ELSEVIER

Journal of Nuclear Materials 290–293 (2001) 1045–1051

journal of nuclear materials

www.elsevier.nl/locate/jnucmat

Section 10. Off-normal heat deposition

Prediction and mitigation of disruptions in ASDEX Upgrade

G. Pautasso [a,*], S. Egorov [b], Ch. Tichmann [a], J.C. Fuchs [a], A. Herrmann [a], M. Maraschek [a], F. Mast [a], V. Mertens [a], I. Perchermeier [a], C.G. Windsor [c], T. Zehetbauer [a], ASDEX Upgrade Team

[a] *Max-Planck-Institut für Plasmaphysik, EURATOM Association,Garching, Germany*
[b] *Technical University,St. Petersburg, Russian Federation*
[c] *Euratom-UKAEA Fusion Association, Culham Science Center, OX14 3DB, UK*

Abstract

Disruptions in tokamaks are instabilities events which can damage the machine components. The avoidance and mitigation of these events is desirable in present machines as well as in Next Step devices (such as ITER). A neural network has been developed to predict the occurrence of disruptions caused by edge cooling mechanisms in ASDEX Upgrade. The network works reliably and is able to predict the majority (85%) of the disruptions. The neural network has been trained to predict the time interval up to the disruption and this makes it suitable to be used on-line either to avoid disruptions (by means of auxiliary heating and reduction of gas puffing) or to mitigate the unavoidable ones. For this last purpose, a solid pellet injector has been developed and tested; the injected impurity pellets have been shown to reduce the vertical forces and the conductive fluxes to the divertor.

Keywords: Plasma properties

1. Introduction

Tokamak plasmas may undergo a major instability called disruption. Independent of its cause, a disruption is triggered by the non-linear growth and interaction of MHD helical modes which leads to the fast and irreversible loss of thermal confinement. Its consequences are a large heat flux to the limiter or divertor structures (thermal quench) and the Ohmic dissipation of the plasma current due to the increased plasma resistivity (current quench). In addition, in elongated configurations, the plasma may lose its vertical stability after the thermal quench or because of control errors. Both the decaying plasma current and the moving current column can induce large currents in the machine structures and cause electromagnetic forces which can damage them. It has also been observed that the large voltage associated with the current quench may generate runaway electrons which, interacting with the plasma facing surfaces, can cause their erosion or failure.

Existing machines have been designed to withstand many disruptions, since one of their purposes is the investigation of disruptive boundaries. Nevertheless the asymmetric distribution of the loads and the limited lifetime of mechanical components are occasionally responsible for damage to the machine. The avoidance and mitigation of disruptions is therefore strongly recommended in the existing tokamaks. In a tokamak-like ITER, where the disruptive thermal load is predicted to evaporate and melt 10–100 μm of localized portions of the divertor target plate [1], avoidance and mitigation systems are necessary to prolong the lifetime of the machine components.

A pre-requisite for any avoidance or mitigation action is a system which is able to recognize a forthcoming plasma disruption. Disruptions have different physical causes and happen in different regions of the operational space; according to the parameter region in which they take place, they are classified as: ideal beta-, density-, low-q limits, locked mode at low density, neoclassical locked mode and low-q, etc. Most of the plasma

* Corresponding author. Tel.: +49-89 3299 1724.
E-mail address: pautasso@ipp.mpg.de (G. Pautasso).

0022-3115/01/$ - see front matter
PII: S0022-3115(00)00546-8

disruptions in ASDEX Upgrade are caused by edge cooling mechanisms [2] (for example, high-electron or/and impurity density) and happen in a plasma parameter regime (poor L-mode confinement), which is far away from the desired operational space (H-mode, high-β). In addition, they are usually announced by well-identified precursors (detachment, MARFE, growth and locking of resistive tearing modes) which can be detected by available diagnostics. Nevertheless none of these precursors alone can assure that a disruption is predicted early enough. For some disruptions, MARFE and locked modes have a very short life and their detection is too late for any avoidance or mitigation action to be able to take place. Besides, they do not necessarily lead to a disruption and may even disappear after control actions are taken.

In this work we make use of experience gathered during several years of operation and from disruption studies in order to train a neural network to recognize a forthcoming disruption.

Along with the neural network, an injector for solid impurity pellets (the so-called 'killer' pellets) has been developed for ASDEX Upgrade. The injector is configured to be triggered by the network when an imminent disruption has been predicted. The injection of impurity pellets in a plasma has already been proved to be a method to reduce mechanical forces and thermal loads during disruptions in present tokamaks [3].

2. Disruption prediction with a neural network

2.1. Artificial neural network

An artificial neural network is a non-linear function mapping one multi-dimensional space, $\{\vec{x}\}$, into another one, $\{\vec{z}\}$. For an overview on the subject see, for example, Ref. [4]. This function has a pre-defined structure but contains several parameters which are going to be determined during the training. The training consists of the evaluation of the parameters which minimize the difference between the target output, $\vec{t}$, and the network output, $\vec{z}$.

Among several possible structures of the network, we use a, so-called, feed-forward multilayer (two layers, in this work) perceptron. This kind of network is known to approximate arbitrarily well any continuous multi-dimensional mapping [4]. The hth component of the vector output ($h = 1, nz$), can be written as

$$z_h = \sum_{i=1}^{ny} WY_{hi} y_i + WY_{h0}, \qquad y_i = \mathscr{F}\left(\sum_{j=1}^{nx} WX_{ij} x_j + WX_{i0} \right), \tag{1}$$

where y_i is the ith component of the output of the first layer; nx, ny and nz the dimension of the input vector, the number of the hidden neurons and the dimension of the network output, respectively, $\mathscr{F}(a)$ is a non-linear function; in this work we chose it to be the sigmoidal function: $(1 + \exp(-a))^{-1}$. The input variable to the first layer is transformed first linearly, by means of the matrix WX; then the first layer generates ny non-linear function which are combined together linearly by the second layer (matrix WY) in order to fit the given target space.

The values of the $(nx + 1) \times ny + (ny + 1) \times nz$ unknown elements of the matrixes WX and WY are found by minimizing an error function of the type

$$E = 0.5 \times \sum_{k=1}^{N} [\vec{z}(\vec{x}_{(k)}, WX, WY) - \vec{t}_{(k)}]^2, \tag{2}$$

where the sum is extended to the whole training set.

Several minimization methods are available; they consist in evaluating the derivatives of E with respect to the elements of the WX and WY matrixes and in correcting the unknown parameters using an appropriate learning rate, δ, in the following way:

$$WX_{ij}^{(n+1)} - WX_{ij}^{(n)} = -\delta \frac{\partial E}{\partial WX_{ij}} \tag{3}$$

(n) is the iteration number.

2.2. Review of related works

Neural networks have already been shown to be suitable for the prediction of disruptions in tokamaks.

The high-β disruption boundary was modelled in DIII-D [5] using 33 plasma parameters as input to a two-hidden-layers feed-forward perceptron. The output of the network was the value of $\beta_N = \beta_t(aB)/I_p$ at the time of disruption for that shot. The neutral beam heating phases from 84 disruptive discharges were used for the training and validation. The ratio between the network output and the actual value of β_N can be used as β-limit detection parameter. The probability of correct disruption prediction and the probability of false alarm can be respectively maximized and minimized by choosing a proper alarm threshold for this ratio. The network was shown to perform rather well for a set of 28 disruptions not used for training, with a 90% probability of disruption detection and less than 20% probability of false alarms.

For JET [6] a disruption alarm was trained to distinguish between pre-disruptive (output = 1) and stable plasmas (output = 0). A one-hidden-layer feed-forward perceptron was trained with 360 disruptive shots. The best performance was achieved with a network using the following seven input parameters: locked mode, density, input power, radiated power, q_{95}, l_i, β_p and derivative of the stored energy. In this case the probability of dis-

ruption detection, tested on the training data, was of 90% with a 5% probability of false alarms.

Unpublished work was carried out with the COMPASS-D database [7]; 215 shots and 7 input variables (current, vertical and radial field, toroidal field, loop voltage, total energy and β_p) were used to train a one-hidden-layer feed-forward perceptron to distinguish between the disruptive plasmas (data selected 10 ms before disruption, output = 1) and non-disruptive conditions (output = 0). The 'leave-one-out' method was employed with 10 hidden units for training and testing: all the disruptions were recognized, and just two of the non-disruptive shots were misclassified.

Exploratory work was done with TEXT disruptions [8]. Using time-series-prediction a network was trained with a soft-X-rays time trace from one shot and then used to predict the disruption in other two discharges (3 ms in advance). A clustering method was used to define a multi-D region, where locked modes are born and lead to disruptions in TCV discharges [9].

2.3. Input and output data for ASDEX Upgrade

For the training a neural network needs a large database of input vectors and associated outputs. The database used to train and validate the network for ASDEX Upgrade consists of discharges in the shot range 10,000–11,400 with a lower-single-null plasma in flat-top. Some shots were neglected owing to incomplete or incorrect measurements. Disruptions caused by injection of killer pellets or following a VDE were also excluded. The pre-disruption phase of these shots and complementary non-disruptive phases were selected to generate the database for the training of the network. The pre-disruption phase was defined as the L-mode phase following the H-mode phase before the disruption or as the phase starting just before a MARFE and ending with the disruption, for a plasma which has been for longer than 0.8 s in L-mode.

The output of the network was chosen to be the time interval up to the disruption; this makes the output a variable that is easy to interpret and which can be used as a flexible trigger to avoid or mitigate a disruption. For the pre-disruption phases the output ranges from 0 to 0.8 s; for the non-disruptive phases it is set equal to 0.8 s.

The input consists of the time histories of several plasma parameters describing the plasma regime in the discharge flat-top. The choice of input variables was the result of a compromise between the physics and the availability of the data in real-time, since the network will be used for the real-time control of the discharge. A preliminary set of 30 input variables was successively reduced during the training by eliminating the variables which did not significantly contribute to the output. This is done firstly by evaluating the magnitude of the derivative of the output with respect to the input variables summed over the whole training set, i.e., $\sum_{k=1}^{N} \partial z_{(k)}/\partial x_i$ [5], which is a statistical measure of the sensitivity of the output to a given input parameter (our output has dimension 1). Secondly the network is trained again with reduced inputs until its performance starts to deteriorate significantly.

The plasma parameters used as input to the network presented in this work are:

1. the safety factor, q_{95};
2. the plasma internal inductivity, l_i;
3. the plasma density divided by the Greenwald limit, n_e/j;
4. the radiated fraction of the input power;
5. the energy confinement time normalized with the time predicted by a scaling law (Ref. [10, Eq. 12]);
6. a MARFE signal, using two divertor bolometer channels, indicating the presence (= 1) or the absence (= 0) of a MARFE;
7. the locked mode signal, using the measurements from the saddle coils, indicating the absence (= 1) or the presence (= 0) of a locked mode;
8. the normalized beta, β_N;
9. the time derivatives of the variables no. 2,3,4,5 and 8.

The input variables were further normalized to the $[0, 1]$ interval, time averaged over 25 ms and fed to the network every 2.5 ms. The database of 15,373 data points from 99 disruptive shots and 12,636 data points from non-disruptive shots was then used for training.

The listed input variables were chosen to be dimensionless in order to allow for the establishment of a inter-machine network for disruption prediction. In fact, a big disadvantage of the use of a network for disruption prediction is the large number of these events needed for the training. This prediction method could become much more attractive if a network trained with data from present tokamaks could be used on a new machine. Such an approach is planned, in collaboration with UKAEA, involving JET and COMPASS-D data, but is not discussed here.

2.4. Training

The training consisted in minimizing the function:

$$E = \sqrt{N^{-1} \sum_{k=1}^{N} \left[1 - \Delta t_{NN(k)}/\Delta t_{(k)}\right]^2}, \tag{4}$$

where Δt and Δt_{NN} are the measured and the predicted time to disruption, respectively.

The training was carried out with the Neural Network Toolbox of Matlab and in particular with a self-modified version of the subroutine employing the

Levenberg–Marquardt algorithm [11] for the minimization of Eq. (4)

Several networks with different number of hidden neurons were trained and tested. The data were divided in a training set (3/4 of the data) and a validation set. The training lasted typically 50–70 cycles and was stopped after checking that the function given by Eq. (4) and calculated for the validation set had reached a minimum or a stationary value. Several networks with a typical RMS error of 15% were trained and validated once more on a large number of shots outside of the training database.

2.5. Network predictions

The performance of several 'good' networks was tested off-line in the flat-top phase of 500 shots, in the shot range 11,475–12,170, which were not included in the training.

Only the disruptions in flat-top or within the first 100 ms of the plasma ramp-down-phase, with an X-point-down during the last 50 ms preceding the disruption, were considered in this analysis. The shots with large impurity gas injection (which lead quickly to disruption) or incomplete measurements and the disruptions caused by VDE or killer pellets were disregarded. All other shots were retained.

The disruption alarm was defined as $\Delta t_{NN} < 50$ ms for 7.5 ms. The prediction of the disruption was defined to be correct if the disruption alarm was activated in the time interval $[t_{\mathrm{disr}} - 500\ \mathrm{ms},\ t_{\mathrm{disr}}]$; a disruption alarm in shots without a disruption or more than 500 ms before a disruption was considered a false alarm; a disruption was not recognized when the disruption alarm was not activated.

The results of the analysis were as follows:

- 85% of the disruptions were recognized; the network did not recognize the disruption in 10 of the 65 disruptive shots (15%);
- the network produced one or more false alarms in 5 of the 500 shots (1%).

Examples of the network prediction are shown in Fig. 1 for (a) a density limit disruption and (b) a disruption following an impurity event. The disruptions, that were not recognized, happened very fast (because of impurity events, UFOs) or had been poorly represented in the training database (impurity accumulation or disruption after a locked mode phase in ITB experiments). The false alarms were caused by pre-disruption phases from which the plasma recovered thanks to an action taken by the feed-back system (increasing the heating power or closing a gas valve) or by plasmas, close to the disruptive boundaries, which did not disrupt (see Fig. 2).

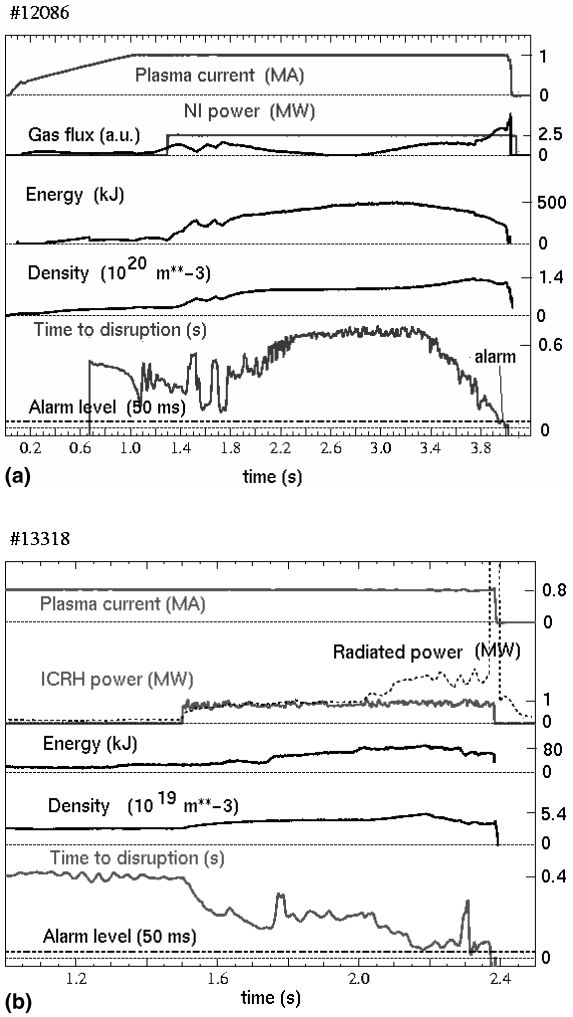

Fig. 1. (a) Network prediction for a density limit disruption; the alarm conditions are reached at $t = 3.955$ s, 65 ms before disruption. (b) Network prediction for a disruption following an impurity event; the alarm conditions are reached at $t = 2.320$ s, 60 ms before disruption.

3. Use of the neural network

The neural network is presently being installed to work on-line as part of the control system of ASDEX Upgrade. Its input data are sampled, time averaged and normalized as it is done in the off-line program. After verifying that the on-line system generates the same result as the off-line program, the network output will be used to avoid or mitigate disruptions.

The performance of the network as a system to avoid disruptions can only be checked by using it on-line. Nevertheless we can already compare the off-line network output with the output of the detachment detection algorithm, which is already employed in ASDEX

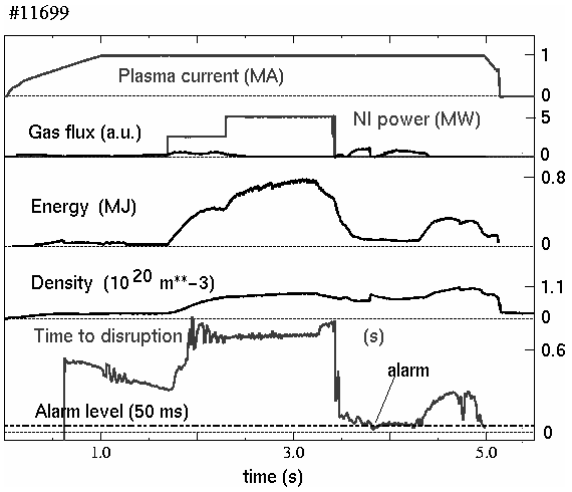

Fig. 2. False alarm due to poor L-mode confinement after turning off the NI power at high density.

Upgrade to avoid disruptions. This algorithm uses the ratio between the H_α emission from the inner and outer divertor plates; this ratio is compared to a threshold value and used to turn on heating and close gas valves. Fig. 3 shows an experiment where, by modulating the gas puff, the detachment phases appeared periodically and the detachment algorithm reacted to these phases by turning off the gas valve (at time 2.95, 3.4 and 4.55 s). The network output remains above the disruption alarm (set to 50 ms) up to 4.8 s and it could have been used to control the gas valve in a similar way, if we had chosen an avoidance threshold of 120 ms.

For disruption mitigation we foresee the use of an impurity pellet injector; we have chosen the limit of

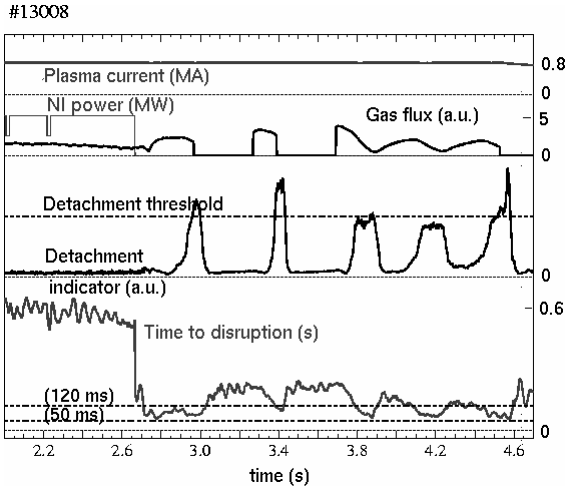

Fig. 3. Comparison between the detachment indicator (used here to control a gas valve) and the network prediction.

50 ms as disruption alarm to trigger it after verifying that, with this threshold alarm, the percentage of false alarms is acceptably low and that the percentage of predicted disruptions remains significant. This value of the alarm threshold takes into account that: (1) the injector is now situated 3.5 m from the plasma and the time delay between trigger and pellet ablation is of the order of 15 ms; (2) the network output must be below 50 ms for at least 7.5 ms before triggering the disruption alarm; (3) a margin of a few tens of ms must be allowed to compensate the error affecting the network prediction.

4. Impurity pellet injection

4.1. Background

The injection of an impurity pellet allows a soft plasma termination. The reduction of the thermal load to the target plates during the thermal quench is mostly due to the cooling of the plasma by the ablation and ionization of the pellet. The reduction of the vertical forces in disruptions followed by VDE is due to the faster decay of the current after pellet injection and to the smaller currents induced in the passive stabilizer and in the halo region [12]. Although impurity pellet injection mitigates disruptions in ASDEX Upgrade and the development of this technique is pursued to protect the machine, its application to ITER is not straightforward. As already discussed in Ref. [13] the injection of high-Z impurities in ITER is expected to generate an unacceptable runaway current. The massive injection of frozen pellets or of a liquid jet of deuterium has been proposed as alternative for an ITER shutdown scenario.

Earlier experiments in ASDEX Upgrade [12] used frozen neon pellets, prepared in a cryostat and injected with a centrifuge. Since a cryo-system is not ideal for a stand-by application, typical of a disruption mitigation system, a new pellet gun for the injection of solid pellets was developed. The injector was designed, built and brought into operation on ASDEX Upgrade by the second author in August 1998 and the experiments were planned on the basis of the experience gathered in the previous years. Since then, a guiding tube and additional fast valves have been installed between injector and vessel in order to reduce the propellant gas, which reaches the plasma before the pellet, and its influence on the plasma.

4.2. Technical specifications

The pellets used in this experiment were made of silicon (Si) or titanium (Ti) powder mixed with molten polyethylene (PE) and then extruded in the form of a cylinder. Spheres of 2 mm of diameter were then obtained by compression and heating. Si and Ti were

Table 1
Weight and composition of the pellets used

Type	Weight (mg)	N_I (I = Si or Ti)	Grains (μm)	N_{CH_2}
Si (60%)-PE	7 (4.2 Si)	0.9×10^{20}	15–20 (Si)	1.2×10^{20}
Ti (72%)-PE	10 (7.2 Ti)	0.9×10^{20}	10 (Ti)	1.2×10^{20}

chosen because of their 'good' radiative properties and their acceptability in the machine; PE because of its relatively low sublimation energy and its acceptable chemical composition. These pellets were injected from the low-field side and with a velocity of in the range 200–250 m/s; their weight and composition are listed in Table 1.

4.3. Experimental data

Several experiments have been carried out with the injector. We subdivide them in two main groups and discuss the experimental observations accordingly:

(1) pellets injected in 'healthy' plasmas at a pre-programmed time, and

(2) pellets injected with the locked mode trigger in plasmas with locked modes and pre-disruption characteristics.

Group (1). At the beginning of the experimental campaign the pellets were injected at a pre-programmed time in order to test the injector and collect data in a variety of plasma conditions. The target plasmas did not have pre-disruption characteristics and had a relatively high thermal energy. Independent of the plasma thermal energy the plasma disrupts typically 2 ms after the appearance of the pellet at the plasma boundary (the pellet goes half way through the plasma). The pellets of this group were injected in plasmas with energies in the range 50–850 kJ. In plasmas with energies above 500 kJ the pellets seem to ablate almost completely before they leave the $q = 1$ surface on the high-field side. In plasmas with an energy below 500 kJ the pellets do not ablate completely; a residual fragment is always seen by the CCD camera on the high-field side. In the plasmas with energies around 100 kJ the pellets poorly ablate; the SXR cameras show ablation limited to the plasma central region. These results indicate that these pellets are oversized for these plasmas. We are particularly interested in these low-energy plasmas because the pre-disruption plasmas have typically a low energy in ASDEX Upgrade. Nevertheless we prefer to work with large pellets which contain enough impurities to mitigate the disruption in the case of high-energy plasmas.

Group (2). In successive experimental campaigns pellets were injected using the locked mode trigger. Several of these pellets were injected in the presence of a locked mode in plasmas with pre-disruption characteristics. Some of these pellets reached the plasma after the thermal quench. The pellets which reached the plasma within a few ms after the onset of the disruption sufficiently ablated during the immediate post-disruption phase due to the presence of a large loop voltage and consequently fast electrons. Nevertheless there is no evidence of long lived runaway electrons generated by the pellets. In very few discharges terminated with killer pellets, the ECE measurements of temperature show a short (1 ms) burst of supra-thermal electrons during the ablation; this is followed by bursts of the hard X-ray emission which last 2–3 ms and do not survive during the whole current quench (5–10 ms). The remaining pellets were injected in plasmas with neoclassical locked modes, high-β and no pre-disruption characteristics; the conclusions for group (1) also apply to these pellets.

4.4. Disruption forces

Disruptions caused by the injection of impurity pellets had reduced halo currents and vertical forces (typically 50% less). It was shown in Ref. [14] that the forces associated with the halo currents are the major contributor to the maximum vertical force (F_z) acting on the vessel support rods in ASDEX Upgrade and that $F_z \simeq 0.4\, I_h B_t$ (I_h is the maximum halo current toroidally averaged and B_t is the toroidal magnetic field). Fig. 4 shows the distribution of the magnitude of the ratio I_h/I_p (I_p is the plasma current) as function of the plasma energy. The

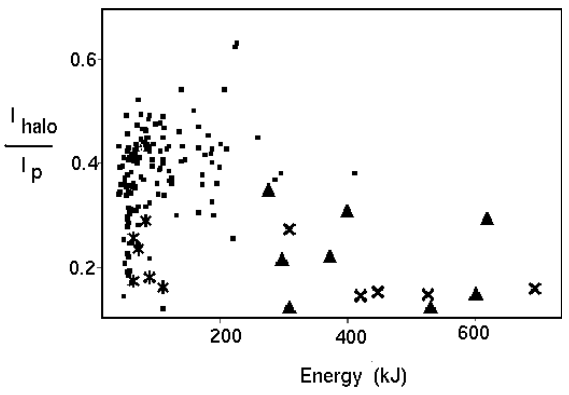

Fig. 4. Ratio between halo current and plasma current versus plasma energy: shots without pellet (box), pellets injected with timer (triangle), with locked mode trigger in non-disruptive plasmas (×), with the locked mode trigger about the time of the thermal quench (star).

picture shows that most of the plasmas disrupted with an energy in the range 0–300 kJ. All the pellets injected with the timer reduced the mechanical forces but were also in an energy range which is not typical for ASDEX Upgrade disruptions. In the shots, in which the pellets reached the plasma during the disruption, the halo currents were also in the lower range. The large scatter in the values of I_h/I_p at low energy is due to the different impurity contamination of the plasma at the disruption time. It has also been observed that strong impurity puff in high-q, low-energy plasmas and impurity events significantly reduce the disruption loads.

4.5. Thermal loads

Up to 100% of the thermal energy is typically deposited on the divertor plates during the thermal quench; during the current quench less than 30% of the remaining magnetic energy reaches the divertor and mostly in form of radiation. The impurity pellets reduce the thermal load to the divertor by suppressing the conductive heat flux during the thermal quench and increasing the radiated power during the whole current quench. Fig. 5 shows the time traces of the power deposited to the divertor and of the radiated power in two density limit disruptions without and with the mitigating effect of a killer pellet. Shot 13302 has a minor disruption at $t = 4.46$ s and the major disruption at $t = 4.5$ s; respectively, 60 and 40 MW of power are seen reaching the divertor plate (shadow region) after these events.

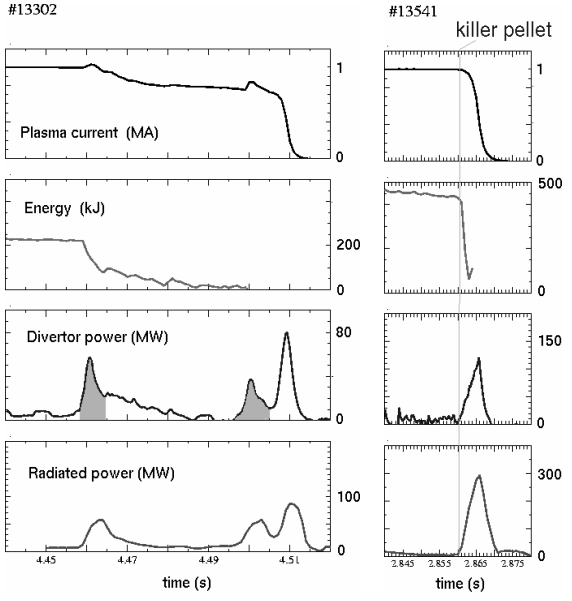

Fig. 5. Plasma current, thermal energy, divertor target power and radiated energy in two density limit disruptions without (13,302) and with (13,541) injection of a killer pellet.

The thermal quench is suppressed in shot 13541, where a killer pellet was injected with the locked mode trigger and ablates at $t = 2.86$ s. Correspondently the radiated power has a maximum of 300 MW.

5. Conclusions

The present work shows that a disruption recognition system based on a neural network is feasible, relatively reliable (85% predicted events and 1% false predictions) and it motivates its wider application. The feed-forward calculation of the output of a trained network is simple and fast, making it suitable for real-time implementation. In addition, the chosen output of the network was the time to disruption, which is a variable that can be used for avoidance as well as for mitigation of disruptions. A new impurity pellet injector has been extensively tested with Si– or Ti–PE pellets: it has been proved to work in a reliable way and to fulfill its purpose of mitigating the disruption loads. The injector is planned to be routinely on stand-by and to be triggered by the disruption alarm system described above.

Acknowledgements

The authors G.P. and Ch.T. acknowledge fruitful discussions with C. Morabito and M. Versaci.

References

[1] G. Federici et al., these Proceedings.
[2] G. Pautasso et al., EPS Conference on Controlled Fusion and Plasma Physics, Prague, 1998.
[3] Chapter 3 of: ITER Physics Expert Group on Disruptions et al., Nucl. Fus. 39 (1999) 2251.
[4] C.M. Bishop, Neural Networks for Pattern Recognition, Clarendon, Oxford.
[5] D. Wroblewki et al., Nucl. Fus. 37 (1997) 725.
[6] F. Milani, PhD thesis, University of Aston, Birmingham, December 1998.
[7] C.G. Windsor, N.M. Ershov, A.M. Popov, T.N. Todd, Disruption prediction in COMPASS-D using a neural network, unpublished report.
[8] A. Vannucci et al., Nucl. Fus. 39 (1999) 255.
[9] Y. Martin et al., 25th EPS Conference on Controlled Fusion and Plasma Physics, Prague, 1998.
[10] P. Franzen et al., Fusion Technol. 33 (1998) 84.
[11] H. Demuth, M. Beale, Neural Network Toolbox User's Guide, Version 3.0, 1998.
[12] G. Pautasso et al., Nucl. Fus. 36 (1996) 1291.
[13] S. Putvinski et al., J. Nucl. Mater. 241–243 (1997) 316.
[14] O. Gruber et al., Plasma Physics and Controlled Nuclear nt. Conference on Fusion Research 1994 (Proceedings of the 15th International Conference, Seville, 1994), vol. 1, IAEA, Vienna, 1995, p. 675.