp}	0,0	q_{95}	3,40	P_{inp}	0,0	P_{inp}	0,00	q_{95}	0,0

Table II: Ranking of the most important signals for disruption prediction as calculated by CART for the various time intervals. The numbers between 0 and 100 simply indicate the relative importance of the signals in the various intervals but do not have any quantitative absolute meaning. The table includes only the primary splitters.

SNR 20		SNR 25		SNR 30		SNR 35		SNR 40		SNR 50	
Variable	Rank	Variable	Rank	Variable	Rank	Variable	Rank	Variable	Rank	Variable	Rank
dW_{dia}/dt	100	dW_{dia}/dt	100	dW_{dia}/dt	100	dW_{dia}/dt	100	dW_{dia}/dt	100	dW_{dia}/dt	100
β_p	10,89	P_{net}	7,73	β_p	17,54	P_{net}	7,68	β_p	17,80	β_p	17,12
P_{net}	8,39	*Ipla*	7,05	*Dens*	6,28	β_p	7,34	l_i	9,03	l_i	9,95
Dens	7,23	β_p	6,87	q_{95}	5,96	*Ipla*	4,59	*Dens*	4,86	*Dens*	4,21
Ipla	3,47	l_i	6,47	l_i	3,52	P_{inp}	4,28	P_{net}	3,13	P_{net}	3,43
l_i	1,63	P_{inp}	0,80	P_{net}	1,55	l_i	4,02	*Ipla*	2,19	*Loca*	2,39
q_{95}	0,00	*Loca*	0,63	*Ipla*	1,31	q_{95}	0,00	q_{95}	1,24	*Ipla*	2,10
P_{inp}	0,00	q_{95}	0,35	P_{inp}	0,70	*Loca*	0,00	*Loca*	0,86	P_{inp}	1,12
Loca	0,00	*Dens*	0,00	*Loca*	0,00	*Dens*	0,00	P_{inp}	0,00	q_{95}	0,00

Table III: Ranking of the various signals (primary splitters) as calculated by CART for various S/N ratios for the interval $[t_D-140, t_D-100]$

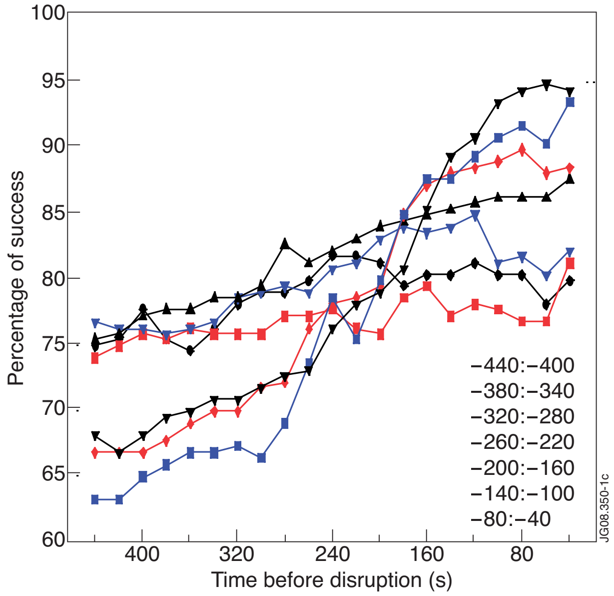

Figure 1: Overall percentage of success (including both safe and disruptive pulses) of the various classification trees versus time to disruption.

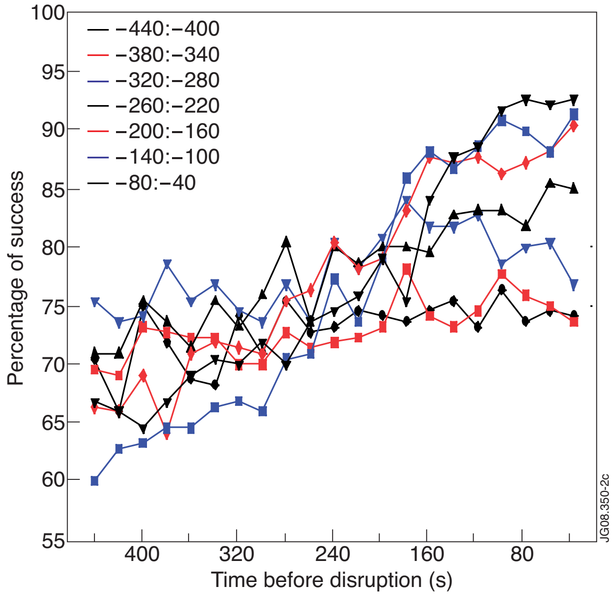

Figure 2: Overall percentage of success (including both safe and disruptive pulses) of the various classification trees versus time to disruption after summing Gaussian white noise to the input signals to obtain a S/N ratio of 30.

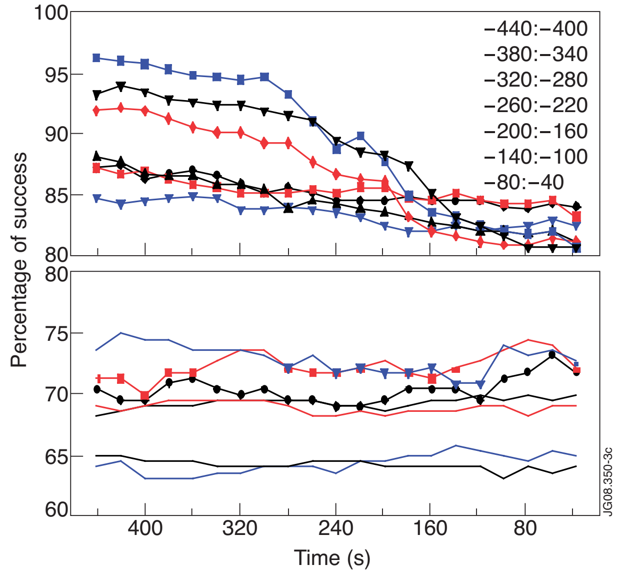

Figure 3: Percentage of missed and false alarms of the various classification trees versus time to disruption for the same database as in figure 1.

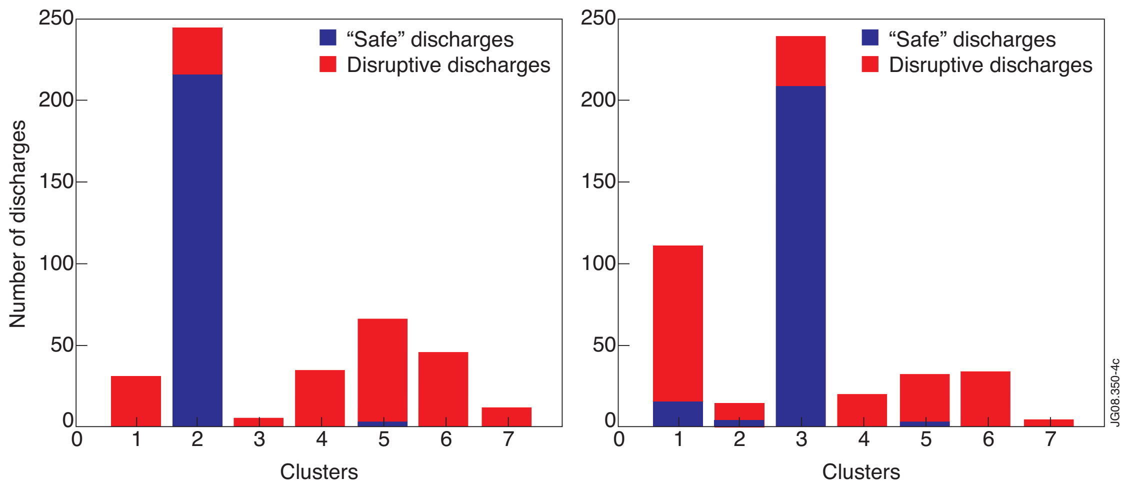

Figure 4: Distribution of the discharges using K-Means for two different time intervals.

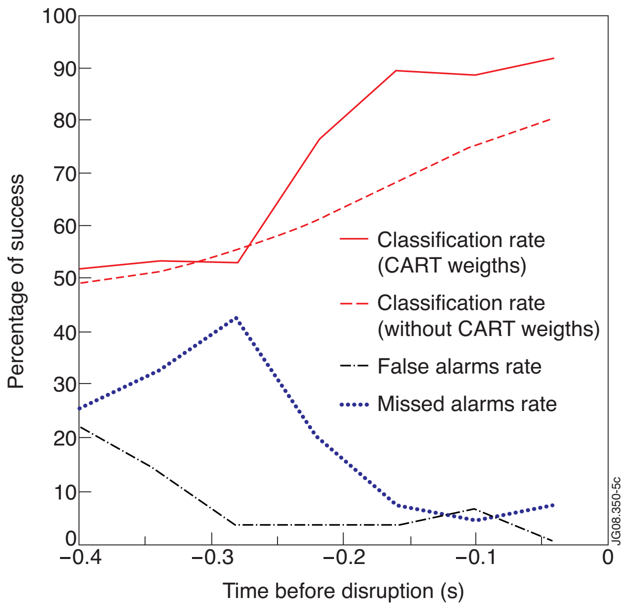

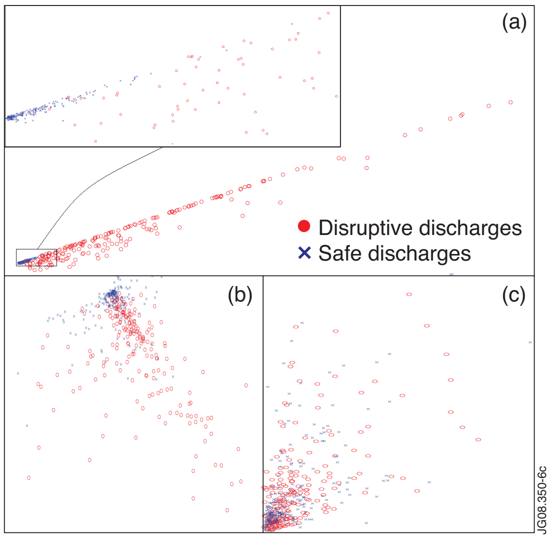

Figure 5: Overall percentage of success of the K-means. The performance decreases significantly for times earlier than 180 ms before the disruption.

Figure 6 Two plane projections of JET database obtained with the pseudo Grand Tour technique. a) time interval [t_D-0.08, t_D-0.04] b) time interval [t_D-0.26, t_D-0.22 s] c) time interval [t_D -0.44, t_D -0.4].

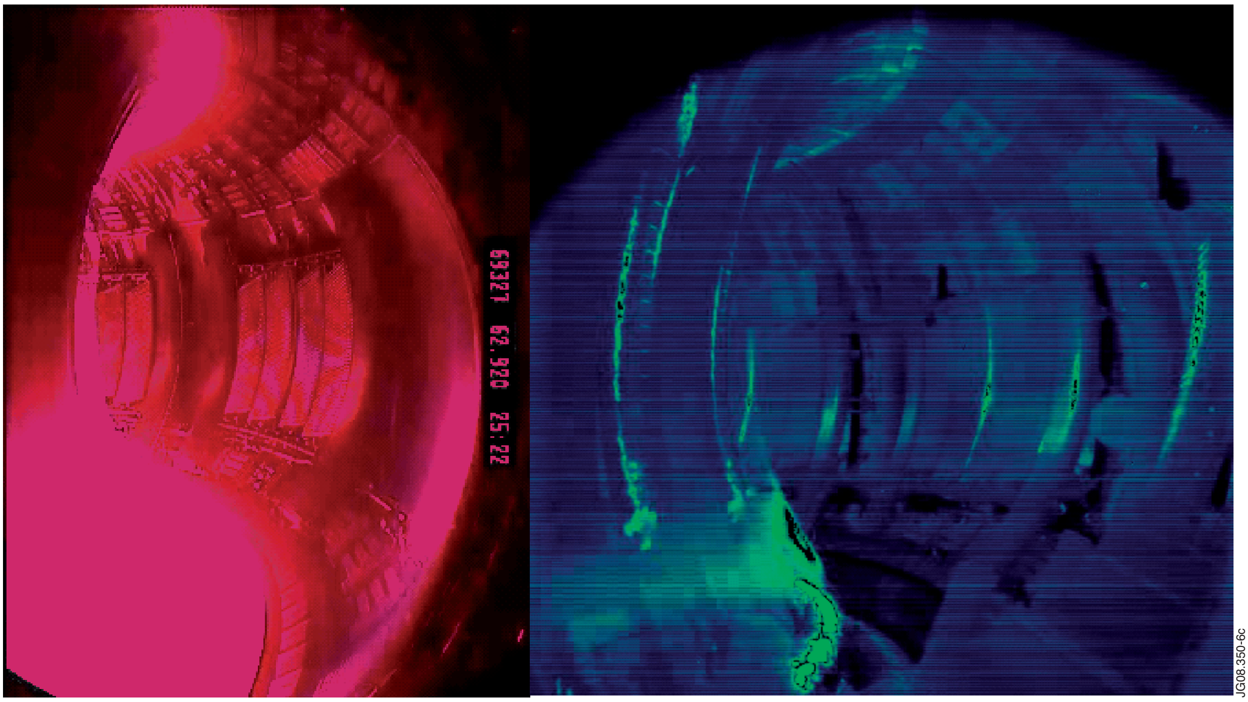

Figure 7: Visible (left) and Infrared (right) views of JET outer region more than one second before a density limit discharge, showing the early phase of a MARFE instability. Automatic detection of these manifestations of a MARFE can be obtained with modern image processing techniques.