


Disruption event characterization and forecasting in tokamaks


Cite as: Phys. Plasmas **30**, 032506 (2023); <https://doi.org/10.1063/5.0133825>

Submitted: 06 November 2022 • Accepted: 13 February 2023 • Published Online: 17 March 2023

 S. A. Sabbagh,  J. W. Berkery,  Y. S. Park, et al.

COLLECTIONS

 This paper was selected as Featured

 This paper was selected as Scilight



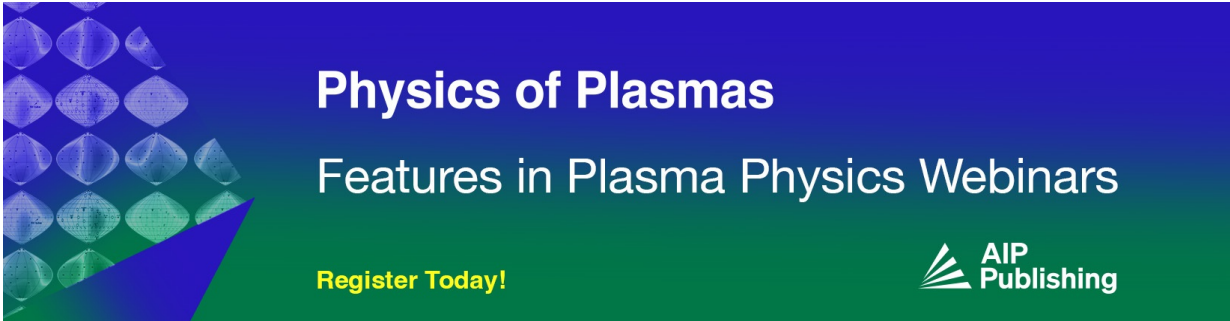
View Online



Export Citation




CrossMark



Physics of Plasmas
Features in Plasma Physics Webinars

Register Today!



Disruption event characterization and forecasting in tokamaks



Cite as: Phys. Plasmas **30**, 032506 (2023); doi: 10.1063/5.0133825

Submitted: 6 November 2022 · Accepted: 13 February 2023 ·

Published Online: 17 March 2023



View Online



Export Citation



CrossMark

S. A. Sabbagh,^{1,a)} J. W. Berkery,² Y. S. Park,¹ J. Butt,¹ J. D. Riquezes,¹ J. G. Bak,³ R. E. Bell,² L. Delgado-Aparicio,² S. P. Gerhardt,² C. J. Ham,⁴ J. Hollocombe,⁴ J. W. Lee,³ J. Kim,³ A. Kirk,⁴ J. Ko,³ W. H. Ko,³ L. Kogan,⁴ B. P. LeBlanc,² J. H. Lee,³ A. Thornton,⁴ and S. W. Yoon³

AFFILIATIONS

¹Department of Applied Physics and Applied Mathematics, Columbia University, New York, New York 10027, USA

²Princeton Plasma Physics Laboratory, Princeton University, Princeton, New Jersey 08543, USA

³Korea Institute of Fusion Energy, Daejeon 34133, South Korea

⁴UKAEA, Abingdon OX14 3DB, United Kingdom

^{a)} Author to whom correspondence should be addressed: sabbagh@pppl.gov

ABSTRACT

Disruption prediction and avoidance is a critical need for next-step tokamaks, such as ITER. Disruption Event Characterization and Forecasting (DECAF) research fully automates analysis of tokamak data to determine chains of events that lead to disruptions and to forecast their evolution allowing sufficient time for mitigation or complete avoidance of the disruption. Disruption event chains related to local rotating or global magnetohydrodynamic (MHD) modes and vertical instability are examined with warnings issued for many off-normal physics events, including density limits, plasma dynamics, confinement transitions, and profile variations. Along with Greenwald density limit evaluation, a local radiative island power balance theory is evaluated and compared to the observation of island growth. Automated decomposition and analysis of rotating tearing modes produce physical event chains leading to disruptions. A total MHD state warning model comprised of 15 separate criteria produces a disruption forecast about 180 ms before a standard locked mode detector warning. Single DECAF event analyses have begun on KSTAR, MAST, and NSTX/-U databases with thousands of shot seconds of device operation using from 0.5 to 1×10^6 tested sample times per device. An initial multi-device database comparison illustrates a highly important result that plasma disruptivity does not need to increase as β_N increases. Global MHD instabilities, such as resistive wall modes (RWMs), can give the briefest time period of warning before disruption compared to other physics events. In an NSTX database with unstable RWMs, the mode onset, loss of boundary and current control, and disruption event warnings are found in all cases and vertical displacement events are found in 91% of cases. An initial time-dependent reduced physics model of kinetic RWM stabilization created to forecast the disruption chain predicts instability 84% of the time for experimentally unstable cases with a relatively low false positive rate. Instances of the disruption event chain analysis illustrate dynamics including H-L back transitions for rotating MHD and global RWM triggering events. Disruption warnings are issued with sufficient time before the disruption (on transport timescales) to potentially allow active profile control for disruption avoidance, active mode control, or mitigation.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0133825>

I. INTRODUCTION

Disruption (DIS) prediction and avoidance is a critical need for next-step tokamaks, such as ITER, since plasma disruptions^{1,2} can place significant thermal heat loads and electromagnetic forces on the device reducing the lifetime of components and can potentially lead to damage from runaway electrons.³ Meeting these challenging goals with the high reliability required for ITER and future tokamaks goes beyond active plasma instability control alone and will require multiple approaches, including an understanding of the connection between

events leading to disruptions, and the ability to forecast such events well before they occur. Studies often aim to predict the onset of a disruption with sufficient time to successfully trigger disruption mitigation systems. This criterion will mitigate potential damage to the device in question; however, it will also terminate the plasma and is expected to require significant time to reset the tokamak to a proper operational state to continue the operation. With sufficiently early forecasting of a potential disruption, appropriate control systems and actuators could be used to alter the plasma state in a way that would

entirely avoid disruption. The Disruption Event Characterization and Forecasting Code (DECAF),^{4,5} under development for this purpose, is used to automate the analysis of tokamak data to determine chains of events that lead to disruptions and to forecast their evolution to inform plasma profile and mode control systems aimed to most preferably avoid plasma disruption entirely, or if needed to mitigate the deleterious effects of a disruption. The DECAF approach also provides the physics understanding required to control the plasma to best avoid unfavorable plasma operational states.

The disruption chain “events” as defined in DECAF largely follow the paradigm established by the analysis performed by de Vries *et al.*⁶ for JET. These studies, in which the results were entirely produced by individual and manual examination of each discharge studied, established a framework for quantifying disruption events. The relative temporal combination of these events was considered in that work, which have the analog of disruption event “chains” in DECAF. The DECAF approach aims to automatically determine the relation of the events and quantify their appearance to characterize the most probable and deleterious event chains. It additionally aims to forecast the onset of the events and chains, especially for events that experimentally manifest in close time proximity to the disruption and would elude disruption avoidance control systems or even disruption mitigation systems. Earlier work by de Vries *et al.*⁷ using the JET database was also an important element of the evolution of disruption research, especially as it showed fairly low disruptivity in the carbon wall operation (dropping to near 5%), but did not address disruption chain events and their interconnection. This and the studies in Ref. 6 contrast the higher disruptivity in JET experienced with the ITER-like wall operation, illustrating the need for a more physics-based assessment to understand the causes of the disruptions more universally.⁸ Other studies have attempted an automatic disruption classification in a statistical manner utilizing generative topographic mapping applied to the JET operation with the ITER-like wall.⁹ Initial studies of real-time assessment allowing early reaction to potential disruption causes, including attempts at disruption avoidance, have been conducted on machines, such as ASDEX-U and TCV.¹⁰

The DECAF approach differs from these studies. It is primarily physics based and aims to provide a quantitative and, importantly, deterministic (rather than a statistical) predictor for disruptions. It also aims to provide an understanding of the dynamics of the events leading to disruptions to best ensure disruption prediction extrapolability to future devices. This is highly important in high fusion power devices, such as ITER in which the production of disruptions to teach purely automated model building approaches is highly restricted. Still, the DECAF approach and code are highly flexible and allow a large range of models from simple empirical comparisons to reduced explicit analytic models based on computationally intensive first-principles physics analysis, machine learning (ML) reductions of first-principles physics models, or hybrid machine learning models that use both physics-based and pure machine-based techniques. DECAF events are not simply labels. Instead, they contain both attributes and methods in an object-oriented programming sense. In this way, the collected understanding of these events can be programmed and continually validated against tokamak data to improve their general validity. To best validate the expanding models being added to DECAF, significant effort is being placed on testing the algorithms against *full* tokamak databases on multiple tokamak devices throughout the

world. As shown later, this approach is required to avoid errant determination of plasma parameters from databases limited to time periods that are only in close time proximity to the disruption. A larger variety of devices also provides essential depth in testing physics models and determining uniqueness and commonality in the events and their chains leading to disruptions. In the present work, the KSTAR, MAST, and NSTX databases are examined, with the analysis expanding to the NSTX-U, MAST-U, ASDEX-Upgrade, DIII-D, and TCV databases that are also available. Importantly, DECAF meets the disruption predictor requirements outlined by Humphreys *et al.*¹¹ that a predictor must (i) predict specific phenomena, (ii) provide a continuous variable quantifying proximity to disruptive states that can trigger actions, (iii) provide sufficient lead time for mitigation or avoidance, (iv) be extrapolable to new devices, and (v) be real-time calculable. To clarify, the goal of the research presented here is the automated characterization of disruptions by event chains and their forecasting using input data and analysis that are not computed in real time. Research conducted to support DECAF aims to use full physics models to first determine understanding of the behavior of the relevant disruption event. However, the implementation and development of the present DECAF code has been performed to efficiently allow analogous analysis to be performed in real-time.⁵ Model reduction and other techniques for this purpose (e.g., neural net representations) are discussed below for certain models shown. When used, reduced models are validated against full models. Future work continues to do this and also to compare the performance of the offline and real-time DECAF analysis.

The following sections examine an important subset of the event analysis in DECAF research and insights gained on the connection of plasma dynamics to the events. Section II describes the DECAF approach further, including the concept of disruption warning levels (with connection to past work). Section III describes the continuing development of the physical models in DECAF, with examples, including density limits and magnetohydrodynamic (MHD) instabilities that include both their automated characterization and present forecasting capabilities. Section IV describes the initial investigation of large, general databases (from KSTAR, NSTX, and MAST), including full disruption event chains, the importance of analyzing long periods of the plasma evolution in the devices, the increasing capability of the code to produce early disruption forecasting, and a brief summary of the disruption prediction model performance evolution. A summary and discussion are included in Sec. V.

II. DISRUPTION CHAIN EVENTS AND WARNING LEVELS

Figure 1 simply illustrates the paradigm that DECAF follows in providing automated understanding of the dynamics leading to a tokamak disruption along with an example from experiment. Continuous tokamak plasma operation at high fusion performance is desired [Fig. 1(a)]. However, at some point, this “normal” operational plasma state can be altered by various events ranging from purely technical considerations (e.g., magnetic field power supply interruption) to more complex reasons, such as the onset of plasma instabilities, loss of heating power balance, or loss of torque balance. This alteration is considered as a chain of individual events, starting with a trigger event and evolving toward the plasma disruption [here labeled by the acronym disruption (DIS) representing the plasma current quench]. The DECAF analysis of device databases aims to automatically determine not only

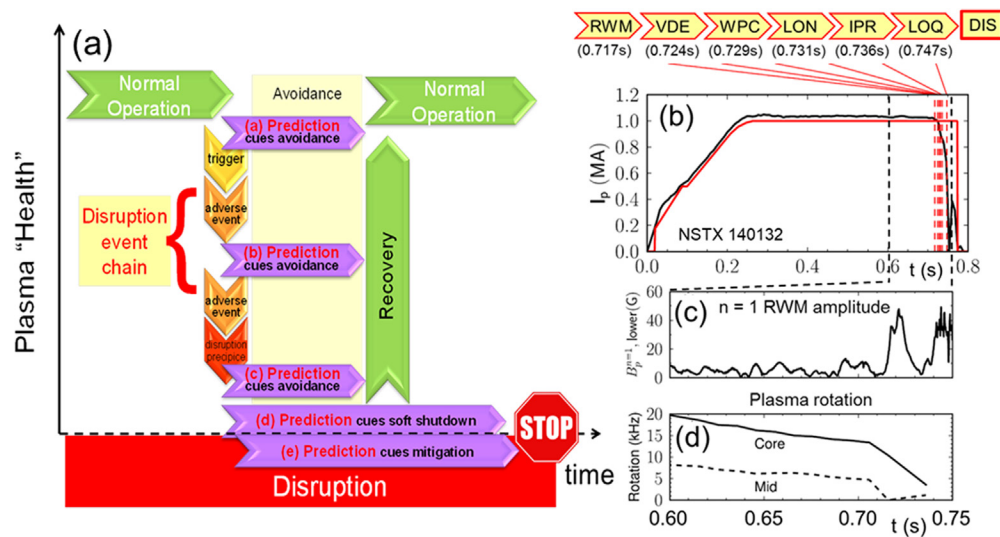


FIG. 1. (a) Schematic diagram illustrating a paradigm of plasma state evolution away from normal operation toward a plasma disruption as a series of events that form a disruption event chain, (b) automated evaluation of the disruption event chain for a plasma discharge with $\beta_N = 4.5$, and (c) and (d) higher time resolution illustration of $n = 1$ RWM amplitude and plasma toroidal rotation as the disruption is approached.

the frequency of occurrence of events but also the understanding of this chain of events. Future real-time implementations of DECAF diagnostic interpretation and forecasting models of such events can then be used to trigger disruption avoidance systems. This expanded strategy can be contrasted to present disruption avoidance systems, e.g., MHD instability control systems¹² that essentially wait until the “disruption precipice” to address avoidance of the oncoming disruption. Figure 1(b) shows a DECAF analysis of a plasma with “mid-range” normalized beta $\beta_N \equiv 10^8 \langle \beta_t \rangle a B_0 / I_p$ (where $\beta_t \equiv 2\mu_0 \langle p \rangle / B_0^2 = 4.5$ toroidal beta, p is the plasma pressure, B_0 is the vacuum toroidal magnetic field at the plasma geometric center, and a is the plasma minor radius at the midplane) in the NSTX¹³ device. In this case, a global magnetohydrodynamic instability (resistive wall mode, RWM) is identified by DECAF as the event chain trigger, which in the past was thought to be the direct precursor to the disruption (DIS). However, DECAF identifies several interceding events. Next in the chain, 7 ms after the RWM event is a vertical displacement event (VDE). The algorithm for determining whether a VDE event has reached a critical warning level is described in Sec. II. Five milliseconds later, a wall proximity warning (WPC) is issued, indicating that the plasma boundary is about to touch the device first wall (and does soon after the warning). A low plasma density warning (LON) is issued 2 ms later, followed 5 ms later by the plasma current request (IPR) event warning that the feedback control target plasma current request is no longer being met. By now, the original separatrix-limited plasma is in contact with the wall and is decreasing in size, with the edge plasma safety factor, q , decreasing as plasma poloidal flux is lost. At 9 ms later, a low q warning (LOQ) is issued. Finally the time of the disruption (DIS) is found, based on the plasma current quench.¹⁴ The DIS event occurs over 30 ms after the trigger RWM is issued, which is expected to be just enough time in ITER to trigger the disruption mitigation system effectively. However, this relatively short-duration disruption chain would be better handled if the RWM event itself was

forecast at an earlier time. DECAF presently has a model to do this, as discussed in Sec. III C 2. At present, DECAF event warning levels are determined by a flexible diagnostic and physics model “point” system similar to that successfully used for NSTX¹⁵ but significantly expanded including results from general analysis of the evolving diagnostic input and plasma equilibria with a continuous warning level determination. A key expansion of the present DECAF approach is that several event criteria can be used in conglomerate to determine combined “levels” that allow DECAF to issue event warnings. For example, at present, 15 separate criteria are used to determine the total MHD warning level for rotating MHD modes (see Sec. III B 1).

III. PHYSICAL MODEL DEVELOPMENT

A profound power of the DECAF approach is the ability to test any physical model developed by the fusion research community for practical use as part of a disruption prediction model ensemble. Models that can quantitatively forecast disruptions more accurately across all devices can then objectively be chosen as being more desirable using quantitative figures of merit. Over 50 disruption chain events are presently identified, with over 20 events that have diagnostic evaluation and physical models providing warning levels. Simpler evaluations examine key diagnostics in combination to compute warning levels, with comparison to critical levels to determine when DECAF issues event warnings. For example, the VDE event combines a comparison of axis position ($|Z|$), axis velocity ($|dZ/dt|$), and $Z dZ/dt$ against threshold levels set in the model. Critical levels of such models will differ for each machine. The validation of the technical and physics-based models for each of the five devices in the present DECAF database now comprises the primary near-term research effort. More desirable are models that more transparently reproduce the behavior of all tokamak devices. The simplest models in this class are empirical models, such as the Greenwald density limit (GDL). A

next level includes models that are more closely based on first-principles physics, examples of which are discussed below.

An important advantage of the physics-based approach of DECAF is that parameters that would logically need to change in models for each device are guided or specifically identified by theory. A simple example is the kink stability beta limit, which can simply be expressed in the form $C_1^* < I_p/aB_0$. Theory has shown that C_1 , often taken as a constant, can be computed using ideal MHD theory.¹⁶ Such an approach is taken when possible. Additionally, extended modeling has shown that this limit can be expanded to include plasma shape, pressure peaking, and plasma internal inductance (also be determined by ideal MHD theory).¹⁷ Where appropriate, those additional parameterizations are used in DECAF (for example, in the resistive wall mode growth rate determination). In this way, theory can be used to alter model parameters for different devices.

A. Density limits

The Greenwald density limit¹⁸ (event GWL) is included in DECAF as a universal empirical model for disruption forecasting. Recently several theories have been developed to explain the observed global Greenwald limit in tokamaks, including a ballooning stability limit at the separatrix¹⁹ and a local island power balance (IPB) theory.^{20,21} In the latter theory, power balance in an island between input Ohmic heating and radiated power results in a maximum local density that scales with local current density. If the density at the island exceeds the limit, or alternatively if the radiated power at the island exceeds the input power ($P_{loss} > P_{input}$), then the island grows and can lead to plasma disruption. The limit can be written either in a form which mimics the global Greenwald density limit in a local form or one that mimics a radiated power fraction localized to the magnetic island surface. This model has been added to the DECAF code, including the radiated power, resistivity, and current density profiles as inputs.

The radiated power profile (P_{loss}) can either be directly measured or can be estimated from density profiles and calculated cooling rates of deuterium and impurities, such as carbon, which depend on electron temperature.²² Figure 2 shows both the measured and calculated profiles for an NSTX discharge. The P_{loss} is calculated as $P_{loss} = n_e \sum n_Z L_Z$, where the species Z considered in this case are limited to deuterium and carbon and the cooling rates L in $W\ m^{-3}$ are given for




FIG. 2. Profiles of calculated (deuterium and carbon) and measured total radiated power density and calculated input power density for NSTX discharge 134 020 at 0.60 s.

deuterium by $L_D = 5.35 \times 10^{-37} T_e^{1/2}$ with electron temperature in keV (Ref. 17), and for carbon by tabulated formulas in Ref. 20. The input power profile that P_{loss} is compared to in Fig. 2 is calculated from $P_{input} = \eta j^2$. The resistivity profile, η , is calculated based on electron temperature and the effective charge Z_{eff} (formula from Ref. 17), which are measured by Thomson scattering and charge exchange recombination spectroscopy. The current density profile used here is the total surface-averaged current density profiles from various sources (Ohmic, bootstrap, beam-driven), which are computed by the TRANSP code.

The power balance model is a local condition for island growth; therefore, mode marginal stability would occur when $P_{loss}/P_{input} > 1$ at the location of the island. This defines the DECAF event “island power balance” (IPB) shown in Fig. 2. The rotating MHD mode growth that arises when $P_{loss}/P_{input} = 1$ (Fig. 3) is measured as having toroidal mode number $n = 1$. The lowest order rational surface in the plasma is $q = 2$, so $m/n = 2/1$ activity is the most likely candidate. Therefore, the local power balance criterion is evaluated at the $q = 2$ surface. The $n = 1$ mode onset in Fig. 3 is highly correlated with the IPB event warning in the plasma shown. Also shown is the computed Greenwald fraction evolution and the DECAF event GWL defined as the line-averaged plasma density equal to the Greenwald density. At the IPB event, the Greenwald fraction is ~ 0.9 . While this correlation is positive, the present state of analysis shows the quantitative evaluation of the IPB event to be sensitive to the accuracy of the local input criterion (e.g., position of the $q = 2$ surface). The present DECAF analysis shows

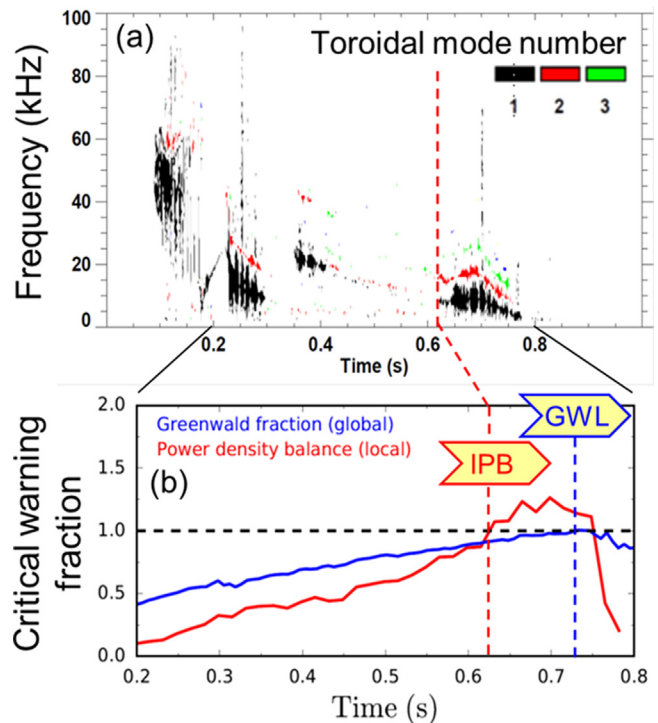


FIG. 3. (a) Spectrogram of rotating MHD activity from a toroidal array of magnetic probes for NSTX discharge 134 020, illustrating $n = 1$ mode growth near the time of the loss of power island power balance (DECAF event IPB). (b) Greenwald fraction and local power balance criterion.

that the local island power balance evolution follows the evolution of the global Greenwald fraction. For 13 discharges tested, the Greenwald fraction ranges from 0.75 to 1.05 at the time of MHD onset and the local island power balance fraction has a range of about 0.60 to 1.50. Continued analysis is focused on reducing this variation and eliminating the need for the full TRANSP analysis for each plasma, for example through neural net evaluation of a representative set of TRANSP runs to determine the required input for the IPB event. In addition, the Greenwald density limit itself has been shown not to be a very stringent limit in auxiliary heated plasmas, and the radiative power model in the DECAF IPB event has also not been proven to be stringent if the variation on the input parameters was reduced. Both of these criteria will be further investigated for larger data sets in future work to examine how they may be made more stringent.

B. Rotating MHD instabilities

An automated analysis of rotating MHD modes with tearing characteristics has started by using a DECAF module to produce physical event chains leading to disruptions through slowing of the modes by resonant field drag mechanisms and subsequent locking.²³ An algorithm portable across tokamaks devices has been developed that processes the spectral decomposition and signal phase matching of magnetic probe signals for mode discrimination. Multiple modes occurring simultaneously are tracked, and bifurcation of the toroidal rotation frequency and locking for each mode due to the loss of torque balance under resonant braking are detected.

1. Disruption event characterization

The information analyzed for these modes along with plasma rotation profile and other plasma measurements produces predictive warnings for the individual modes, along with a total MHD event

warning signal showing initial success as a disruption forecaster. These capabilities are illustrated in Fig. 4 for the same plasma shown in Fig. 3. In the plasma illustrated, rotating MHD instabilities thought to be nonlinearly saturated and slowly evolving resistive modes are found using a generalized phase matching algorithm in DECAF using an array of toroidal magnetic probes typically available in tokamaks. The code discriminates the toroidal mode number of the instabilities and tracks all modes that have greater than a specified amplitude. Modes that approach the disruption are indicated by the chevrons in the diagram (which show the mode n number). DECAF events based on the mode evolution are also shown, including the bifurcation of the modes (BIF-n1,2) (loss of torque balance leading to rapid loss of mode rotation), and events marking the locking of the modes in the laboratory frame of reference (LTM-n1,2). A single “total” MHD warning signal that varies with time is also shown. This warning is created by a set of criteria and can be used as a disruption predictor, as described in Sec. III B 2.

2. Forecasting

A significant part of DECAF research is determining the best criteria to create predictive warnings. The warning model shown in Fig. 4 is comprised of 15 separate criteria, also shown in the figure displayed as a heat map. The conditions fall into seven groups: (i) detected mode frequency very low to high (conditions 2 and 3), (ii) mode frequency torque balance condition (condition 4), (iii) mode strength (condition 5), (iv) decreasing β_N and less safe β_N/l_i level (conditions 6 and 7), (v) decreasing plasma rotation at three radial positions (conditions 8–10), plasma rotation too low at three radial positions (conditions 11–13), (vi) sufficiently strong locked mode signal (condition 14), and (vii) l_i below the (zero-beta) current drive kink limit (condition 15). This summary illustration of the criteria used to produce the warning level provides a useful display of how the total warning level reaches high

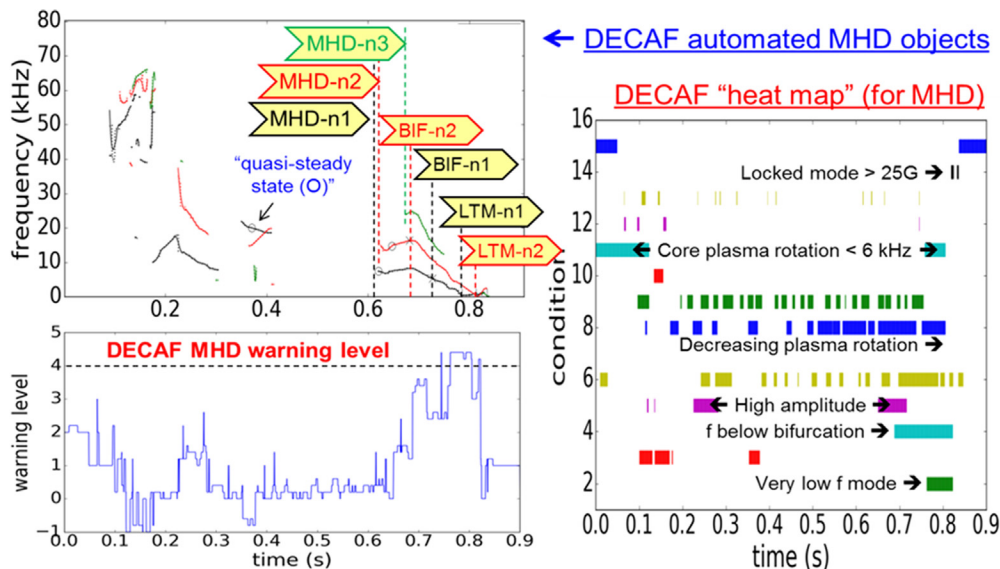


FIG. 4. Rotating MHD mode discrimination capabilities in DECAF. The upper left frame shows the mode discrimination and decomposition into DECAF events. The lower left panel shows a total MHD warning level that increases as the disruption is approached. The right panel shows a heat map illustration of 15 event criteria that comprise the total MHD warning level (colors label different conditions).

values, indicating a disruption onset. A total warning level of 4 indicates close proximity to the disruption for this model. At present, the criteria and threshold levels used in the model are based on experimental experience in determining the evolution of parameters, such as mode frequency and plasma rotation profile evolution, to produce the warning level evolution, which is then compared to the success of binary classification of the plasma as having or not having a disruption. Future work will aim to numerically automate this process while constraints are set for the plasma parameters considered in the model. The heat map also gives us an understanding of what is happening in the plasma to create the undesirable plasma states approaching the disruption. Early in the discharge (before $t = 0.25$ s), MHD modes are also found, and core plasma rotation is low as the plasma starts up and typically transitions from counter-neutral beam injection (NBI) rotation to co-NBI rotation. However, the mode frequencies are relatively high at this time, which is generally a safe condition. Later, near $t = 0.25$ s, the MHD warning level increases as modes are again found but now with decreased and decreasing rotation frequency. However, these frequencies are not critically low (no mode bifurcations are found) and plasma rotation is not low, so the warning level remains low. However, after $t \sim 0.6$ s, the heat map clearly shows more negative (destabilizing) criteria occurring simultaneously including an increased mode amplitude, decreasing mode frequencies, and decreasing plasma rotation across the profile. Near $t \sim 0.7$ s, more negative criteria occur: mode frequencies are now below a recent evaluation of mode bifurcation frequency levels, the modes drop to very low frequency, and core plasma rotation (three channels of the plasma rotation profile are considered—core, middle, and outer) is critically low. Late in the evolution in close time proximity to the disruption ($t \sim 0.8$ s), a critical level of locked mode amplitude occurs. Such a locked mode detector signal is typically used as the primary, and occasionally the only indicator to predict a possible disruption, but this indication occurs very late in the evolution. We see here that the DECAF analysis starts to show a significant change in the total MHD warning level about 180 ms earlier, providing far better advanced notice of the potential disruption allowing the potential for control systems to alter plasma stability to avoid disruption. Additionally, and of critical importance, the DECAF analysis provides physical understanding of the negative evolution of the plasma state as it moves toward the disruption. Further forecasting of resistive MHD stability using the resistive DCON code is being investigated through supporting KSTAR research.^{24,25}

C. Global MHD instabilities

Global MHD instabilities, such as external kink/ballooning modes or resistive wall modes (RWMs),²⁶ can cause the most rapid disruptions (e.g., Fig. 1) and give the least amount of pre-disruption warning time. Therefore, attention needs to be put toward forecasting such events to cue profile control systems well before instability develops.

1. Disruption event characterization

To examine disruption event chains with global MHD triggers, the DECAF analysis was performed on a database of 45 NSTX discharges that were pre-determined to have unstable RWMs that lead to disruptions. Tearing modes were stable during these discharges to focus on global MHD in this analysis. A typical disruption event chain

with an RWM trigger is shown in Fig. 1. In this database, the RWM, loss of boundary control (WPC), LOQ, IPR, and DIS events are found in 100% of the plasmas, and VDE events are found in 91% of the plasmas. A pressure peaking warning (PRP) occurred on a majority of the discharges analyzed (35 of 45), but typically occurred with or after the RWM, not before in this database. The GWL event warning is found in a few cases when the warning level is set at a Greenwald fraction of 0.9. Interestingly, GWL can start the RWM disruption chain and is explained physically by the correlation of reduced plasma rotation caused by increasing plasma density, leading to RWM instability by a destabilizing change in the plasma rotation profile, discovered in NSTX.²⁷ Analysis shows that 61% of RWM events in a shot occur within 20 conducting wall current diffusion times, τ_w , of the disruption. The other RWM events found occur earlier but are not false positives as they cause significant thermal collapses or “minor disruptions” of the plasma with subsequent recovery (plasma stored energy can drop by 30% or more over tens of ms, much larger than the largest edge-localized modes (ELMs) in tokamaks that cause far smaller stored energy decreases up to $\sim 6\%$). It is useful to examine which events are commonly associated with this dynamic. For example, the low safety factor condition (LOQ) was clearly seen to happen often in close conjunction (either just before or just after) the designated time of disruption (DIS). One way of seeing this is to examine a histogram of some of the timing of the events before the time of disruption (DIS), shown in Fig. 5. Here, only the events within $14 \tau_w$ of the disruption are shown, where τ_w is taken to be 5 ms; there are some RWM events at earlier times which are not shown here. It is clear that LOQ and IPR events occur close to the time of disruption, and these are often preceded by VDE and RWM events, which peak around 30 ms prior to the disruption. Associations between DECAF events are powerful conclusions of the analysis shown here, and it is expected that a greater number of DECAF events associations will be discovered in the future analysis of much larger databases.

Examining the common chains of events more closely can provide insight into how to cue avoidance systems to return to normal plasma operations (see Fig. 1). For example, if the RWM can be detected in real-time by a growing mode amplitude signal from an array of external magnetic sensors, it is useful to know what the typical

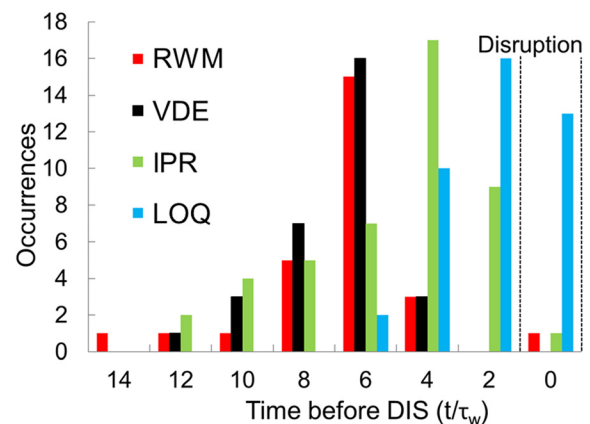


FIG. 5. Histogram of the timing of RWM event triggered disruption chain events before the time of disruption.

routes of plasma behavior directly follow the RWM so that plasma control systems may be employed to avoid them. Even with the limited dataset examined here, we can find interesting trends. Of the 26 RWM events occurring within 100 ms of the disruption, they were followed immediately by VDE (related to bulk plasma motion) 15 times. Furthermore, looking at the two-event chains that happened directly after this set of RWMs, we find that even though there are theoretically 56 two-event combinations that could occur from the eight currently tested, just two two-event chains accounted for 50% of the cases (VDE → PRP and VDE → WPC) and five accounted for 77%.

2. Forecasting

Kinetic RWM analysis has shown high success over years of quantitative comparison to experiment to determine the mode marginal stability allowed through plasma precession drift and bounce orbit resonances, collisionality, ν , and energetic particle effects.^{25,28} To allow rapid processing, full kinetic RWM computations (using the MISK code^{25,26}) that have been used successfully to predict mode stability on NSTX and DIII-D have been used to create a reduced model of the kinetic RWM stability growth rate in DECAF [Fig. 6(a)]. This allows the plasma stability to be forecasted well before the mode ever appears. Gaussian functions with parameters fit from full MISK calculations of NSTX marginally stable equilibria are used to define the kinetic energy functional δW_K as functions of $E \times B$ frequency and collisionality. The model also incorporates expressions dependent on plasma pressure peaking, internal inductance, and aspect ratio for the ideal MHD no-wall and with-wall beta limits computed from thousands of DCON calculations using experimental equilibria.²⁹ The modeled growth rate can be used to forecast RWM instability based on plasma equilibrium reconstructions and rotation measurements and is time-dependent based on the equilibrium evolution. Figure 6(a) illustrates the evolution of a high normalized beta NSTX experimental plasma as it becomes RWM unstable near a predicted marginal stability contour [while not shown, the growth rate contours on Fig. 6(a) change as the plasma evolves]. This reduced kinetic RWM stability model in DECAF, detailed in Ref. 27 performed well in its first incarnation against a larger database of plasmas to determine the proximity of discharges to marginal stability [Fig. 6(b)]. The model correctly

predicted that 16% of the plasmas were stable. The model predicted instability 84% of the time (stringent marginal stability evaluation) for experimentally unstable cases with a relatively low false positive rate (7% of plasmas predicted to have disrupted were actually stable plasmas). DECAF also showed 44% of plasmas were predicted unstable within 320 ms ($\sim 60 \tau_w$) of the disruption time, and 33% were predicted unstable within 100 ms of a minor disruption. Stability was predicted in 77% of experimentally stable cases. The evolution of discharges that were RWM stable was notably separate on the (ExB frequency, collisionality) stability map, not crossing the computed marginal stability contour. An additional positive aspect of the reduced kinetic MHD RWM growth rate model is that it is real-time calculable. This contrasts the full model, which takes several minutes to compute each point in rotation and collisionality space.

IV. INITIAL INVESTIGATION OF GENERAL DATABASES

A. Individual disruption chain events

The DECAF code has produced an initial analysis of large databases for multiple tokamak devices for a small set of disruption characterization events. The analysis is conducted over the full duration of the planned plasma current flat-top, rather than a limited period near the disruption time as might be available from a disruption database. Thousands of shot seconds are available in the databases, with upward of $0.5\text{--}1 \times 10^6$ tested sample times per database. For example, if the DIS event is used, the analysis produces the equivalent of “disruptivity diagrams” showing the probability of a disruption occurring within a given parameter space of tokamak operation. These diagrams are shown for NSTX, MAST, and KSTAR in Fig. 7 expressed as standard stability operational space (l_i, β_N) figures (l_i is the plasma internal inductance). This multi-device comparison illustrates a highly important and still largely unappreciated result separately published for DIII-D and NSTX³⁰ for smaller datasets that plasma disruptivity does not increase (and can actually decrease) as β_N increases. This is not due to lower operation probability in this space and can be explained in NSTX by increased kinetic MHD RWM stabilization in this regime.^{27,28} However, as will be shown in Sec. IV B, the high beta regions of low disruptivity are in fact key areas for DECAF algorithms to analyze events that can lead to disruptions.

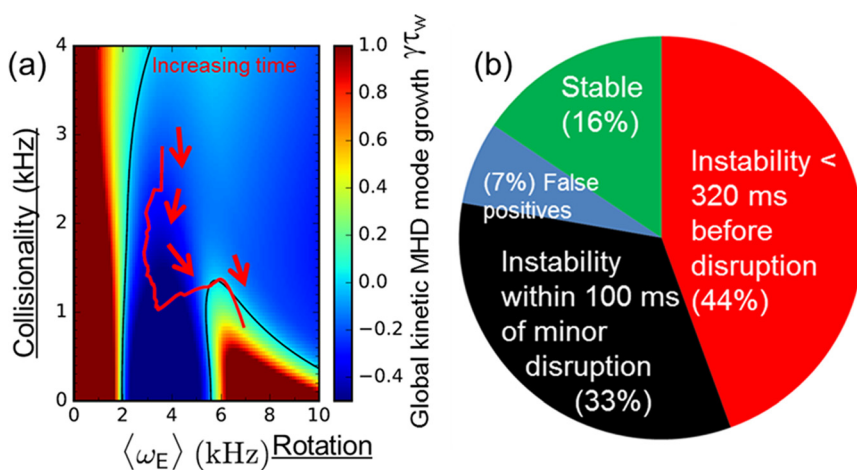


FIG. 6. (a) Stability map vs ExB frequency and collisionality from DECAF reduced kinetic RWM stability model; (b) statistics illustrating results of the model in forecasting instability for RWM unstable NSTX plasmas.




FIG. 7. Event probability diagram of DECAF event DIS during I_p flat-top in large databases from the (a) NSTX, (b) MAST, and (c) KSTAR tokamaks showing that disruption probability does not have to increase as plasma normalized beta is increased.

Unlike standard disruptivity plots, DECAF can provide critical insight by illustrating where in parameter space events other than DIS happen. For example, the VDE event detects the loss of vertical stability. When plotted in the parameter space of elongation, κ , vs l_i , it becomes clear that vertical stability shows a strong dependence on these parameters [Fig. 8(a)] and that, additionally, the location in parameter space of an event preceding the disruption (like VDE) can be far from where the actual disruption event occurs [DIS, shown in Fig. 8(b)]. Therefore, the often-used analysis of a disruptivity plot, such as Fig. 8(b) is shown to be misleading. The correct figure to use to understand disruptions due to VDE is Fig. 8(a).

B. Disruption event chain analysis for arbitrary discharges

DECAF event characterization and event chain analysis show that the disruption forecasting analysis often starts during plasma states that can appear safe. This is illustrated using the disruptivity

database plot shown in Fig. 7(a) and the figures in this section. The regions of high disruptivity in Fig. 7(a) may be thought to be the most important based on human inspection. However, an apparent problem is that the region of high disruptivity at low β_N and mid-range l_i is not physically understood to be a dangerous operational region. The enigma is resolved by understanding that the plasma state can significantly evolve from more usual high performance parameters to the point at which the disruption actually occurs. This fact is completely missed, for example, by disruption database studies that only process data near the disruption time. Even worse, such studies may parameterize disruptive limits based on these terminal states. In contrast, DECAF disruption event chain analysis of two discharges in Fig. 7(a) that disrupts (DIS event in DECAF, marked by red and green X's in the figure) shows that the *start* of the event chains appears in the region indicated by the red and green circles—which are *far* from what might be expected. This also illustrates why the use of a numerical tool that would focus only on the regions of high disruptivity (such as a “black box” machine learning approach)




FIG. 8. Event probability diagrams of DECAF events VDE (a) and DIS (b) for a large database from the NSTX tokamak.

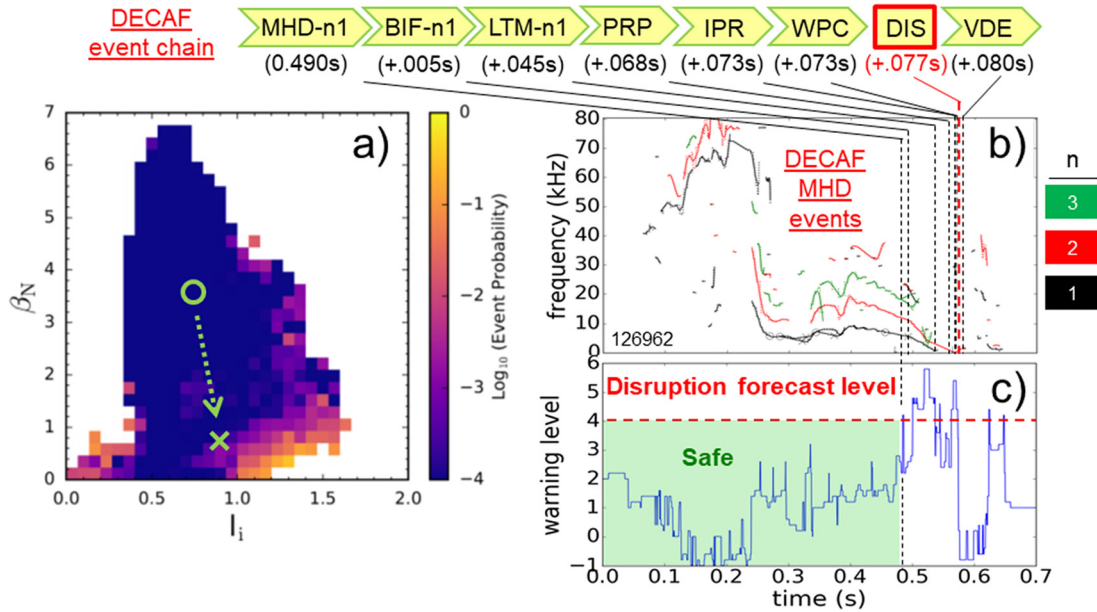


FIG. 9. (a) Evolution of plasma state from a low disruptivity region toward disruption due to a rotating MHD mode locking dynamic in NSTX (circle: time of DECAF forecasted trigger event; X: state at time of disruption), (b) DECAF decomposition of rotating MHD in relatively slow evolutions toward disruption, (c) computed DECAF rotating MHD warning signal. The DECAF event chain leading to disruption is shown on the top of the figure.

would produce an inaccurate assessment of the plasma states that produce early disruption forecasts. The disruption event chains for these plasmas in Fig. 7(a) are shown in Figs. 9 and 10 along with the DECAF MHD mode decomposition and total MHD warning level.

As before, we see this warning level rising toward and past the critical value of 4.0 as the disruption is approached. The DECAF mode decomposition adds information, showing that the mode evolution toward lower rotation frequencies is relatively slow in Fig. 9. This is

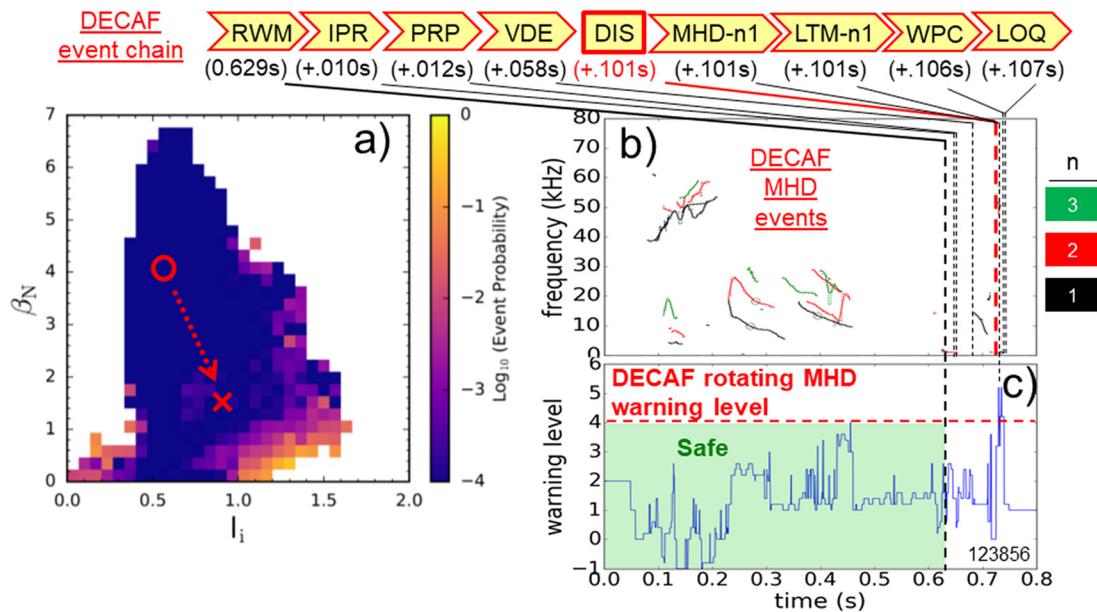


FIG. 10. (a) Evolution of plasma state from a low disruptivity region toward disruption due to an unstable RWM in NSTX (circle: time of DECAF forecasted trigger event; X: state at time of disruption), (b) DECAF decomposition of rotating MHD modes, (c) computed DECAF rotating MHD warning signal. The DECAF event chain leading to disruption is shown on the top of the figure.

one reason why the plasma disrupts far from the plasma state at the trigger event.

The DECAF event chain in Fig. 9 provides a wealth of information. In Fig. 9, we see a critical warning for the individual $n = 1$ rotating MHD mode (MHD-n1) as a starting point for the chain. Note from the top frame that a low frequency $n = 1$ mode was detected far earlier—near $t \sim 0.22$ s. However, the warning level for the activity was not determined to be sufficiently high then to warn of a disruption. The mode bifurcation (event BIF-n1) occurs 5 ms later. The mode locks (event LTM-n1) 45 ms after the bifurcation. Then, a different dynamic occurs, as DECAF finds a pressure peaking event warning (PRP) happening 23 ms after the mode lock. While the warning literally flags that the pressure peaking factor is exceedingly high, it also importantly indicates that an H–L energy confinement back-transition has occurred, the H-mode pedestal is lost, and the neutral beams have better penetration increasing the plasma pressure peakedness. The IPR warning occurs 5 ms after PRP, and simultaneously, the plasma makes a close approach to the vessel wall (WPC). Finally, the plasma disrupts 4 ms after the WPC event. It is also interesting that the VDE event warning occurs 3 ms after DIS. Usually the events are reversed in time. This indicates that the plasma mainly remains on the midplane during the evolution, uncharacteristic of NSTX disruptions. Figure 9 shows a relatively slow RWM-triggered disruption [i.e., compared to Fig. 1(b)]. In this disruption event chain, the PRP warning again indicates an H–L back transition and a VDE is produced approximately $10 \tau_w$ after the RWM trigger occurs. As shown, the disruptions in these two plasmas occur 77 and 101 ms after the initial DECAF warnings. These intervals represent transport timescales (a few energy confinement times) and so would allow sufficient time for active profile control for disruption avoidance, or easily allow time for active mode control or disruption mitigation.

The DECAF event chain in Fig. 10 is quite a bit different than shown in Fig. 9, which is understood by the physical evolution of the MHD mode activity. The DECAF decomposed MHD mode evolution shows clear rotating mode activity through $t = 0.45$ s, and the corresponding DECAF warning level increases toward the critical value. However, before reaching that level, the rotating MHD signatures vanish, and the DECAF warning level drops. This happens since in the experiment, the plasma evolution was changed to avoid rotating MHD later in the plasma discharge in order to study RWMs, which do not become unstable if rotating MHD modes of sufficient amplitude are present. As shown in Fig. 10, RWM is the eventual trigger event of the DECAF disruption event chain, and not the rotating MHD warning level. Another interesting observation of that the DECAF RWM event reaching the critical level shown will usually cause a plasma disruption within several RWM growth times $< 5\tau_w \sim 45$ ms. The longer time period between the RWM trigger event and the DIS event is understood when examining the DECAF event chain. The IPR event occurs 10 ms after RWM, indicating that plasma is under significant distress. Then, a PRP event occurs 12 ms after the RWM event, indicating that the plasma made a backtransition from H-mode to L-mode confinement that will cause a drop in β_N . That effect, counterbalanced with the rise in the plasma pressure peaking can lead to a longer growth time of the RWM amplitude, which in this case leads to a somewhat longer than usual disruption onset with VDE occurring 58 ms after RWM, and DIS happening 101 ms later.

C. Disruption prediction performance

Disruption prediction research using the DECAF approach also importantly allows *quantifiable* figures of merit (most importantly the plasma disruptivity) to assess any prediction models produced. This figure of merit allows an objective assessment of the relative performance of different models and allows an assessment of how close the predictor would come compared to ITER needs. Figure 11 shows a progression of DECAF disruption forecasting models. The earliest models included about ten events and were run on databases for which the events that led to the disruption were known. For such databases, DECAF produced very high performance (e.g., 100% true positives). A next evaluation of models focused on earlier forecasting once the first forecasting model was implemented in the DECAF code. True positives were found to be $\sim 84\%$, which was a measure mainly of the single forecasting model. The addition of more forecasting models (such as the MHD analysis shown) could improve that performance with further development. The more recent testing of the code has been on large databases $\sim 10\,000$ shot \times seconds of plasma run time tested. This was done with a smaller number of events due to computer RAM limitations. With five events, applied to all plasma shots from an NSTX database, DECAF has produced performance levels of over 91% true positive disruption predictions. False positives in this analysis reached 8.7%, which is fairly high. However, further code development that allows the events to poll each other will reduce this level considerably. For example on the NSTX device, a large portion of the false positive rate are due to the RWM event. Experimental experience shows that if a sufficiently large rotating MHD mode is present in the plasma, then an RWM will not become unstable. Thus, present code development includes the ability of the RWM to poll the MHD event for the existence of such mode activity. The addition of that criterion will significantly reduce the false positive percentage. Furthermore, restrictions on the analysis scope due to computer hardware are presently being largely eliminated through software improvements and access to more powerful computer clusters.

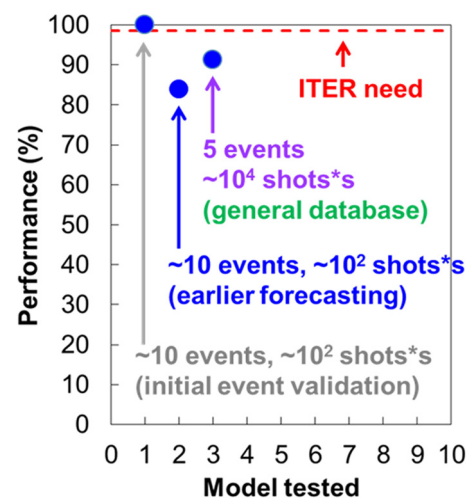


FIG. 11. DECAF model performance evolution (true positive disruption forecast).

V. SUMMARY AND DISCUSSION

The Disruption Event Characterization and Forecasting Code is a physics-based, fully automated analysis paradigm that continues to compile the knowledge of many years of tokamak research by implementing a collection of models to solve the critical issue of plasma disruptions in tokamaks. The approach considers the evolution from normal tokamak operation toward a plasma disruption as occurring through a linked series, or chain of physical or technical events, most or all being “off-normal.” The approach has several advantages, including an analysis that produces greater understanding of the events as a collective and allowing control systems to be guided by the events and event chains in taking early actions (e.g., on transport timescales) to avoid disruptions before they happen via the forecasting element of the analysis. The DECAF approach is inherently deterministic rather than statistical. This is important, as the performance of approaches that use a set of statistical elements, rather than deterministic events, will produce a warning level equal to the product of each of the individual warning elements. For example, ten elements that each have a 99% probability of success would have a total performance level of $(0.99)^{10} = 90.4\%$, while ITER is seeking a 98% true positive rate. In contrast, the deterministic nature of DECAF events and the critical event chains will predict a plasma to have a disruption even if all but one event in the chain fails to be correctly assessed at the critical level. In this way, DECAF events and the critical event chains are far better at reaching high true positive disruption prediction values. Attention does need to be placed on keeping false positive levels low. The DECAF approach meets the disruption predictor requirements outlined by Humphreys *et al.*¹¹ Four of the five requirements have been demonstrated in this paper. The DECAF events comprise both attributes and methods, which include the ability to predict the specific plasma physics phenomena covered by the scope of each event. The events in combination provide a continuous variable quantifying the proximity to disruptive states. The events themselves and more powerfully the event chains can be used to more intelligently trigger actions by control systems (e.g., by using gradients of the underlying physical processes to determine how to change the plasma state) than a system that only provides a binary classification of the plasma state. The DECAF event chains produce event relationships that can inform what control algorithms to use, as well as what actuators should be used. Also, relationships established by the event chains can be further analyzed by modern graph theory software constructs to produce understanding of the resulting DECAF analysis appropriate for automated, rapid computer responses when the chains appear in real time. The present model and forecasting development has produced analysis that forecasts the onset of a disruption on plasma transport timescales, thereby allowing sufficient lead time to avoid plasma disruption through plasma control techniques (e.g., profile control, shape control, plasma stored energy control, and plasma mode control). The DECAF paradigm puts high priority on the ability to characterize and forecast plasma disruptions from the full databases of multiple tokamaks to allow more stringent validation of the underlying physics models across several devices, thereby producing superior extrapolability to future devices.

While DECAF connects to the databases of several tokamaks, once read into the code, the data used are abstracted from the specific local data names to general names [e.g., charge exchange spectroscopy (CES) toroidal plasma rotation profile on KSTAR³¹ and charge

exchange recombination spectroscopy (CHERS) plasma rotation profile on NSTX³² becomes “plasma rotation” in DECAF], with a generalized input describing the measurement geometries. This allows straightforward mapping between quantities and consistency with data management systems, such as the integrated modeling & analysis suite (IMAS) for ITER.³³ Note that it is a common practice to create a full DECAF model (e.g., all events) for a new device, for example, MAST-U, by starting from a well-established model aimed at a similar device, for example, NSTX. Note that in such a porting of the DECAF model, the underlying physics models made need to have parameters change based on theory (e.g., ideal MHD stability limits). The changes are guided by theoretical understanding of the models. Future studies will include a systematic evaluation of superior DECAF models that have high accuracy in disruption prediction and forecasting across all plasma data from all available databases. Reaching high prediction accuracy levels required for ITER (~98%) and future devices is an ultimate goal of DECAF research.

Modern machine learning (ML) techniques were envisioned to be used in specific ways in the DECAF analysis since its origin once they were found to be needed. The first instance of this was thought to be the reproduction of the mode growth rate results for the kinetic MHD global mode model detailed in Ref. 27 by a deep learning neural net approach.³⁴ However, as shown in the reference, basis functions that produced good fits to the analysis results were found based on the underlying physics solutions and so that approach was utilized. Having analytic forms for the models are superior to purely numerical approaches in that gradients of the key quantities used for forecasting and control algorithms can be computed analytically, producing smooth results to reduce analysis output noise. However, finding such basis functions is not straightforward and so is not possible in general. Therefore, we have now started to implement machine learning for DECAF,^{35,36} but in significantly different ways than have been used for disruption prediction research to date, which typically approach the problem treating a disruption database as a binary classification problem, and applying deep learning or other techniques to the database as a black box. Machine learning in DECAF follows a philosophy that is more amenable to produce human understanding of the results and allow greater flexibility for use in control systems. Specifically, we are presently adopting three approaches of using machine learning to support DECAF. First is the reduction of results from certain complex physical models by deep learning neural nets to allow rapid (including real time) determination of quantities used in DECAF models. Two such machine learning techniques following this approach above have already been applied as analysis supporting the DECAF code, specifically, deep learning neural nets, and non-linear random forest regression analysis. These were used to train on DCON ideal MHD stability code calculations.³³ Second, a “hybrid model” approach is envisioned to comprise the bulk of ML use in DECAF. In this approach, both a physics model and a database that the physics model is meant to reproduce are given. The difference between the computed physical quantities produced by the model and the data, δ_e , is typically interpreted as the error between these two, which represents the part of the physical model that is not known. The hybrid model approach focusses on reproducing δ_e , rather than the entire physical model. This approach has two advantages: (i) the number of independent variables needed to produce the ML model of δ_e will be a smaller set than required to reproduce the entire database using ML alone, and (ii) the

ML techniques applied to δ_e have a greater probability of exposing key physics variables needed to reproduce δ_e because of this reduction of independent variables. This aids artificial intelligence (AI) to more easily produce human understanding of the remaining part of the physical model. A simple example: a database of the position of spheres dropping in a gravitational field with an atmosphere, and as a model a set of kinematic equations in a gravitational field, but in a vacuum. In this case, the δ_e would represent the key missing physics component: air resistance. It would not be difficult for the hybrid AI approach to determine that δ_e is a function of the square of the sphere's velocity. The new physics discovered that reproduces δ_e can be added to the original physical model, with the opportunity to iterate this process to discover more of the underlying physical model. The approach of addressing δ_e is analogous to the use of an optimum observer in control system design.³⁷ Finally, ML techniques will be used to evaluate more general linkages between DECAF events.

There is one final disruption predictor requirement remaining to be demonstrated by DECAF—that it be real-time calculable. This characteristic can also be met by all models discussed in this paper and that are presently available in DECAF. For models that are not easily reproduced using analytic basis functions, results from a full physics model can be functionally represented to allow real-time computation through ML techniques, for example, Deep Learning. This neural network technique is equivalent to fitting analysis results as a function of the key independent variables without assuming analytic basis function models. The present offline modeling and analysis is starting to be implemented for real-time use in the KSTAR device. In such an implementation, only causal analysis techniques based on real-time measurements and real-time equilibrium reconstruction quantities can be used. In fact, a recent set of experiments on the KSTAR device have been run with an initial implementation of DECAF installed in the plasma control system. Comparison of disruption prediction performance results between the two approaches will be made once the analysis of the initial real-time results is completed.

ACKNOWLEDGMENTS

This research was supported by U.S. DOE Grant Nos. DE-SC0020415, DE-SC0016614, DE-SC0018623, and DE-FG02-99ER54524, KFE Grant No. EN2301-14, and part-funded by the EPSRC Energy Programme Grant No. EP/W006839/1.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

S. A. Sabbagh: Writing – original draft (equal). **Jonathan Hollocombe:** Resources (supporting). **Jeongwon Lee:** Resources (supporting). **Jayhyun Kim:** Resources (supporting). **Andrew Kirk:** Resources (supporting). **Jinseok Ko:** Resources (supporting). **Won-Ha Ko:** Resources (supporting). **Lucy Kogan:** Resources (supporting). **Benoit P. LeBlanc:** Resources (supporting). **Jong-ha Lee:** Resources (supporting). **Andrew J. Thornton:** Resources (supporting). **John William Berkery:** Investigation (supporting); Writing – review & editing (supporting). **Siwoo Yoon:** Resources (supporting). **Young-Seok**

Park: Investigation (supporting); Writing – review & editing (supporting). **Jalal Butt:** Investigation (supporting); Writing – review & editing (supporting). **Juan Riquezes:** Investigation (supporting); Writing – review & editing (supporting). **Jun-Gyo Bak:** Resources (supporting). **Ronald E. Bell:** Resources (supporting). **Luis F. Delgado-Aparicio:** Resources (supporting). **Christopher Ham:** Resources (supporting); Writing – review & editing (supporting).

DATA AVAILABILITY

The data that support the findings of this study are openly available in disruption event characterization and forecasting in tokamaks at <http://arks.princeton.edu/ark:/88435/dsp018p58pg29j>, Ref. 38.

REFERENCES

- T. C. Hender, J. C. Wesley, J. M. Bialek, A. Bondeson, A. H. Boozer, R. J. Buttery, A. Garofalo, T. P. Goodman, R. S. Granetz, Y. Gribov *et al.*, *Nucl. Fusion* **47**, S128 (2007).
- N. W. Eidietis, S. P. Gerhardt, R. S. Granetz, Y. Kawano, M. Lehnen, J. B. Lister, G. Pautasso, V. Riccardo, R. L. Tanna, A. J. Thornton, and ITPA Disruption Database Participants, *Nucl. Fusion* **55**, 063030 (2015).
- M. Sugihara, S. Putvinski, D. Campbell, S. Carpentier-Chouchana, F. Escourbiac, S. Gerasimov, Y. Gribov, T. Hender, T. Hirai, K. Ioki *et al.*, in Proceedings of the 24th International Conference on Fusion Energy, San Diego, CA, 2012.
- S. A. Sabbagh, J. W. Berkery, Y. S. Park, J. H. Ahn, Y. Jiang, J. D. Riquezes, R. E. Bell, M. D. Boyer, B. P. LeBlanc, C. E. Myers *et al.*, in Proceedings of the 27th International Conference on Fusion Energy, Gandhinagar, India, 2018.
- S. A. Sabbagh, J. W. Berkery, Y. S. Park, J. H. Ahn, Y. Jiang, J. D. Riquezes, J. Butt, J. Bialek, J. G. Bak, M. J. Choi *et al.*, in Proceedings of 28th International Conference on Fusion Energy (IAEA, 2021), Paper No. IAEA-CN-286/1025.
- P. C. de Vries, M. F. Johnson, B. Alper, P. Buratti, T. C. Hender, H. R. Koslowski, V. Riccardo, and JET-EFDA Contributors, *Nucl. Fusion* **51**, 053018 (2011).
- P. C. de Vries, M. F. Johnson, I. Segui, and JET EFDA Contributors, *Nucl. Fusion* **49**, 055011 (2009).
- P. C. de Vries, M. Baruzzo, G. M. D. Hogewij, S. Jachmich, E. Joffrin, P. J. Lomas, G. F. Matthews, A. Murari, I. Nunes, T. Pütterich, C. Reux, J. Vega, and JET-EFDA Contributors, *Phys. Plasmas* **21**, 056101 (2014).
- B. Cannas, P. C. de Vries, A. Fanni, A. Murari, A. Pau, G. Sias, and JET Contributors, *Plasma Phys. Controlled Fusion* **57**, 125003 (2015).
- M. Maraschek, A. Gude, V. Igochine, H. Zohm, E. Alessi, M. Bernert, C. Cianfarani, S. Coda, B. Duval, B. Esposito *et al.*, *Plasma Phys. Controlled Fusion* **60**, 014047 (2018).
- D. Humphreys, G. Ambrosino, P. de Vries, F. Felici, S. H. Kim, G. Jackson, A. Kallenbach, E. Kolemen, J. Lister, D. Moreau *et al.*, *Phys. Plasmas* **22**, 021806 (2015).
- S. A. Sabbagh, R. E. Bell, J. E. Menard, D. A. Gates, A. C. Sontag, J. M. Bialek, B. P. LeBlanc, F. M. Levinton, K. Tritz, and H. Yuh, *Phys. Rev. Lett.* **97**, 045004 (2006).
- M. Ono, S. M. Kaye, Y.-K. M. Peng, G. Barnes, W. Blanchard, M. D. Carter, J. Chrzanowski, L. Dudek, R. Ewig, D. Gates *et al.*, *Nucl. Fusion* **40**, 557 (2000).
- C. E. Myers, N. W. Eidietis, S. N. Gerasimov, S. P. Gerhardt, R. S. Granetz, T. C. Hender, G. Pautasso, and JET Contributors, *Nucl. Fusion* **58**, 016050 (2018).
- S. P. Gerhardt, D. S. Darrow, R. E. Bell, B. P. LeBlanc, J. E. Menard, D. Mueller, A. L. Roquemore, S. A. Sabbagh, H. Yuh *et al.*, *Nucl. Fusion* **53**, 063021 (2013).
- F. Troyon and R. Gruber, *Phys. Lett.* **110**, 29 (1985).
- J. W. Berkery, S. A. Sabbagh, R. E. Bell, S. P. Gerhardt, B. P. LeBlanc, and J. E. Menard, *Nucl. Fusion* **55**, 123007 (2015).
- M. Greenwald, J. L. Terry, S. M. Wolfe, S. Ejima, M. G. Bell, S. M. Kaye, and G. H. Neilson, *Nucl. Fusion* **28**, 2199 (1988).

- ¹⁹T. Eich, R. J. Goldston, A. Kallenbach, B. Sieglin, and H. J. Sun, ASDEX Upgrade Team, and JET Contributors, *Nucl. Fusion* **58**, 034001 (2018).
- ²⁰D. A. Gates and L. Delgado-Aparicio, *Phys. Rev. Lett.* **108**, 165004 (2012).
- ²¹Q. Teng, D. P. Brennan, L. Delgado-Aparicio, D. A. Gates, J. Swerdlow, and R. B. White, *Nucl. Fusion* **56**, 106001 (2016).
- ²²D. E. Post, R. V. Jensen, C. B. Tarter, W. H. Grasberger, and W. E. Lokke, *At. Data Nucl. Data Tables* **20**, 397 (1977).
- ²³R. Fitzpatrick, *Nucl. Fusion* **33**, 1049 (1993).
- ²⁴Y. S. Park, S. A. Sabbagh, J. H. Ahn, J. W. Berkery, Y. Jiang, J. M. Bialek, J. Kim, H. S. Han, S. H. Hahn, Y. M. Jeon *et al.*, in Proceedings of 27th International Conference on Fusion Energy, Gandhinagar, India, 2018.
- ²⁵Y. Jiang, S. A. Sabbagh, Y. S. Park, J. W. Berkery, J. H. Ahn, J. D. Riquezes, J. G. Bak, W. H. Ko, J. Ko, J. H. Lee *et al.*, *Nucl. Fusion* **61**, 116033 (2021).
- ²⁶A. Bondeson and D. Ward, *Phys. Rev. Lett.* **72**, 2709 (1994).
- ²⁷J. W. Berkery, S. A. Sabbagh, R. Betti, B. Hu, R. E. Bell, S. P. Gerhardt, J. Manickam, and K. Tritz, *Phys. Rev. Lett.* **104**, 035003 (2010).
- ²⁸S. A. Sabbagh, J. W. Berkery, R. E. Bell, J. M. Bialek, S. P. Gerhardt, J. E. Menard, R. Betti, D. A. Gates, B. Hu, O. N. Katsuro-Hopkins *et al.*, *Nucl. Fusion* **50**, 025020 (2010).
- ²⁹J. W. Berkery, S. A. Sabbagh, R. E. Bell, S. P. Gerhardt, and B. P. LeBlanc, *Phys. Plasmas* **24**, 056103 (2017).
- ³⁰J. W. Berkery, S. A. Sabbagh, A. Balbaky, R. E. Bell, R. Betti, A. Diallo, S. P. Gerhard, B. P. LeBlanc, J. Manickam, J. E. Menard, and M. Podestà, *Phys. Plasmas* **21**, 056112 (2014).
- ³¹W. Ko, H. Lee, D. Seo, and M. Kwon, *Rev. Sci. Instrum.* **81**, 10D740 (2010).
- ³²B. Stratton, R. J. Fonck, K. P. Jaehnig, N. Schechtman, and E. J. Synakowski, in *Proceedings of the IAEA Technical Committee Meeting on Time Resolved Two- and Three-Dimensional Plasma Diagnostics, Najoya, Japan* (IAEA, Vienna, 1991).
- ³³F. Imbeaux, S. D. Pinches, J. B. Lister, Y. Buravand, T. Casper, B. Duval, B. Guillerminet, M. Hosokawa, W. Houlberg, P. Huynh *et al.*, *Nucl. Fusion* **55**, 123006 (2015).
- ³⁴F. Chollet, *Deep Learning with Python* (Manning Publications Co., 2018).
- ³⁵A. Piccione, J. W. Berkery, S. A. Sabbagh, and Y. Andreopoulos, *Nucl. Fusion* **60**, 046033 (2020).
- ³⁶A. Piccione, J. W. Berkery, S. A. Sabbagh, and Y. Andreopoulos, *Nucl. Fusion* **62**, 036002 (2022).
- ³⁷B. Friedland, *Control System Design: An Introduction to State-Space Methods* (Dover Publications, Inc., 1986).
- ³⁸S. A. Sabbagh and J. W. Berkery, “Disruption event characterization and forecasting in tokamaks,” (2023), see <http://arks.princeton.edu/ark:/88435/dsp018p58pg29j>.