MASTER

Real-time, model-based event detection in tokamaks

Akkermans, Gijs R.A.

Award date:
2017

TU/e Technische Universiteit
Eindhoven
University of Technology

Department of Applied Physics

Real-time, Model-based Event Detection in Tokamaks.

*Thesis for the Masters Science &
Technology of Nuclear Fusion and Applied
Physics (plasma track).*

Gijs R. A. Akkermans, BSc.

Supervisors:
dr. ir. Federico Felici
dr.ir. Jan van Dijk
ir. Thomas Blanken

Eindhoven, January 2017

Abstract

In order to enable real-time control of low-to-high confinent (L-H) transitions and Edge Localised Modes (ELMs) in tokamak plasmas, a real-time detector for L-H transitions and ELMs is needed. In this work, a detector is designed which can detect transitions between the L-mode and H-mode with a typical timing error of 10 ms and which can detect the timing and energy loss of 95% of ELMs in a set of 14 ELMy and 4 quiescent H-mode discharges on the Tokamak à Configuration Variable (TCV). The detector is based on an Interacting Multiple Model Kalman filter, where the constituent models are simple descriptions of signal features that are expected during transitions and ELMs. The detector uses measurements of the intensity of H_α line radiation measured with a photo-diode detector, density measurements from the Far Infra-Red interferometer, and magnetic measurements of the plasma diamagnetic flux, related to the internal energy. The losses of plasma internal energy during ELMs are determined from the diamagnetic flux and the ELM timing from the detector. Moreover, the performance of the detector has been assessed by an objective cost function, and optimised using single-parameter sweeps of important detector parameters. The results show good detector performance, almost on par with post-shot visual data analysis, and motivate the implementation of the detector on real-time tokamak control systems.

Table of Contents

List of Acronyms

LHT   Low-to-High confinement Transition
HLT   High-to-Low confinement Transition
ELM   Edge Localised Mode
ETB   Edge Transport Barrier
LCFS   Last Closed Flux Surface
NBI   Neutral Beam Injection
PD   Photo-Diode
FIR   Far Infra-Red interferometer
DML   Dia-Magnetic Loop
XTe   X-ray electron Temperature
MHD   Magneto-HydroDynamics
GSE   Grad-Shafranov Equation
ZF   Zonal Flow
LCO   Limit-Cycle Oscillation
RV   Random Variable
PDF   Probability Density Function
MJLS   Markov Jump-Linear System
MM   Multiple Model
AMM   Autonomous Multiple Model
CMM   Cooperating Multiple Model
HR   Hypothesis Reduction
GPB   Generalised Pseudo-Bayes
IMM   Interacting Multiple Model
TCV   Tokamak à Configuration Variable
JET   Joint European Torus
ITER   International Thermonuclear Experimental Reactor

Chapter 1

Introduction

In the face of increasing worldwide energy consumption, finite availability of fossil fuel resources, and ever-growing concerns about global warming, the development of renewable, clean, and safe energy sources is presently an important scientific objective. Nuclear fusion, in particular, offers the advantages of year-round energy production without need for storage, lack of CO_2-emissions and abundant fuel reserves likely lasting tens of thousands of years for 'easy' D-T fusion and billions of years for the more difficult D-D fusion[1]. Also, fusion compares favorably to conventional nuclear fission; fusion reactors are intrinsically immune to catastrophic accidents like a 'meltdown' and produce only moderate amounts of short-lived nuclear waste, which is safe to handle within a few hundred years [2].

The challenge of fusion is a technological one. The fuel has to to be heated to a temperature of around 165 million °C, at which point it has entered the plasma state. Then, the hot plasma must be confined at a high enough pressure for a long enough time to produce more power than is required for external heating. The most widely researched strategy to confine a hot fusion plasma is to use a machine called a 'tokamak', a toroidal device featuring a helically wound toroidal magnetic field which limits energy transport in the outward direction.

In order to optimise the operations of tokamaks towards maximum stable energy production at minimal cost and maintenance time, real-time control is a necessary tool. For example, control is used to avoid or stabilise various plasma instabilities [3, 4], and could be used to steer plasma parameters and profiles towards regimes of interest for scientific study or optimise power production in (mainly future) power-producing tokamaks.

An increasingly common practice is to base the control strategy on an estimate of the plasma state as given by a dynamic state observer [5, 6]. Such dynamic state observers base their state estimate not only on measurements but also on a model of the system. This remedies the necessity for ill-conditioned static profile fits, and promotes accuracy of estimated profiles by inclusion of models of e.g. transport, heating, current drive and fuelling sources. At the same time, systematic discrepancies between real measurements and expected measurements can be used to correct the estimated state evolution.

In some cases, large discrepancies between an observer's expected measurements and the real measurements are a reason for worry. For example, they may point to diagnostic faults or model-reality mismatches. Some specific 'events' in tokamaks produce typical deviations in diagnostic signals but are generally not modelled by state observers because they are too unpredictable. Some of these do signal a real change in the plasma state and/or its dynamics, and their detection may

be a reason for the real-time control system to switch to a different controller or to a different internal model for one of its observers.

Two events that present a large change in plasma dynamics are the L-H transition and ELMs. The former is the transition from the standard 'low confinement' L-mode to the 'high confinement' H-mode, and signals an improvement of confinement near the plasma edge. This causes an increase in plasma temperature and density, and a decrease in recycling leading to lower neutral concentrations. The L-H transition can be triggered in many existing tokamaks by the application of sufficient heating power. The scaling of the power threshold with main plasma parameters is known from experiments, but it is also known to depend on details like which external heating system is used using which injection geometry, as well as the geometry of the magnetic field. Therefore, when a heating power close to the scaling law value is applied, it is not always clear beforehand whether an L-H transition will be triggered. Especially the timing of the transition can be hard to predict with high time resolution.

During an H-mode, Edge Localised Modes (ELMs) are periodically triggered under some conditions. These peeling-ballooning mode are driven unstable when the edge pressure gradient and/or current density grow too large. ELMs cause periodic expulsions of particles and energy, which also causes a burst of recycled neutrals. On the whole, ELMs deteriorate energy and particle confinement, which has a negative effect on the fuel density and temperature, but also allows the removal of helium 'ash' and other impurities out of the plasma, which would otherwise limit performance. The severity of the confinement deterioration increases with the ELM frequency. ELMs also present a pulsed heat load to the divertor, which is especially at risk of being damaged by 'large' ELMs (in terms of expelled energy). Since low-frequency ELMs are typically larger, higher ELM frequencies are better for the divertor.

Apart from allowing observers (and controllers) to use models that more accurately reflect the plasma dynamics in different operating regimes, real-time detection of L-H transitions (and H-L back-transitions) and ELMs is useful in a number of situations. Firstly, it can be used to verify whether an applied heating scheme has the desired effect of triggering an L-H transition. If an accidental back-transition is detected, this may be a reason to increase the heating power or decrease the plasma density, so that the plasma can transition back into H mode. When the plasma enters an ELM-free H-mode, the improved particle confinement may lead to an uncontrolled density increase, eventually leading to a high-density disruption. Therefore, the detection of an L-H transition without the subsequent detection of ELMs may be a reason to actively induce ELMs to limit the density increase, or to reduce heating in order to cause a back-transition rather than a disruption. In a typical ELMy H-mode, detection (and control) of the ELM frequency and amplitude (in terms of expelled energy) is needed to ensure an optimum trade-off between confinement degradation, divertor damage, and impurity removal.

Summarising, given the typically available diagnostics, the L-H transition and ELMs are not predictable to a high time resolution in real-time, and their detection in real-time would be useful to tokamak operation. Hence, this thesis adresses the research question *"can specific events in the TCV tokamak, particularly, L-H transitions and ELMs, be detected in real-time using a model-based approach?"*. For L-H transitions (and H-L back-transitions), only the times at which they occur are of interest. For ELMs, the timing as well as the expelled energy is of interest.

In order to be useful for real-time applications, the detection algorithms used should be causal and real-time capable. It is also required that events are detected accurately within 10 ms, the typical thermal time scale of the Tokamak 'a Configuration Variable (TCV). Real-time capability places a limit on the computational cost of the algorithm; each computational step must be performed within one sampling time. The first sub-question is: *Which diagnostic signals have to be taken into account to detect the events?*

Furthermore, a model-based approach is chosen, mainly because it offers more flexibility to the device-specific diagnostics that are available. This means that if a model-based event detector designed for one tokamak is applied to another tokamak with different diagnostics measuring the same plasma parameters, essentially only the model of the measurement system would need to be changed, while the model for the plasma's internal evolution only needs to be tweaked to the possible differences in plasma parameters.

Both the L-H transition and ELMs require complicated physics modelling to be reproduced in self-consistent simulations, but can be detected relatively easily by visual analysis of measurement data. A full time-dependent model of the L-H transition requires accurate modelling of the evolution of density, temperature and momentum due to turbulence. The effect of plasma kinetic profiles and flows on turbulence levels and thus transport coefficients has to be taken into account, in principle by analysing the (in)stability of different modes from linearised MHD (magneto-hydrodynamics) equations for small perturbations [7], or by sufficient simplifications [8, 9]. Finally, a trigger mechanism should be included through which the turbulence levels and/or kinetic profiles affect the plasma flow and/or the radial electric field in a way that significantly decreases turbulence, for example through Reynolds stress [9] or interaction with edge neutrals [7]. Especially on this last point, modeling approaches vary and consensus about which model best describes the L-H trigger mechanism is still forming [10].

Meanwhile, experts can often accurately pinpoint the time at which an L-H transition actually occurs by looking at diagnostic data after the shot has concluded. They identify the moment at which diagnostics indicate a sudden decrease in recycled neutrals and/or a sudden increase in the time derivative of particle density, or they can compare radial kinetic profiles at different time steps to look for the formation of a pedestal. In [11], Gonzalez et al. present a neural network which also uses a data-based approach and is trained to automatically recognise L-H transitions in post-shot analysis. However, this neural network (NN) is not causal, because a classification of the plasma state at some point in time depends on measurement data in a time window $\pm 100\,\text{ms}$ or $\pm 500\,\text{ms}$ around that time, depending on which classifier is used. This implies a time delay of at least $100\,\text{ms}$ in real-time applications. Therefore, the work of Gonzalez et al. is not necessarily a good starting point for designing a real-time LHT detection algorithm.

The modelling of a single ELM requires non-linear MHD simulations, and due to the large separation of time scales, few modelers attempt to model multiple ELM cycles. Simulation results typically agree qualitatively with diagnostic measurements, but quantitative agreement can only be obtained for a limited range of plasma edge parameters [12]. On the other hand, ELMs are clearly visible as spikes on recyled neutral diagnostics.

Instead of using full models, we use simple models which do not necessarily contain all of the physics, but whose time-evolution corresponds to the diagnostic features that are to be recognised. This allows the models to be evaluated in real-time. A dynamic state estimator is based on each of these models, and each estimator attempts to track the plasma state. The estimator whose model consistently expects measurements that are closer to the real measurements than any other estimator's prediction is taken to represent the mode that the system is in. This leads to the second sub-question: *How should the evolution of the underlying plasma parameters producing the measured signals be modelled during different modes of operation?*

The multi-model (MM) state estimation method is expected to offer some advantages over other possible forms of signal pattern recognition. For example, a simple algorithm with detection tresholds on various signal levels and their time derivatives would likely be less robust to signal noise or to differences in signal behavior between discharges, compared to MM methods. A full-model approach would lead to large computational loads when implemented in real-time. A neural network (NN) would be less transparent in terms of algorithm operation, which may lead to more difficulties in debugging. Also, the variables in a NN would have little to no physical meaning,

so there is no possibility of improving the algorithm as better physics models become available at lower computational cost.

In a multi-model system where the active mode can change during the shot, the amount of possible mode histories grows exponentially as a function of time, similar to a tree that branches at every time step. In order to come to an optimal estimate, a multi-model filter should take into account each possible mode history, which quickly leads to prohibitively large computational requirements. Possible remedies are to prune unlikely branches of the model history tree, or to combine multiple branches into a single history. After every time step, the state estimates of the different estimators are combined in a way that reflects how the model history tree was pruned or merged. This leads to the third sub-question: *Which hypothesis reduction strategy should be used to combine the state estimates based on different models?*

For this research project, a training dataset is used consisting of measurements from 13 discharges on the TCV tokamak. The detector will be optimised to detect the events in these shots. Its performance will then be evaluated on the basis of a test dataset, consisting of data from another 5 TCV discharges. If the detector shows satisfactory performance, it will eventually be implemented on the TCV real-time control system to assist with feedback control of L-H transitions and ELM-induced heat loads.

Chapter 2 gives a theoretical background about the events to be detected and the detection strategies used. In Chapter 3, simple models are formulated which reproduce signal features accompanying the events in question. In Chapter 4, a detector based on these models is used to detect events in TCV measurement data. Chapter 5 draws some conclusions to answer the research question, and presents an outlook to future research.

Chapter 2

Theory

This chapter presents a theoretical background on relevant plasma events and detection strategies. Section 2.1 focusses on physics aspects of the H-mode, especially those that have shaped the event detector's design. Section 2.2 details the working of the multiple model Kalman filter, the format of the designed detector. Finally, Section 2.3 is about the analysis of the diamagnetic loop signal, which can be used in conjunction with detector results to analyse the energy losses during ELMs.

2.1 The High confinement mode

In this section, a number of features of the high confinement mode are discussed, such as Edge Localised Modes in Section 2.1.1 and the scaling of threshold power in Section 2.1.2. Detection of low to high confinement transitions is treated in Section 2.1.3, and Section 2.1.4 describes possible physical mechanisms that may trigger a transition to the H-mode.

The 'high confinement' H-mode is an operational regime that has been achieved in most tokamaks. It greatly increases both energy and particle confinement when a threshold of auxiliary heating power is crossed. This in turn increases the plasma density and temperature, and hence the kinetic pressure, which makes the H-mode a very attractive mode of operation for power generation. For example, the record amount of fusion power obtained in the Joint European Torus (JET) was produced in H-mode [13], and the International Thermonuclear Experimental Reactor (ITER) under construction in southern France will rely on the H-mode to reach its goal of producing significant net fusion power [14].

Typically, a plasma's energy confinement time and particle confinement time are increased by a factor of about two by the appearance of the H-mode [15], allowing the plasma to store more thermal energy for the same input power. During a transition from the standard 'low confinement' L-mode to the 'high confinement' H-mode, only the confinement near the plasma edge improves. There, a large pressure gradient is established, which is mostly due to a large density gradient. This edge pressure gradient is apparently caused by a local improvement in confinement, a so-called Edge Transport Barrier (ETB). On the order of the confinement time, this ETB causes an increase in density to spread from the edge to the centre of the plasma. Figure 2.1a shows density measurements performed in ASDEX around the time of an L-H transition, showing the formation of the ETB and inward spread of density increase.
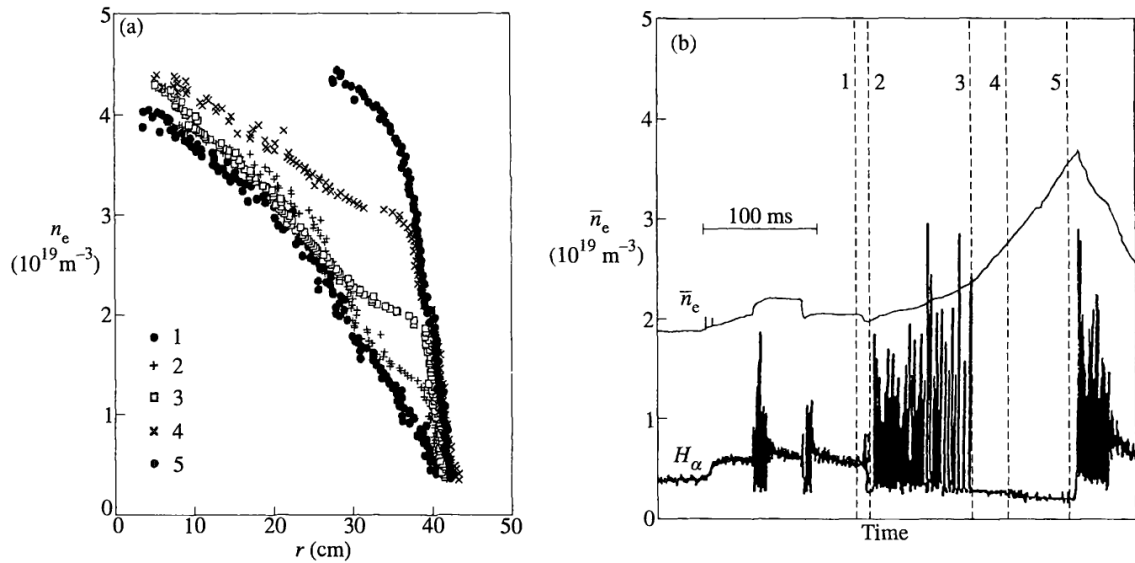

Figure 2.1: (a) The density profile measured at ASDEX at different times around an L-H transition, showing the formation of an Edge Transport Barrier. (b) The line-averaged density and H_α signal as function of time. 5 time instances are marked, to which the profiles under (a) correspond. [15, 16]

The H-mode also has certain negative aspects. For one, the increase in density due to improved particle confinement is uncontrolled and may lead to a back-transition to the L-mode or even a total loss of confinement due to high-density disruptions of the plasma [17]. Also, the increased impurity confinement time may lead to an accumulation of impurity ions in the plasma, which increases radiation heat losses and dilutes the concentration of hydrogen isotopes, lowering the fusion power in (future) power-producing devices.

2.1.1 Edge Localised Modes

A final typical feature of the H-mode are the so-called Edge Localised Modes (ELMs), which, in a simplified view, are peeling-ballooning modes in magneto-hydrodynamics (MHD), although more unstable mechanisms may be at play depending on, e.g., the ELM type. The high confinement of the H-mode causes a large edge pressure gradient and associated current density (so-called neoclassical bootstrap current), and these can drive the ELMs unstable. ELMs periodically expel energy and particles from the plasma. Between ELMs, the pressure continuously increases due to the good confinement, until the stability limit is reached again. As such, ELMs ameliorate the problems of uncontrolled density and impurity accumulation in the quiescent (ELM-free) H-mode, but also deteriorate confinement during H-mode. Furthermore, ELMs present a pulsed heat load to the divertor (i.e., the exhaust), which may lead to unacceptable damage to or a decrease in the lifetime of the divertor .

ELMs are generally categorised in three types. Type I ELMs cause large, infrequent heat loads, type III ELMs cause an almost continuous heat load, and type II ELMs are intermediate in terms of both frequency and size [15, 18]. For a higher ELM frequency, the problems of impurity accumulation and high transient heat loads to the divertor are generally diminished, but overall confinement deteriorates. In order to ensure an acceptable trade-off between these different factors, observation and possibly control of the ELM frequency and amplitude is neccesary. Experimentally, the most precise way to determine which ELM type is present is to observe the dependence

of their frequency on heating power. For type I ELMs, the repetition frequency increases with increased power crossing the separatrix, while for type III ELMs, the repetition frequency decreases with increased power crossing the separatrix [19]

2.1.2 Scaling laws for threshold power

Experiments have been performed on different tokamaks regarding the threshold power needed to enter the H-mode and its dependence on different parameters such as plasma density, plasma size and magnetic field strength. For most tokamaks, auxiliary heating is required to overcome the threshold, but on TCV, L-H transitions have been achieved using only ohmic heating for a variety of plasma conditions [20, 21]. The results of threshold power experiments on different tokamaks can be summarised by 'scaling laws', log-linear fits of the threshold power with main plasma parameters, see e.g. [22, 23]. Such scaling laws show that increased density, plasma size, and magnetic field all require a higher threshold power. Scaling laws have some value in providing extrapolations to next-generation devices, but due to differences between machines, the measured threshold powers generally differ from those predicted by the scaling by about 30% [22], limiting the confidence in these predictions.

The threshold power is also known to depend on other aspects. For example, for a single-null divertor configuration in which the ion $\mathbf{B}\times\nabla B$ drift points towards the divertor, the threshold power is roughly twice as low as when the ion $\mathbf{B}\times\nabla B$ drift points away from a single-null divertor, or when a double-null divertor geometry is used. The threshold also lies a factor two lower in deuterium plasmas compared to hydrogen plasmas, and is lower when impurity levels and particle recycling are reduced [22].

2.1.3 Detection of the low to high confinement transition

The occurence of an L-H transition and ELMs is most visible as a decrease in H_α line radiation, which can be measured with e.g. photodiodes. When hydrogen ions impact on the divertor, part of them are recycled - they recombine with an electron and are released again, re-entering the plasma as neutral hydrogen. Due to collisions with plasma electrons, some of these are excited to higher energy levels and will emit photons upon relaxation. The intensity of the Balmer-α or H_α line corresponding to the relaxation from the third lowest to the second lowest hydrogen energy level is the most intense of these lines. It is measured in many tokamaks to monitor the rate of hydrogen recycling.

During an L-H transition, edge particle transport and recycling suddenly decrease, causing a sudden drop in the H_α signal. This can be seen in Figure 2.1b, just before the second marked time instant. Conversely, each ELM causes a sudden burst of particle recycling, accompanied by a burst in the H_α signal. Such bursts are seen in Figure 2.1b between the times marked 2 and 3.

Determination of L-H transition times is almost exclusively done post-shot, and often manually by visual inspection of the acquired data. A sudden drop in the H_α signal occuring closely after the application of significant heating power is sought. If the H_α drop due to the L-H transition is 'hidden' among other patterns in the H_α signal, one can examine other signals, e.g., look for a sudden change in the time derivative of the plasma density, or the formation of a pressure pedestal using e.g. localised electron density measurements.

Such an approach was formalised into a neural networking algorithm in [11]: Wavelets were used to locate a clear drop in the H_α signal. For discharges where such a drop could not be

found, a Support Vector Machine (SVM) classification algorithm was used, which classified time instances as L or H mode based on a number of diagnostic signals selected by a feature extractor. The SVM algorithm was trained using the discharges where a clear H_α drop was found.

2.1.4 Low to high confinement transition mechanisms

Many theories aiming to explain the L-H transition exist, see [24] or [10] for a review. A common element in many of them is the reduction of turbulent transport by a sheared poloidal flow, as developed theoretically by Biglari, Diamond and Terry [25]. The sheared poloidal flow rips apart turbulent eddies as visualised in Figure 2.2. The result is a reduction in turbulent transport, i.e., an Edge Transport Barrier.

After the onset of the H-mode, a sheared $\mathbf{E} \times \mathbf{B}$ flow is sustained by the radial electric field E_r and the toroidal magnetic field B_φ. The flow is sheared because of the appearance of a deep radial electric field well near the plasma edge, at the time of the L-H transition. E_r has a large gradient on both sides of this well.

The most important ingredients to the radial electric field are [10]:

$$E_r = \frac{1}{Zen_e}\nabla p_i - v_\theta B_\varphi + v_\varphi B_\theta, \tag{2.1}$$

where Z is the ion charge number, e the elementary charge, n_e the electron density, p_i the ion kinetic pressure, and $\mathbf{v}$ the fluid flow veclocity. The first term is the diamagnetic term and contains the ion pressure gradient. This is actually also a $\mathbf{v}7\mathbf{B}$ term, given by the ion diamagnetic drift velocity $\mathbf{v_{dia}} = \frac{1}{Zen_e}\frac{\nabla p_i 7\mathbf{B}}{B^2}$ [27] crossed with the magnetic field $\mathbf{B}$. This diamagnetic drift is a result of the Larmor orbits of charged particles, also affected by the force due to a pressure gradient. The diamagnetic contribution to E_r grows after the L-H transition with the formation of the ETB, deepening the electric field well and stabilising the H-mode. The $v_\varphi B_\theta$ component is negative when a toroidal flow is present in the opposite direction to the plasma current, for example when neutral beams are injected in the counter-current direction. Indeed, beam injection geometry has been found to affect the power threshold [28]. The evolution of the poloidal flow velocity v_θ is determined by the poloidal momentum balance, containing the Lorentz force, poloidal viscous dampening, friction with neutrals, and turbulent Reynolds stress (the average effect of turbulent fluctuations on the mean flow). An important feature of L-H transition models is the mechanism which allows the formation of an electric field well [10].

A popular trigger mechanism in L-H transition theories [10, 26, 29] is a reduction in turbulence by sheared 'Zonal Flows' (ZFs). These ZFs are supposedly formed spontaneously out of microscopic turbulence by the Reynolds stress, which couples energy from small scales to larger scales. The flow velocity in a ZF oscillates as a function of space, meaning the flow is sheared
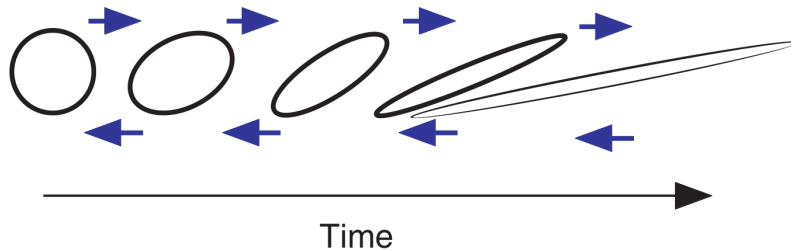

Figure 2.2: A turbulent eddy is ripped apart by sheared flow [26].

and can dampen the turbulence that caused it. If a large enough pressure gradient is built up before the ZF is dampened, the resulting radial electric field well can drive a sheared mean flow that stabilises the increased pressure gradient. Until a large pressure gradient develops, ZFs can repeatedly be generated by turbulence, decrease the turbulence level, be dampened as their driving force vanishes, and be regenerated after the turbulence level increases again. These Limit Cycle Oscillations (LCOs) of Zonal Flows and turbulence levels can cause minor periodic expulsion of particles and energy close to the L-H treshold, and can lead to an 'Intermediate' phase or 'dithering' phase between the L and H-mode [30, 31]. Several recent experimental results suggest ZFs as a likely trigger mechanism for the L-H transition [32–35].

2.2 State estimation

Using insights into H-mode physics, seperate models for L-mode and H-mode plasmas will be formulated in Chapter 3. This section discusses state estimation for dynamical systems, and formulates algorithms that can be used to detect transitions between L and H mode in diagnostic data, using the aforementioned models. These algorithms are applied to TCV diagnostic data in Chapter 4.

In systems theory for the control of dynamical systems, the internal *state* of a system is the smallest subset of internal variables that fully reflects the system's internal condition. Hence, for given system dynamics and future external influences (inputs, noise), the evolution of the system is uniquely determined by the current state. Real-time knowledge of a system's state is useful both for verifying the effectiveness of a past control action and for choosing a present control action to steer the state in the right direction (feedback control). Two possible strategies of estimating the state are to reconstruct it using the measurements or to model its evolution using a state-space model (see Section 2.2.2). However, both these stategies suffer from some serious drawbacks.

Basing the state estimate only on measurements introduces difficulties when the state to measurement relation is not invertible or the inverse relation is ill-conditioned. Loosely speaking, this means that 'not all states are measured directly' (as is the case for many real systems) and it results in the inability to distinguish between two different states that produce (almost) the same output. Also, the measurements may be subject to noise that cannot be filtered optimally without more knowledge of the system dynamics.

Basing the state estimate only on the state-space model introduces the problem of choosing an initial condition. A simulated state based on a poor initial estimate may converge toward the real state of the system, but only if the system is stable and only on a time scale specified by the system dynamics. Even when the initial condition is estimated correctly, process noise and possible mismatch between the model and the actual dynamics will cause changes to the state, which cannot be accounted for in the model without taking into account measurements from the real system. Also, events that are difficult to anticipate because of their random nature cannot be accounted for using only limited models.

To come to an optimal estimate of a system's state, both measurements and the state-space model should be taken into account. This makes intuitive sense, since both the model and the measurements contain information about the system.

Section 2.2.1 summarises some results from probabilty theory and introduces the terminology used in later sections. In Section 2.2.2, state-space systems are introduced. State-space is a description of physical systems for which effective algorithms regarding state estimation have been formulated. Subsection 2.2.3 describes the Kalman filter, which is a minimum-error-covariance

filter for a linear state-space system under stochastic noise, using both a model and measurements. Section 2.2.4 describes multiple model state estimation, which is useful when the system at hand may transition between a number of modes that exhibit different dynamics.

2.2.1 Probability theory

A state-space system is often assumed to be subject to stochastic noise processes which degrade the state estimate. On the other hand, information in new measurements can be incorporated in the state estimate by updating or conditioning the state estimate on these new measurements, which will improve the state estimate. This section discusses the probability-theoretical background of both conditional probabilities and stochastic noise.

In a random experiment with a given number of possible outcomes, the probability $P(A)$ of event A can be defined as the relative frequency of the occurence of event A, when the experiment is conducted a large (infinite) number of times:

$$P(A) \equiv \frac{\text{Number of times } A \text{ occurs}}{\text{Total number of outcomes}} \qquad (2.2)$$

where $0 \leq P(A) \leq 1$ holds.

The *conditional probability* of A given the occurence of B is denoted $P(A|B)$, and is also known as the probability of A *conditioned on* the occurence of B. When $P(B)$ is nonzero, the conditional probability is given as:

$$P(A|B) = \frac{P(A \cup B)}{P(B)} \qquad (2.3)$$

where $P(A \cup B)$ is the joint probability of A and B (i.e., the probability of both A and B occuring), and $P(A)$ and $P(B)$ are so-called *a priori* or *prior* probabilities. These prior probabilities apply when nothing is known about the outcome (yet). In contrast, $P(A|B)$ is called the *a posteriori* or *posterior* probability of A since it is the probability of A after some possibly related piece of information (the occurence of B) has become known.

Random variables

When the possible outcomes of a random experiment are mapped to a range of real values, the numerical product of the experiment is a *random variable* (RV). When the experiment is actually performed, this produces a *realisation* of the random variable. This realisation of the RV is a certain ordinary value, and is hence not an RV itself.

A random variable X is characterised by its *probability density function* (PDF) $f_X(x)$ defined by:

$$P(a < X \leq b) = \int_a^b f_X(x)dx. \qquad (2.4)$$

The PDF is a nonnegative function, normalised in the sense that $\int_{-\infty}^{\infty} f_X(x)dx = 1$.

The conditional PDF $f_X(x|A)$ of the RV X given the occurence of the event A is given by:

$$P(a < X \leq b|A) = \int_a^b f_X(x|A)dx. \qquad (2.5)$$

and the conditional PDF of the RV X_1 given the realisation x_2 of the RV X_2 is given by:

$$f_{X_1|X_2}(x_1|x_2) = \frac{f_{X_1 \cup X_2}(x_1 \cup x_2)}{f_{X_2}(x_2)}, \tag{2.6}$$

where $f_{X_1 \cup X_2}(x_1 \cup x_2)$ is the joint PDF of X_1 and X_2, which yields a probability when integrated over both x_1 and x_2:

$$P(a_1 < X_1 \leq b_1, a_2 < X_2 \leq b_2) = \int_{a_2}^{b_2} \int_{a_1}^{b_1} f_{X_1,X_2}(x_1, x_2)dx_1 dx_2. \tag{2.7}$$

The expected value, or mean of X is denoted $E(X)$ or $\hat{x}$ and is the average of the realisations of X when the experiment is repeated a very large (infinite) number of times:

$$E(X) = \int_{-\infty}^{\infty} x f_X(x)dx. \tag{2.8}$$

The variance of X is denoted σ^2 and is a measure of the expected variation of X from its mean:

$$\sigma^2 = E\left[(X - \hat{x})^2\right]. \tag{2.9}$$

The notation $X \sim (\hat{x}, \sigma^2)$ signifies that X is a random variable with mean $\hat{x}$ and variance σ^2. Conditional means and variances are obtained by using the conditional PDF in the integral.

The PDF of a Gaussian RV X is given by:

$$f_X(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-(x - \hat{x})^2}{\sigma^2}\right], \tag{2.10}$$

which can be characterised fully by its mean $\hat{x}$ and variance σ^2. The notation $X \sim \mathcal{N}(\hat{x}, \sigma^2)$ indictates that X is a Gaussian RV with mean $\hat{x}$ and variance σ^2.

The covariance of two scalar RVs is given by:

$$C_{XY} = E\left[(X - \hat{x})(Y - \hat{y})\right], \tag{2.11}$$

and is a measure of the interdependence between their variations from their respective means.

The correlation between two scalar RVs is given by:

$$R_{XY} = E(XY). \tag{2.12}$$

If $R_{XY} = E(X)E(Y)$, X and Y are uncorrelated.

Covariance and correlation can be generalised for RVs that are vectors with each entry having its own PDF. Both covariance and correlation are then matrices. Say X is an $n \times 1$ RV and Y is an $m \times 1$ RV. Their covariance matrix is then given by:

$$C_{XY} = E\left[(X - \hat{x})(Y - \hat{y})^T\right] = \begin{bmatrix} E\left[(X_1 - \hat{x}_1)(Y_1 - \hat{y}_1)\right] & \cdots & E\left[(X_1 - \hat{x}_1)(Y_m - \hat{y}_m)\right] \\ \vdots & & \vdots \\ E\left[(X_n - \hat{x}_n)(Y_1 - \hat{y}_1)\right] & \cdots & E\left[(X_n - \hat{x}_n)(Y_m - \hat{y}_m)\right] \end{bmatrix}, \tag{2.13}$$

and their correlation matrix by:

$$R_{XY} = E(XY^T) = \begin{bmatrix} E(X_1 Y_1) & \cdots & E(X_1 Y_m) \\ \vdots & & \vdots \\ E(X_n Y_1) & \cdots & E(X_n Y_m) \end{bmatrix}. \tag{2.14}$$

The covariance matrix of a vector RV X with itself is called its auto-covariance matrix C_{XX}, which has ordinary variances on the diagonal and covariances everywhere else. The auto-correlation matrix R_{XX} of X with itself has $E(X_i^2)$ on the diagonal and correlations everywhere else. Furthermore, C_{XX} and R_{XX} are always symmetric and positive semi-definite.

The Gaussian distribution can also be generalised into the Gaussian multi-variate distribution. The notation $X \sim \mathcal{N}(\hat{x}, C)$ indicates that X is a vector Gaussian RV with PDF:

$$f_X(x) = \frac{\exp(-\frac{1}{2}(x-\hat{x})^T C^{-1}(x-\hat{x}))}{\sqrt{\det(2\pi C)}}, \tag{2.15}$$

where C^{-1} is the inverse and $\det(C)$ the determinant of the autocovariance matrix C.

A stochastic process $X(t)$ is an RV that changes in time. This means that its PDF is explicitly a function of time. Therefore, $X(t_1)$ and $X(t_2)$ at two different instances in time are two different RVs, and their realisations $x(t_1)$ and $x(t_2)$ are sampled from different PDFs so they may have different expected values (mean of the respective PDFs). The auto-correlation matrix of the stochastic process $X(t)$ is the correlation matrix of the RVs $X(t_1)$ and $X(t_2)$:

$$R_X(t_1, t_2) = E\left[X(t_1)X^T(t_2)\right]. \tag{2.16}$$

A stochastic process is white if $X(t_i)$ is independent of any $X(t_j)$ with $i \neq j$. This is satisfied if the autocorrelation is zero for any two different times:

$$R_X(t_i, t_j) = R(t_i)\delta_{i-j} = \begin{cases} R(t_i) & \text{if } i = j, \\ 0 & \text{if } i \neq j \end{cases}, \tag{2.17}$$

where the Kronecker δ_k function is defined by the above equation.

2.2.2 Linear state-space systems

A class of systems for which effective algorithms regarding state estimation have been formulated are finite-dimensional, linear *state-space systems*. This work focusses on discrete-time state-space systems. The *state vector* x_k, an $n \times 1$ column vector, is the smallest possible subset of internal variables which fully and uniquely reflects the system's internal condition at time k. The *state-space* is the multi-dimensional space of all states that the system can reach. Multiple trajectories through state-space (caused by different external influences in the past) can lead to the same x_k. However, for given external influences, the state x_k uniquely determines the future trajectory through state-space.

The sequence of all state vectors up to time k are denoted $\{x_k\} = [x_1, x_2, \dots, x_k]$. The state vector evolves in discrete time steps according to linear dynamics:

$$x_k = F_{k-1}x_{k-1} + G_{k-1}u_{k-1} + w_{k-1}. \tag{2.18}$$

Here, F_{k-1} is the $n \times n$ system matrix, which determines the autonomous evolution of the state vector. The second term depends on the input vector u_{k-1} (an $l \times 1$ column vector), which lists the input signal values that are externally imposed to the process at the previous time step $k-1$. This input vector is multiplied by the input matrix G_{k-1}, which is $n \times l$ and describes in what way the inputs affect the states. w_{k-1} is the process noise, an $n \times 1$ column vector random variable, which is assumed to be sampled from a zero-mean stochastic process. w_k can be interpreted as disturbances to the system's state which are not explicitly modelled.

Information regarding the system's state is extracted by performing measurements. The measurement vector y_k is an $m \times 1$ column vector consisting of the output values of the different sensors used to observe the system. These measurements depend linearly on the state variables:

$$y_k = H_k x_k + v_k, \tag{2.19}$$

where H_k is the $m \times n$ output matrix specifying this dependence and v_k is the measurement noise, which has dimension $m \times 1$, and is also assumed to be sampled from a zero-mean stochastic process. v_k contains the unmodelled disturbances which do deteriorate the measurements but have no further effect on the state of the system itself. F_k, G_k, and H_k may depend explicitly on the time instant k, but are constant for many systems. The latter systems are called time-invariant.

For a linear state-space system, the general relations $x_k = f(x_{k-1}, u_{k-1}, w_{k-1}, k)$ and $y_k = h(x_k, v_k, k)$ which hold for a general non-linear system can be written in a linear form. The Kalman filter which is introduced in the next section depends on this linearity.

Using the state-space model for a system, the system or 'plant' can be simulated as depicted in Figure 2.3, where the 'model evolution' block and the 'synthetic diagnostics' block represent Eqs. 2.18 and 2.19, respectively. If the process and measurement noise covariances, the model-reality mismatch, and the error in the assumed initial condition are all small, this simulation will output synthetic measurements close to the real measurements.
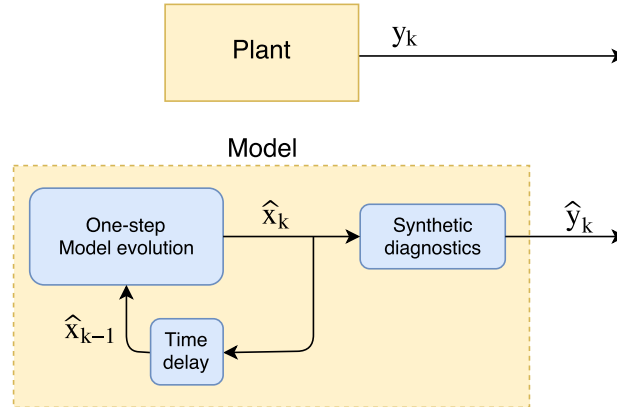
Figure 2.3: Block diagram of a state-space model simulating a real system. At every time step, the simulated state $\hat{x}_k$ is evolved using the model dynamics, and synthetic measurements $\hat{y}_k$ are produced based on the state, which may be compared to the plant measurements y_k to validate the model.

2.2.3 The discrete-time Kalman filter

The Kalman filter incorporates both the system model and measurements in order to come to an optimal state estimate. Here, optimal refers to the property of minimum variance estimation error, given stochastic noise with known covariance [36]. The Kalman filter is iterative and causal, making it suitable for real-time implementation. Suppose that the system whose states are to be estimated is given by:

$$x_k = F_{k-1} x_{k-1} + G_{k-1} u_{k-1} + w_{k-1}, \tag{2.20a}$$

$$y_k = H_k x_k + v_k, \tag{2.20b}$$

where the noise processes $\{w_k\}$ and $\{v_k\}$ are white, zero-mean, and uncorrelated, and have known covariance matrices Q_k and R_k, respectively:

$$w_k \sim (0, Q_k), \tag{2.21a}$$

$$v_k \sim (0, R_k), \tag{2.21b}$$

$$E(w_k w_j^T) = Q_k \delta_{k-j}, \tag{2.21c}$$

$$E(v_k v_j^T) = R_k \delta_{k-j}, \tag{2.21d}$$

$$E(v_k w_j^T) = 0. \tag{2.21e}$$

Again, for many systems, F, G, and H are time-invariant, but the Kalman filter can be derived for the general case of explicitly time-dependent system dynamics.

A priori and a posteriori state estimates

Because the state is influenced by stochastic noise, x_k is a random variable (RV). In real-time applications, the goal at time k is to estimate the state x_k using the model and all available (past) measurements. Just before the measurement y_k becomes available, an estimate $\hat{x}_{k|k-1}$ of x_k is taken as the expectated value of the RV x_k, which can be conditioned at most on all past measurements:

$$\hat{x}_{k|k-1} = E(x_k|y_1, \ldots, y_{k-1}) = E(x_k|\{y_{k-1}\}). \tag{2.22}$$

$\hat{x}_{k|k-1}$ is named the *a priori* estimate or simply *prior* estimate of x_k, where the subscript $k|k-1$ denotes that it is only conditioned on prior measurements. After the measurement y_k becomes available, the estimate of x_k can be improved by also conditioning the expected value of x_k on y_k:

$$\hat{x}_{k|k} = E(x_k|y_1, \ldots y_k) = E(x_k|\{y_k\}). \tag{2.23}$$

$\hat{x}_{k|k}$ is the *a posteriori* or *posterior* estimate of x_k. Since the first measurement is taken at time $k = 1$, there are no measurements on which to condition $\hat{x}_{0|0}$. Instead, it is reasonable to take the unconditioned expected value: $\hat{x}_{0|0} = E(x_0)$.

P_k denotes the error covariance of a state estimate. It is taken as the auto-covariance of the RV x_k and represents the uncertainty in the estimate $\hat{x}_k$. $P_{k|k-1}$ is the covariance of the estimation error of $\hat{x}_{k|k-1}$ and $P_{k|k}$ is the covariance of the estimation error in $\hat{x}_{k|k}$:

$$P_{k|k-1} = E\left[(x_k - \hat{x}_{k|k-1})(x_k - \hat{x}_{k|k-1})^T\right], \tag{2.24a}$$

$$P_{k|k} = E\left[(x_k - \hat{x}_{k|k})(x_k - \hat{x}_{k|k})^T\right]. \tag{2.24b}$$

$P_{0|0} = P_0$ is the error covariance of the initial state.

Kalman filter equations

Between measurements, the prior state estimate and error covariance (Eqs. 2.22 and 2.24a) for the new time step k are calculated on the basis of the posterior state estimate and error covariance (Eqs. 2.23 and 2.24b) of the previous time step $k-1$, using the system dynamics (Eq. 2.20a) [36]:

$$\hat{x}_{k|k-1} = F_{k-1}\hat{x}_{k-1|k-1} + G_{k-1}u_{k-1}, \tag{2.25a}$$

$$P_{k|k-1} = F_{k-1}P_{k-1|k-1}F_{k-1}^T + Q_{k-1}. \tag{2.25b}$$

These are the time-update equations for $\hat{x}$ and P. The state estimate transforms just as the state, except without noise, and the estimation error covariance is determined by the state's autonomous evolution and process noise.

Each time step, the posterior state estimate $\hat{x}_{k|k}$ and its estimation error covariance $P_{k|k}$ are calculated by updating the prior state estimate $\hat{x}_{k|k-1}$ and its covariance $P_{k|k-1}$ on the basis of the measurement y_k as [36]:

$$K_k = P_{k|k-1}H_k^T(H_kP_{k|k-1}H_k^T + R_k)^{-1}, \tag{2.26a}$$

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k(y_k - H_k\hat{x}_{k|k-1}), \tag{2.26b}$$

$$P_{k|k} = (I - K_kH_k)P_{k|k-1}(I - K_kH_k)^T + K_kR_kK_k^T, \tag{2.26c}$$

where K_k is the Kalman gain. Effectively, the PDF of the RV x_k is conditioned on the measurements, causing a shift towards states that would cause such measurements, and the mean $(\hat{x}_{k|k})$ and covariance $(P_{k|k})$ of the conditional PDF are calculated. These are the measurement update equations for $\hat{x}$ and P. At every cycle of the observer, evaluation of the Kalman filter equations (Eqs. 2.25 and 2.26) yields the a posteriori estimate for that time step.

The Kalman gain K_k determines how strongly the posterior state estimate is modified compared to the prior estimate on the basis of the difference between the real measurement y_k and the expected measurement $\hat{y}_{k|k-1} = H_k\hat{x}_{k|k-1}$. This difference $y_k - \hat{y}_{k|k-1}$ is called the innovations z_k, and is the part of the measurements containing new information. It can be shown that the innovations is zero-mean and has covariance $S_k = H_kP_{k|k-1}H_k^T + R_k$ [36]. Note that it is precisely S_k which is inverted when calculating the Kalman gain in Equation 2.26a. During measurements, the innovations can be calculated, and the statistical mean and covariance can be compared to the theoretical values to verify filter performance.

The Kalman filter is depicted in Figure 2.4. Compared to the model-only simulation of Figure 2.3, not only the model dynamics are used to evolve the state estimate, but also the innovations are used to refine the estimate during the measurement update.
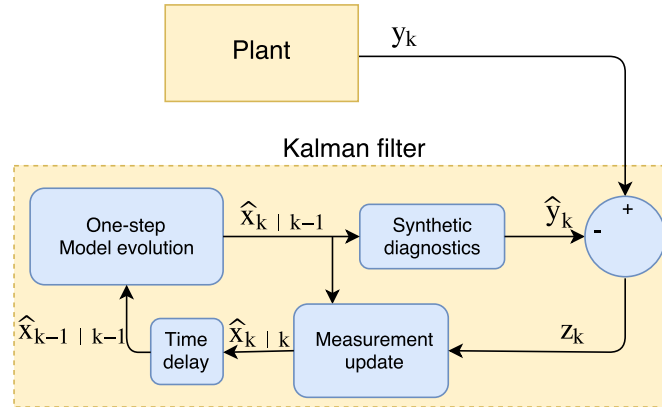
Figure 2.4: Block diagram of a Kalman filter estimating the state of a real system. At every time step, the state estimate $\hat{x}_k$ is evolved using the model dynamics, and synthetic measurements $\hat{y}_k$ are produced based on the state. At each time step, $\hat{y}_k$ is compared to the plant measurements y_k, and the difference is used to update $\hat{x}_k$.

The Kalman gain K_k is large when the measurement noise covariance R_k is small and the estimation error covariance $P_{k|k-1}$ is large. In other words, the filter relies more heavily on measurements when these are reliable while the state estimate is not. An unreliable state estimate may be caused by large levels of process noise (large Q_k) or within the first few time steps, by an inaccurate initial condition (large P_0).

The posterior estimation error covariance $P_{k|k}$ depends on the prior covariance $P_{k|k-1}$ and R_k, and is small when K_kH_k is close to unity. K_kH_k is close to unity when H_k's rank is close to the

state dimension n and R_k is small (loosely, when many of the states are measured independently and accurately).

If the model is linear, the expressions for K_k, $P_{k|k-1}$ and $P_{k|k}$ depend only on the system model and each other. Therefore, they can be precomputed and stored in memory to save computational power during operation. However, the expression for $\hat{x}_{k|k}$ depends on y_k and therefore the state update equations can not be evaluated before the measurements become available.

The Kalman filter described above is the optimal causal filter when the process noise $\{w_k\}$ and measurement noise $\{v_k\}$ are Gaussian, zero-mean, uncorrelated, and white. Here, optimal means that the Kalman filter minimises any weighted two-norm $\tilde{x}_k^T W_k \tilde{x}_k$ of the estimation error $\tilde{x}_k = x_k - \hat{x}_k$, where W_k is an arbitrary positive definite weight matrix. When the Gaussian assumption fails, the Kalman filter still minimises $\tilde{x}_k^T W_k \tilde{x}_k$ within the class of *linear* causal filters, i.e., the class of filters which outputs a linear combination of past measurements. However, non-linear filters may in this case have a lower estimation error [37].

2.2.4   Multiple model state estimation

A class of problems to which the standard Kalman may not present an acceptable solution is that in which a system jumps at unknown times between a total of M different 'modes' m which exhibit distinct dynamics [38]. $m_k^{(i)}$ signifies that mode $i \in 1, 2, \ldots, M$ is in effect at time k. In each of these different modes, the system dynamics may be described by different models. In this case, one may benefit from using all these different models in an estimate of the system's state, and from estimating the mode that is in effect.

If the mode sequence $\{m_k\}$ is a Markov chain, the system is termed a Markov Jump System (MJS) [39]. In such a Markov chain, the probability $P(m_{k+1}^{(j)})$ of the system assuming mode j at time $k+1$ depends only on the mode $m_k^{(i)}$ that was active at the previous time k: $P(m_{k+1}^{(j)}|m_k^{(i)}) = \pi_{ij,k}$. In other words, the probability $p_{ij,k}$ that the system transitions from mode i to mode j between time steps k and $k+1$ is a fixed function of the time index k and does not depend on e.g. the long-term mode history. Furthermore, if the dynamics of each mode of a MJS are linear, the system is termed a Markov Jump-Linear System (MJLS) [39]. The complete state $\xi = (x, m)$ of such a MJLS consists of both the base state x and the mode m, since both are necessary to characterise the system's internal condition.

A multiple model or multi-model (MM) estimator for a MJLS is based on a number M of different linear state-space models, with one corresponding to each mode m. The i-th model is constituted by the following equations:

$$x_k = F_{k-1}^{(i)} x_{k-1} + G_{k-1}^{(i)} u_{k-1} + w_{k-1}^{(i)}, \tag{2.27a}$$

$$y_k = H_k^{(i)} x_k + v_k^{(i)}, \tag{2.27b}$$

$$w_k^{(i)} \sim \mathcal{N}(0, Q_k^{(i)}), \qquad v_k^{(i)} \sim \mathcal{N}(0, R_k^{(i)}), \tag{2.27c}$$

$$E(w_k^{(i)}(w_j^{(i)})^T) = Q_k^{(i)} \delta_{k-j}, \quad E(v_k^{(i)}(v_j^{(i)})^T) = R_k^{(i)} \delta_{k-j}, \quad E(v_k^{(i)}(w_j^{(i)})^T) = 0. \tag{2.27d}$$

The evolution of the MM algorithm's confidence in the different models is tracked in the form of *mode probabilities*: $\mu_k^{(i)} = P(m_k^{(i)})$ is the probability that model i correctly describes the active mode. A Kalman filter is formulated based on each of the M models, and each filter predicts a certain state evolution and corresponding measurements, one time step ahead. On the basis of how close these predictions are to the real measurements, the mode probabilities are updated to find out which filter best describes the system dynamics.

If the different filters operate autonomously (i.e., they each use their own updated state estimate as the initial condition for the next time step), the MM filter is termed an Autonomous MM (AMM) filter [39]. Such a filter is optimal under the assumption that the system stays in the same mode throughout the experiment, i.e., $\Pi_k = [\pi_{ij,k}]$ equals the identity matrix. However, the AMM approach is not optimal when a finite probability of mode transitions is allowed, because the quality of the measurement prediction of a particular filter depends on its initial condition for that time step. A filter that presently assumes the correct mode, but (due to a recent mode transition in the system) formerly assumed an incorrect mode, may start the present time step with an incorrect initial condition, and produce poor measurement predictions during the first few time steps after the actual transition to its assumed mode.

In order to produce an optimal estimate of the system's complete state, an MM estimator should in principle take into account all possible sequences of mode histories $\{m_k^{(i)}\} = (m_0^{(i)}, \ldots, m_k^{(i)})$. In that case, one assumed mode history will always be correct, and the sequence of filters corresponding to that mode history will produce a minimum-variance estimate. However, in this case, the amount of possible model histories at time k and hence the amount of filter operations that have to be performed at time k grows exponentially with time as M^k. This means the computational requirements quickly become prohibitively large. A number of MM algorithms aim at reducing the amount of filtering operations required by limiting the number of mode history hypotheses that are taken into account. This can be done either by merging similar model sequences together or by pruning unlikely model sequences. Effectively, this limits the amount of mode history hypotheses that the MM filter takes into account, so the different methods are termed Hypothesis Reduction (HR) strategies. Practically, each HR strategy leads to the component filters sharing their state estimates in different ways, which ensures that each filter starts the next time step with a reasonable state estimate. For this reason, such algorithms are termed Cooperating MM (CMM) algorithms [39].

For example, one may decide that only model sequences that end on the mode with the highest probability are likely enough to consider, and prune all other sequences. This can be accomplished by using the posterior state estimate of last step's best filter as an initial condition for the new time step. This is the 1-Best algorithm. Alternatively, all past mode histories can be effectively merged into one by using the fused estimate (average of all filters' estimates with mode probabilities as weights) as an initial condition for the next time step. This is the so-called generalised pseudo-Bayes algorithm of order 1 (GPB1). A third option is to choose a unique mixed estimate for each filter, where weights are based not only on mode probabilities but also the transition probabilities to the mode described by each filter. This is the Interacting Multi-Model (IMM) algorithm. These Hypothesis Reduction (HR) strategies differ in computational complexity and in results, and which HR strategy is the most useful depends on the MM estimation problem [39].

One cycle of the IMM algorithm at time k comprises a number of steps. These steps are described in words below, in formulas in Table 2.1, and depicted as a block diagram in Figure 2.5.

- *Estimate mixing step*: To ensure each filter starts the time step with a reasonable state estimate, a mixed state estimate $\bar{x}_{k-1|k-1}^{(i)}$ is calculated for each filter $i = 1, \ldots, M$. The state estimates $\hat{x}_{k-1|k-1}^{(j)}$ of the previous time step are combined with the mode probabilities $\mu_{k-1|k-1}^{(i)}$ and the transition probabilities Π_k to produce these individualised mixed state estimates $\bar{x}_{k-1|k-1}^{(i)}$. The mixed estimate for filter i is the expected value of x_k, conditioned on all past measurements (as usual) *and* the fact that model i correctly describes the system in the current time step k: $\bar{x}_{k|k}^{(i)} = E(x_k, |\{z_k\}, m_k^{(i)})$.

To this end, the prior mode probabilities $\mu_{k|k-1}^{(i)}$ are first calculated using the (old) posterior

mode probabilities $\mu_{k-1|k-1}^{(i)}$ and transition probabilties Π_k (Eq. 2.28). This can be seen as a time-update for the mode probabilities, similar to the state time-update in the Kalman filter. What follows is the calculation of the mixing weights $\mu_{k|k-1}^{j|i}$ in Eq. 2.29, which determine the relative contributions of the (old) posterior state estimates $\hat{x}_{k-1|k-1}^{(j)}$ of all filters to the individualised mixed estimate $\bar{x}_{k-1|k-1}^{(i)}$ of filter i. These weights can be interpreted as $P(m_{k-1}^{(j)}|\{z_k\}, m_k^{(i)})$, the probability that the system was in mode j at time step $k-1$, assuming that it is now in mode i. In Eqs. 2.30 and 2.31, the mixing weights are used to produce an individualised mixed estimate and estimation error covariance for each filter.

- *Filtering step*: For each model $i = 1, \ldots, M$, the relevant Kalman filter operates on the relevant mixed state estimate $\bar{x}_{k-1|k-1}^{(i)}$ and error covariance $\bar{P}_{k-1|k-1}^{(i)}$. A time-update step (Eq. 2.32) and measurement-update step (Eq. 2.37) are performed in order to produce the updated state $\hat{x}_{k|k}^{(i)}$ and its error covariance $\hat{P}_{k|k}^{(i)}$. To this end, the innovations $z_k^{(i)} = y_k - \hat{y}_k^{(i)}$ (Eq. 2.34) , its covariance $S_k^{(i)} = E\left[z_k^{(i)}(z_k^{(i)})^T\right]$ (Eq. 2.35) , and the Kalman gain $K_k^{(i)}$ (Eq. 2.36) are also calculated.

- *Mode probability update step*: On the basis of how well each filter predicted the current time step's measurements, the mode probabilities are adjusted. For each model $i = 1, \ldots, M$, the prior likelihood $L_k^{(i)}$ that model i assigned to the actual measurement y_k is evaluated as the prior probability density that the innovations has a realisation $z_k^{(i)}$ (Eq. 2.39). This likelihood is higher for a filter that has smaller innovations (i.e. more accurate measurement prediction), compared to the innovations covariance. Because the innovations is a zero-mean Gaussian RV with covariance S_k, the multi-variate Gaussian distribution (Eq. 2.15 for $z_k \sim \mathcal{N}(0, S_k)$) is used.

  On the basis of the likelihoods $L_k^{(i)}$ and the prior mode probabilities $\mu_{k|k-1}^{(i)}$, the posterior mode probabilities $\mu_{k|k}^{(i)}$ are calculated, using Bayes' law for updating beliefs under evidence (Eq. 2.40). The mode probabilties are multiplied by the respective likelihoods, and the results are renormalised to a sum of 1. This can be seen as a measurement update for the mode probabilities, and is conceptually similar to the state measurement update in the Kalman filter.

The filtering step and the probability update step are the same for all MM filters, while the mixing step is unique to the IMM (of course, the prior mode probability update is also used regardless of the HR strategy). In the GBP1 algorithm, the mixing step is not used, being replaced by the estimate fusion step.

Estimate fusion step: An overall state estimate $\hat{x}_{k|k}$ and associated error covariance $\hat{P}_{k|k}$ are calculated, where the posterior mode probabilities are used as weights. These are used as the initial condition for all filters in the next filtering step.

In the 1-Best algorithm, neither the mixing step nor the fusion step is used, and instead a choosing step is used:

Estimate choosing: The model $j \in 1, \ldots, M$ with the highest posterior mode probability $\mu_{k|k}^{(j)}$ is selected, and its state estimate $\hat{x}_{k|k}^{(j)}$ and error covariance $P_{k|k}^{(j)}$ are used by all filters in the next filtering step.

Table 2.1: One time step of the Interacting Multiple Model algorithm [39]. The estimate fusion step is not necessary for filter operation, but allows the filter to report an estimate to the outside world which is not conditioned on which model is correct. In the present notation of sub- and superscripts, $\mu_{k|k-1}^{(i)}$ represents the mode probability μ of filter i at time step k, conditioned on measurements up to and including time step $k-1$.

|		
Estimate mixing (for $i = 1, .., M$)		
Prior mode probability:	$\mu_{k|k-1}^{(i)} = \sum\limits_{j=1}^{M} \pi_{ji,k} \mu_{k-1|k-1}^{(j)}$	(2.28)
Mixing weight (for $j = 1, .., M$):	$\mu_{k|k-1}^{j|i} = \pi_{ij,k} \mu_{k-1|k-1}^{(j)} / \mu_{k|k-1}^{(i)}$	(2.29)
Mixed estimate:	$\bar{x}_{k-1|k-1}^{(i)} = \sum\limits_{j} \hat{x}_{k-1|k-1}^{(j)} \mu_{k|k-1}^{j|i}$	(2.30)
Mixed error covariance:		
$\bar{P}_{k-1|k-1}^{(i)} = \sum\limits_{j} \left[P_{k-1|k-1}^{(j)} + \left(\bar{x}_{k-1|k-1}^{(i)} - \hat{x}_{k-1|k-1}^{(j)} \right) \left(\bar{x}_{k-1|k-1}^{(i)} - \hat{x}_{k-1|k-1}^{(j)} \right)^T \right] \mu_{k|k-1}^{j|i}$		(2.31)
Time-update (for $i = 1, .., M$)		
Prior estimate:	$\hat{x}_{k|k-1}^{(i)} = F_{k-1}^{(i)} \bar{x}_{k-1|k-1}^{(i)} + G_{k-1}^{(i)} u_{k-1}$	(2.32)
Prior error covariance:	$P_{k|k-1}^{(i)} = F_{k-1}^{(i)} \bar{P}_{k-1|k-1}^{(i)} (F_{k-1}^{(i)})^T + Q_k^{(i)}$	(2.33)
Measurement-update (for $i = 1, .., M$)		
Innovations:	$z_k^{(i)} = y_k - H_k^{(i)} \hat{x}_{k|k-1}^{(i)}$	(2.34)
Innovations covariance:	$S_k^{(i)} = H_k^{(i)} P_{k|k-1}^{(i)} (H_k^{(i)})^T + R_k^{(i)}$	(2.35)
Filter gain:	$K_{k|k-1}^{(i)} = P_{k|k-1}^{(i)} (H_k^{(i)})^T (S_k^{(i)})^{-1}$	(2.36)
Posterior estimate:	$\hat{x}_{k|k}^{(i)} = \hat{x}_{k|k-1}^{(i)} + K_k^{(i)} z_k^{(i)}$	(2.37)
Posterior error covariance:	$P_{k|k}^{(i)} = P_{k|k-1}^{(i)} - K_k^{(i)} S_k^{(i)} (K_k^{(i)})^T$	(2.38)
Mode probability update (for $i = 1, .., M$)		
Model likelihood:	$L_k^{(i)} = \dfrac{1}{\sqrt{\det(2\pi S_k^{(i)})}} \exp\left[-\tfrac{1}{2} \left(z_k^{(i)} \right)^T \left(S_k^{(i)} \right)^{-1} z_k^{(i)} \right]$	(2.39)
Posterior mode probability:	$\mu_{k|k}^{(i)} = \dfrac{\mu_{k|k-1}^{(i)} L_k^{(i)}}{\sum\limits_{j} \mu_{k|k-1}^{(j)} L_k^{(j)}}$	(2.40)
Estimate fusion		
Overall estimate:	$\hat{x}_{k|k} = \sum\limits_{i} \hat{x}_{k|k}^{(i)} \mu_{k|k}^{(i)}$	(2.41)
Overall error covariance:	$P_{k|k} = \sum\limits_{i} \left[P_{k|k}^{(i)} + \left(\hat{x}_{k|k} - \hat{x}_{k|k}^{(i)} \right) \left(\hat{x}_{k|k} - \hat{x}_{k|k}^{(i)} \right)^T \right] \mu_{k|k}^{(i)}$	(2.42)
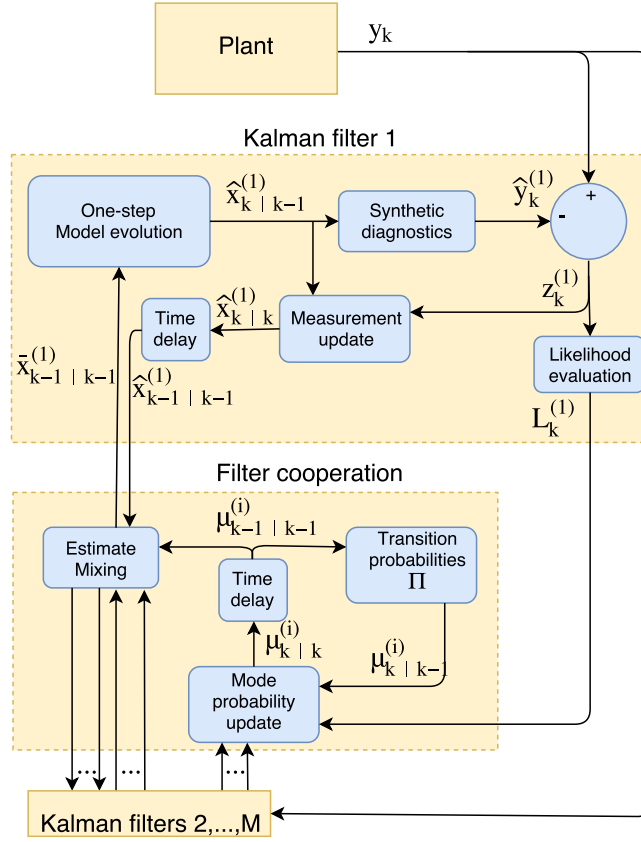
Figure 2.5: Block diagram of an Interacting Multiple Model Kalman filter estimating the state and mode of a real system. One of the M component Kalman filters is depicted in detail. Apart from its usual operation, it evaluates the prior likelihood $L_k^{(i)}$ of the occurence of the real measurement y_k. The likelihoods of all filters are used to update the mode probabilities $\mu_k^{(i)}$ that each filter correctly assumed the mode the system is in. $\mu_k^{(i)}$ evolves autonomously according to the transition probabilities Π_k. $\mu_k^{(i)}$ and Π_k also determine the weights with which the different posterior state estimates $\hat{x}_{k-1|k-1}^{(i)}$ are mixed to obtain the mixed state estimates $\bar{x}_{k-1|k-1}^{(i)}$ on which the model dynamics will operate in the next time step.

2.2.5 Conclusions on state estimation theory

The multi-model filtering theory of Section 2.2.4 provides a framework in which events can be detected from diagnostic data. Multiple Kalman filters update their state estimate on the basis of both their internal models and the measurements. The difference between the internal models is representative of the difference in system dynamics before, after, and during events to be detected, and will cause the different filters to produce different measurement predictions. The models that are used in this work are formulated in Chapter 3. A component filter that produces accurate measurement predictions obtains a higher mode probability, allowing the detection of events by shifts in mode probability, as is done for LHTs, HLTs, and ELMs in Chapter 4. Different Hypothesis Reduction (HR) strategies can be used to ensure that each filter starts each time step with a reasonable initial estimate, and some different HR strategies will be investigated in Chapter 4.

2.3 Diamagnetic loop signal analysis

The multi-model filtering theory of Section 2.2.4 allows one to detect the timing of LHTs, HLTs, and ELMs, on the basis of simple models for these events. However, as stated in Chapter 1, not only the timing of ELMs, but also the plasma energy loss they cause is of interest. This section focuses on the reconstruction of plasma energy losses during ELMs, using the DiaMagnetic Loop (DML) signal.

The diamagnetic flux ϕ_p allows one to calculate the normalised pressure β. Together with information on ELM timing - provided by the event detector - this allows an analysis of how much plasma energy is expelled during an ELM. Section 2.3.1 shows how the DML signal is linked to the plasma pressure. Some useful approximations to this relation are considered in Sections 2.3.2 and 2.3.3, and the relation is rewritten into a normalised form in Section 2.3.4. In our calculations, output of the equilibrium reconstruction code LIUQE is used as input and the results are compared to different LIUQE outputs. Hence, a short overview of LIUQE is given in Section 2.3.5. Finally, Section 2.3.6 shows how the ϕ_p-β relation can be used to actually calculate β. In Section 4.4 of the results, this relation will be used to calculate the energy losses during ELMs.

The DML diagnostic measures the total toroidal magnetic flux ϕ_p produced by the plasma, given by:

$$\phi_p \equiv \int_S B_{\varphi p} dS = \int_S B_\varphi - B_{\varphi v} dS, \tag{2.43}$$

where the plasma's contribution $B_{\varphi p}$ to the toroidal magnetic field is the difference of the total toroidal field B_φ and the vacuum contribution $B_{\varphi v}$ (produced by the toroidal field coils), and S the surface in the poloidal plane enclosed by the diamagnetic loop. Various hardware and software corrections are made to compensate for the vacuum toroidal magnetic field produced by the toroidal field coils, and eddy currents in the conducting vessel wall [40]. A short overview of the DML hardware is given in Section 3.1.3

2.3.1 Relating the diamagnetic loop signal to plasma pressure

The magnetic field produced by the plasma contributes to the total magnetic energy. The plasma's contribution is given by

$$W_M \equiv \int_V \frac{B_\varphi^2 - B_{\varphi v}^2}{2\mu_0} dV, \tag{2.44}$$

with μ_0 the vacuum permeability. This magnetic energy W_M can be linked to the plasma (perpendicular) kinetic energy

$$W_{kin\perp} \equiv \frac{1}{2} \int_V 2p_\perp dV \tag{2.45}$$

(with $p_\perp$ the pressure perpendicular to the magnetic field) through a derivation in magnetohydrodynamics (MHD). The starting point, approach, and result of this derivation is given below, but the interested reader is referred to [41] for the details.

Starting from Maxwell's equations and a Boltzmann equation for each plasma species (i.e., the ideal MHD model), the momentum balance equation can be written compactly as a function of the stress tensor $\mathbb{T}$:

$$\partial \frac{\rho_m \mathbf{u}}{\partial t} + \nabla \cdot \mathbb{T} = \mathbf{0}, \tag{2.46}$$

where ρ_m is the plasma mass density (dominated by the ion mass density), and $\mathbf{u}$ is the mean fluid velocity (again dominated by the ions). Considering an MHD *equilibrium*, the first term with $\frac{\partial}{\partial t}$

in Eq. 2.46 vanishes and the momentum balance reads

$$\nabla \cdot \mathbb{T} = \mathbf{0}, \tag{2.47}$$

where

$$\mathbb{T} \equiv \rho_m \mathbf{u}\mathbf{u} + \left(p_\perp + \frac{B^2}{2\mu_0}\right)\mathbb{I} - \sigma\frac{\mathbf{B}\mathbf{B}}{\mu_0} \quad \text{with} \quad \sigma \equiv 1 - \frac{\mu_0(p_\parallel - p_\perp)}{B^2}, \tag{2.48}$$

where σ is the pressure anisotropy parameter and $\mathbb{I}$ the identity tensor. Since $\nabla \cdot \mathbb{T} = \mathbf{0}$, any of its moments must also be 0. This means that one may take the inner product of $\nabla \cdot \mathbb{T}$ with any vector field $\mathbf{Q}$, integrate the resulting scalar field (with value zero) over any volume V, and arrive at a result of 0. No matter how the integral is rewritten, the result of 0 remains. For simplicity, an axisymmetric V encompassing the plasma volume V_{pl} and an axisymmetric vector field $\mathbf{Q}$ can be taken.

For example, taking $\mathbf{Q}_1 = \hat{\mathbf{R}}$, the unit vector field pointing radially outward from the toroidal revolution axis results in

$$I_1 = \int_V (\nabla \cdot \mathbb{T}) \cdot \hat{\mathbf{R}} dV = 0, \tag{2.49}$$

and likewise, taking $\mathbf{Q}_2 = R\hat{\mathbf{R}} + Z\hat{\mathbf{Z}}$, the position vector relative to the origin (which lies at some random altitude on the toroidal rotation axis) yields

$$I_2 = \int_V (\nabla \cdot \mathbb{T}) \cdot (R\hat{\mathbf{R}} + Z\hat{\mathbf{Z}}) dV = 0. \tag{2.50}$$

Using some vector and tensor algebra relations, these integrals may be rewritten into a simpler form. For a full derivation, see appendix B of [41]. Then, a specific linear combination of I_1 and I_2 can be equated to 0, yielding the pressure balance equation [41]:

$$W_{kin\perp} = \frac{1}{2}[S_1 + \delta_R S_2] - W_{rot,\theta} - W_M, \tag{2.51}$$

where $W_{rot,\theta}$ is the plasma poloidal rotational energy, S_1 and S_2 are the unnormalised Shafranov integrals, and δ_R is a normalised, pressure-weighted radial shift of the plasma. $W_{rot,\theta}$ is given by:

$$W_{rot,\theta} \equiv \int_V \frac{1}{2}\rho_m u_\theta^2 dV, \tag{2.52}$$

with u_θ the poloidal fluid velocity. The Shafranov integrals are related to the flux surface geometry [42]:

$$S_1 \equiv \frac{1}{\mu_0} \oint_{\partial V} \left[\frac{B_p^2}{2}\hat{\mathbf{n}} \cdot \rho\hat{\boldsymbol{\rho}} - (\hat{\mathbf{n}} \cdot \mathbf{B}_p)(\mathbf{B}_p \cdot \rho\hat{\boldsymbol{\rho}})\right] dA \tag{2.53}$$

$$S_2 \equiv \frac{1}{\mu_0} \oint_{\partial V} \left[\frac{B_p^2}{2}\hat{\mathbf{n}} \cdot R'\hat{\mathbf{R}} - (\hat{\mathbf{n}} \cdot \mathbf{B}_p)(\mathbf{B}_p \cdot R'\hat{\mathbf{R}})\right] dA \tag{2.54}$$

Here, $\mathbf{B}_p$ is the poloidal magnetic field, V is an axisymmetric volume encompassing the plasma volume V_{pl}, $\hat{\mathbf{n}}$ is the outward normal vector on a surface element dA on the border ∂V of V, and ρ is the radial coordinate of a polar coordinate system centered around arbitrary major radius R', see also Figure 2.6. Finally, δ_R in the pressure balance (Eq. 2.51) is a normalised, pressure-weighted radial shift of the plasma given by:

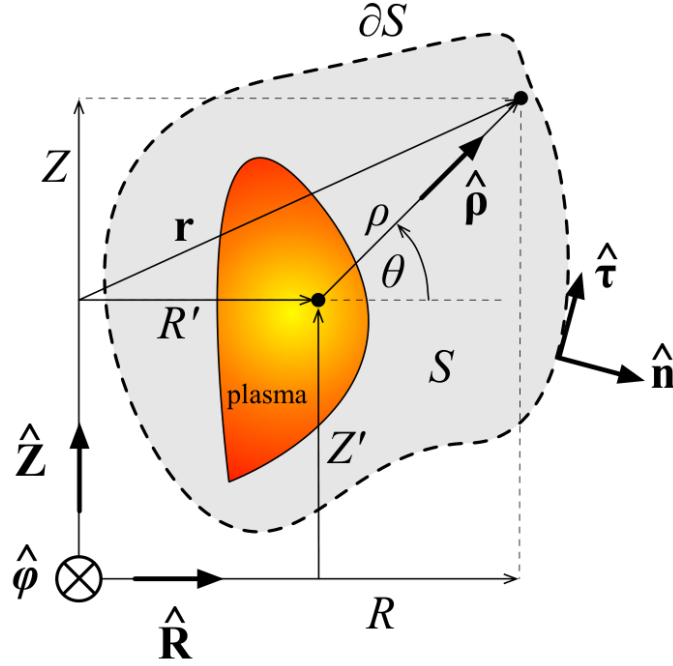$$\delta_R \equiv \frac{R' - R_T}{R'} \quad \text{with} \quad R_T = \frac{\int_S gRdS}{\int_S gdS}, \tag{2.55}$$

Figure 2.6: Geometry of the Shafranov integrals. V is an axisymmetric volume encompassing the plasma bordered by ∂V and S is the poloidal cross-section of V. $\hat{\mathbf{n}}$ is the outward normal vector of ∂V. Any position $\mathbf{r}$ within S can be described by either the cartesian coordinates (R, Z) or the polar coordinates (ρ, θ) centered on the arbitrary (cartesian) coordinates (R', Z'). From [41].

where g is a weighting factor with the unit of pressure:

$$g \equiv \rho u_\varphi^2 + p_\parallel + \frac{B_p^2}{2\mu_0} - \frac{B_\varphi^2(p_\parallel - p_\perp)}{B^2} - \frac{B_\varphi^2 - B_{\varphi v}^2}{2\mu_0} \tag{2.56}$$

Note that Equation 2.51 holds for any choice of V encompassing V_{pl} and (R', Z'). In the derivation of Eq. 2.51, it is assumed that the pressure, rotational velocity, and diamagnetic field vanish outside the plasma: $p_\parallel = p_\perp = 0$, $u_\varphi = 0$, and $B_\varphi = B_{\varphi v}$ in $V \backslash V_{pl}$. Moreover, the pressure anisotropy term is neglected everywhere, since it is always much smaller than either the pressure term (near the plasma centre) or the magnetic pressure term (near the plasma edge).

In both Eqs. 2.53 and 2.54, the second term is zero when the Last Closed Flux Surface (LCFS) is chosen as the contour of S, in which case $\mathbf{B}_p \perp \hat{\mathbf{n}}$ since $\hat{\mathbf{n}} \parallel \nabla\Psi$. However, for a general contour this term may contribute to the integrals.

2.3.2 Static plasma approximation

In the absence of external momentum input, the plasma's rotational energy is negligible compared to its perpendicular kinetic energy; $W_{rot} \approx 0$. Typical toroidal and poloidal flow velocities in the TCV tokamak are $u \simeq 10\,\text{km/s}$ [43]. Hence, for a typical TCV deuterium plasma with major radius $R = 0.88\,\text{m}$, minor radius $a = 0.25\,\text{m}$, elongation $\kappa = 1.5$, and density $n = 1 \times 10^{20}\,\text{m}^{-3}$, the total rotational energy is approximately

$$W_{rot} = \frac{1}{2}nm_D V u^2 \simeq 50\,\text{J}, \tag{2.57}$$

while for a temperature of $1\,\mathrm{keV}$, the plasma perpendicular kinetic energy is:

$$W_{kin\perp} = 2p_\perp V \simeq 50\,\mathrm{kJ}. \tag{2.58}$$

Hence, in the absence of external momentum input like Neutral Beam Injection (NBI), the fluid rotational energy is typically three orders of magnitude smaller than the perpendicular kinetic energy and can be safely neglected in this analysis.

2.3.3 Small diamagnetism approximation

During standard tokamak operation, the toroidal flux produced by the TCV plasma is typically much lower than the vacuum toroidal flux. At the highest plasma current, $I_p = 1\,\mathrm{MA}$, the diamagnetic flux is of the order $\phi_p \simeq 40\,\mathrm{mWb}$, while the toroidal field coils produce a vacuum magnetic flux of around $\phi_v \simeq 2\,\mathrm{Wb}$ [41]. This motivates the small diamagnetism approximation:

$$|B_{\varphi pl}| = |B_\varphi - B_{\varphi v}| \ll |B_{\varphi v}|. \tag{2.59}$$

Under this approximation, W_M as defined in Eq. 2.44 can be rewritten to

$$\begin{aligned} W_M &= \frac{B_\varphi^2 - B_{\varphi v}^2}{2\mu_0} dV = \int_V \frac{(B_\varphi + B_{\varphi v})(B_\varphi - B_{\varphi v})}{2\mu_0} dV \simeq \int_V \frac{2B_{\varphi v}B_{\varphi p}}{2\mu_0} dV \\ &= \int_S \frac{2B_{\varphi v}B_{\varphi vp}}{2\mu_0} 2\pi R\, dS = \frac{2\pi}{\mu_0} RB_{\varphi v} \int_S B_{\varphi p} dS = \frac{2\pi}{\mu_0} RB_{\varphi v}\phi_p, \end{aligned} \tag{2.60}$$

where we sequentially use the notable product $a^2 - b^2 = (a+b)(a-b)$, the approximation $B_\varphi + B_{\varphi v} \approx B_{\varphi v} + B_{\varphi v}$ and the fact that $RB_{\varphi v}$ is constant. The result is that the magnetic energy can be directly related to the measured diamagnetic flux.

2.3.4 Normalisation of the pressure balance

In the limit of a static plasma, small diamagnetism, and isotropic pressure, the pressure balance (Eq. 2.51), simplifies to

$$W_{kin\perp} = \frac{1}{2}[S_1 + \delta_R S_2] - \frac{2\pi}{\mu_0} RB_{\varphi v}\phi_p, \tag{2.61}$$

while Eq. 2.56 simplifies to

$$g = p + \frac{B_p^2}{\mu_0}. \tag{2.62}$$

The kinetic energy is often normalised to obtain the normalised pressure or 'poloidal beta',

$$\beta_{p\perp} \equiv \frac{W_{kin\perp}}{\bar{B}_p^2/2\mu_0}, \tag{2.63}$$

with $\bar{B}_p$ the average poloidal field over the contour of S. However, in in the equilibrium reconstruction code TCV LIUQE (see Section 2.3.5 and [44]), this normalisation factor is evaluated in the cylindrical limit,

$$\beta_{p\perp} = \frac{W_{kin\perp}}{\mu_0 R_0 I_p^2/4}. \tag{2.64}$$

In this work, the cylindrical normalisation factor is also used. Dividing all the terms in Eq. 2.61 by the same normalisation factor yields the equivalent relation

$$\beta_{p\perp} = \frac{1}{2}[s_1 + \delta_R s_2] - \frac{8\pi B_{\varphi v}}{\mu_0^2 I_p^2}\phi_p = s - \frac{\phi_p}{\phi_0}, \tag{2.65}$$

where $s = \frac{1}{2}[s_1 + \delta_R s_2]$ is the shape term, s_1 and s_2 are the normalised Shafranov integrals and $\phi_0 = \frac{\mu_0^2 I_p^2}{8\pi B_{\varphi v}}$ is the diamagnetic flux corresponding to $\beta = 0$ (for a cylindrical plasma).

Note that if a cylindrical approximation is used for the Shafranov integrals as well as the normalisation factor, $s_1 = 2$ and $\delta_R = 0$, meaning s is simply 1. However, in the presence of plasma shaping and/or the outward radial shift due to finite-β effects, this approximation introduces an error of typically 10% in β_p, which is unacceptably large. Using the simple normalisation factor may have a similar effect on β_p, but this procedure differs in that it simply alters the definition of what β_p is being calculated; in principle, the normalisation factor is arbitrary.

2.3.5 Overview of the LIUQE equilibrium reconstruction code

In order to calculate the Shafranov integrals, knowledge of the poloidal magnetic field $\mathbf{B}_p$ is required. Since no magnetic measurements are available inside the vacuum vessel, use is made of the reconstructed magnetic fields given by the equilibrium reconstruction code LIUQE. A short overview of LIUQE is given in this section. See also [44].

Equilibrium reconstruction consists of identifying a plasma current and pressure distribution, along with a flux surface geometry, that satisfy the MHD equations and closely match experimental measurements. Such an MHD equilibrium is described by the pressure balance and Maxwell's equations. With the assumptions of axisymmetry, isotropic pressure, and no net fluid velocity, these read:

$$\mathbf{j} \times \mathbf{B} = \nabla p, \tag{2.66a}$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{j}, \tag{2.66b}$$

$$\nabla \cdot \mathbf{B} = 0. \tag{2.66c}$$

The magnetic field B can be written in cylindrical coordinates R, φ, z as:

$$\mathbf{B} = \frac{1}{2\pi R} \frac{\partial \Psi}{\partial z} \nabla R + \frac{1}{2\pi R} \frac{\partial \Psi}{\partial R} \nabla z + T \nabla \varphi, \tag{2.67}$$

where $\Psi(R, z)$ is the poloidal magnetic flux which determines the flux surface geometry - contours of constant Ψ are the flux surfaces - and $T = R B_\varphi$ describes the toroidal magnetic field. Using Eq. 2.67 and some vector algebra, Eqs. 2.66 can be combined into the Grad-Shafranov equation (GSE) [44]:

$$\nabla^* \Psi = -4\pi^2 \mu_0 R (R p') \frac{T T'}{\mu_0 R}), \tag{2.68}$$

where the elliptical operator is defined as $\nabla^* \equiv R \frac{\partial}{\partial R} \frac{\partial}{R \partial R} + \frac{\partial^2}{\partial z^2}$. p and T are functions of Ψ only, so the accents indicate the derivative with respect to Ψ.

As an equilibrium reconstruction code, LIUQE's goal is to find the poloidal flux $\Psi(R, z)$ and the functions $p(\Psi)$ and $T(\Psi)$ which satisfy equation 2.68 and which best reproduce the experimental measurements of physical parameters related to the quantities behind these functions. Measurements that are used to constrain LIUQE results include active coil currents, magnetic flux loops, poloidal magnetic probes, and the diamagnetic loop [44]. The solution strategy in LIUQE is as follows:

0 Parametrise p' and TT' as a linear combination of selected base functions, with coefficients a_g.

1 Start with an estimate of the coefficients a_g (i.e., an initial guess of p' and TT'), and solve the GSE (Eq. 2.68) for $\Psi(R, z)$.

2 Given this solution for Ψ, determine the coefficients a_g that yield the functions p' and TT' which most closely match the experimental measurements.

Iterate between step 1 and 2 until the solutions converge.

2.3.6 Solving the pressure balance for normalised pressure

From the previous section, it may be noted that the functions $p'(\Psi)$ and $\Psi(R, z)$ contain all the information needed to calculate the normalised pressure β_p. Indeed, a reconstructed β_p is available as LIUQE output data. Still, Eq. 2.65 is relevant because it allows us to calculate β_p semi-independently from LIUQE, which can then be compared to the reconstructed β_p value.

In order to link the diamagnetic flux to β_p, the plasma shape term s in Eq. 2.65 has to be evaluated. Using the poloidal flux Ψ calculated by LIUQE, the poloidal field $\mathbf{B}_p$ can be determined using $\mathbf{B}_p = \nabla\Psi \times \hat{\boldsymbol{\varphi}}$. Both Shafranov integrals can then be calculated. One difficulty is that the weighting factor g used to calculate the radial shift δ_R depends on p, which is unavailable in the LIUQE datasets used in this work and contributes to an integral spanning at least the whole plasma surface. Assuming p does not have a large contribution to the integral for δ_R - i.e., the magnetic field structure alone accurately reflects the radial shift of the plasma- this difficulty can be circumvented using an iterative strategy.

The iterative strategy works as follows: a fixed shape for the pressure profile as a function of the poloidal flux coordinate is assumed, which is multiplied by an estimate of the central pressure p_0. This fixed shape is chosen to be parabolic in the poloidal flux coordinate:

$$
\begin{aligned}
p &= p_0 \left[1 - (1 - \Psi/\Psi_0)^2 \right] && \text{inside the LCFS} \\
p &= 0 && \text{outside the LCFS}
\end{aligned}
\tag{2.69}
$$

with Ψ_0 the (maximum) poloidal flux at the magnetic axis and $\Psi_a = 0$ at the edge. This roughly corresponds to parabolic temperature and density profiles in the spatial radial coordinate. By integrating the pressure profile shape over the plasma volume, one can obtain the constant V' (with the unit of volume) which links p_0 to the total kinetic energy via $W_{kin} = V'p_0$ and hence to the normalised pressure via $\beta_p = \frac{V'}{N}p_0$ (with N the cylindrical normalisation factor).

The strategy is then to start with an estimate of p_0, e.g. $p_0 = 0$, iteratively evaluate β_p and use it to update the estimate for p_0.

Since Equation 2.65 holds when the integrals are taken over any volume encompassing the plasma, it is useful to integrate over the plasma volume itself, as depicted in Figure 2.7. The last closed flux surface (LCFS) contour is specified by LIUQE, as a series of ~ 50 pairs of (R, z)-coordinates on the LCFS.

LIUQE also specifies the poloidal flux on the (R, z)-grid shown in Figure 2.7. Using the numerical gradient, $\mathbf{B}_p = \nabla\Psi \times \hat{\boldsymbol{\varphi}}$ can be calculated on this grid. A first-order numerical integration scheme is used in the integrals, meaning that $\mathbf{B}_p$ is interpolated on the midpoints between the specified LCFS poins, and multiplied by the distances Δl between the LCFS points. ρ and $\hat{\mathbf{n}}$ are also evaluated at the midpoints.
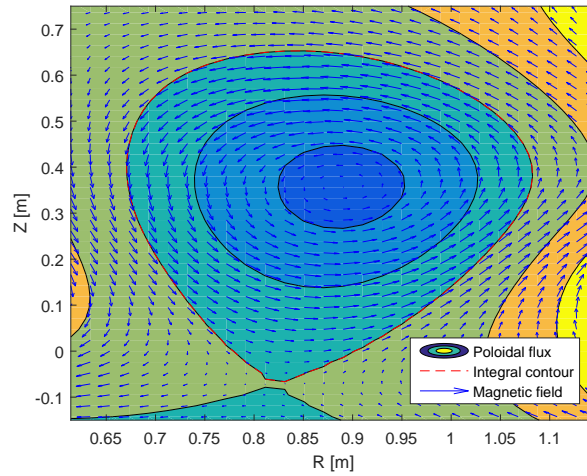
Figure 2.7: Geometry for determining the Shafranov integrals over the last closed flux surface.

2.3.7 Conclusions on diamagnetic loop theory

A relation between the measured plasma diamagnetic flux ϕ_p and the normalised pressure β_p is derived and simplified to the form of Eq. 2.65. In this form, the ϕ_p-β_p relation depends only on the magnetic field strucure $\mathbf{B_p}$ (through the Shafranov integrals s_1 and s_2), the plasma pressure p (through the pressure-weighted radial shift δ_R), and the plasma current I_p. As the equilibrium reconstruction data available for this work do not contain information on the plasma pressure, an iterative strategy is formulated which allows the calculation of β_p when only ϕ_p, $\mathbf{B_p}$, and I_p are given.

This relation is used to calculate β_p from TCV discharge data in Section 4.4, and the results are compared to the β_p obtained from equilibrium reconstruction. In principle, those results can be combined with the ELM timing specified by the event detector to obtain the plasma energy losses during ELMs. This is also done in Section 4.4, but the equilibrium reconstruction results for β_p are used directly in a way which decreases the computational load.

Chapter 3

Models for measurement signal features during events

In this chapter, models are formulated that reproduce typical features seen in diagnostic signals during transitions between L and H mode and during ELMs. These models together form a Markov jump-linear system of the plasma for the purpose of event detection. A multi-model filter for this jump-linear system will be used to detect events in Chapter 4. Section 3.1 discusses what features are to be detected, and the models themselves are formulated in Section 3.2.

3.1 Features to be detected

This section discusses the signal features that are typically observed during L-H transitions, H-L transitions, and ELMs, on which the detector's models are based. Sections 3.1.1 through 3.1.4 discuss features in the photodiode signal which measures the H_α line intensity, the interferometer signal which measures density, the diamagnetic loop signal, and the X-ray electron temperature signal, respectively.

3.1.1 Features in the photodiode signal

The L-H transition (LHT) is most prominently visible by measuring the intensity of H_α line radiation coming from the plasma edge. On TCV, the intensity of different spectral lines is measured using PhotoDiodes (PDs) equipped with different wavelength filters, at a sample rate of 50 kHz. An overview of the viewing lines for the different TCV PDs is given in Figure 3.1. PDs 1-4, 7, and 8 measure the H_α intensity at a wavelength of 656.3 nm over various lines of sight through the plasma. PD 9 observes the H_α intensity at the Neutral Beam duct. PDs 11-18 measure various other line intensities over the vertical line of sight. PDs 5, 6 and 10 are currently not installed.

The PD signal corresponding to a vertical line of sight was found to most clearly show the occurence of an LHT, and will be denoted the H_α signal from this point on. In ideal cases, the H_α signal is at an approximately constant level before an LHT, and suddenly drops by about 50% during the LHT. For example, as in shot 45105 on the TCV tokamak, shown in Figure 3.2
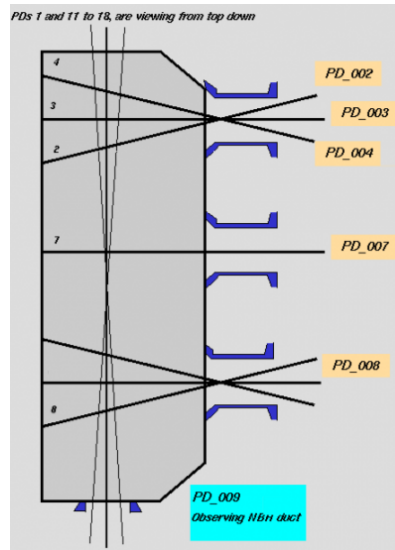
Figure 3.1: Poloidal cross-section of the TCV vessel, showing the view lines for the installed photodiodes. Most important to this work is PD 1, which measures the H_α intensity over a vertical line of sight. Image by Umar Sheikh.

at 1.34 s. The same figure shows that the appearance of ELMs (starting at 1.35 s)is also clearly visible in the H_α signal as sudden spikes, followed by a relaxation to the original baseline level.

During L-mode, ELM-like spikes in the H_α signal are sometimes observed shortly before an L-H transition. These likely signal the occurence of some minor event that causes a short burst of recycled particles. After the intial burst, the H_α signal drops like in an L-H transition, but to the same baseline as before the spike. Figure 3.3 shows two such spikes occuring before an LHT. In some observed discharges, LHTs coincide with the H_α spikes. A transition coinciding with a spike is shown in Figure 3.4.

Some further distinctions can be made between the different transitions. For one, the time-scale at which the H_α signal reaches its new baseline varies between discharges. For example, in Figures 3.3 and 3.4, the H_α signal reaches its new baseline significantly faster than in Figure 3.2. In other words, the latter H_α signature contains a slow tail. Also, the level with which the signal baseline drops varies from almost unnoticable values of $\sim 5\%$ to a well-visible $\sim 50\%$.

The H_α signatures of all LHTs from the test dataset were classified by visual inspection ofthe data, according to the three aforementioned criteria: whether or not they start with a spike, whether they end in a tail, and the approximate relative drop in signal baseline. This classification is presented in Table 3.1.

The H-L transitions (HLTs) are, conversely, visible as an sudden increase in the baseline of the H_α signal. In the case of a transition from an ELMy H-mode to L-mode, this often means that the H_α signal peaks as in a regular ELM, but the signal then relaxes to the higher baseline, as shown in Figure 3.5.

In shots where the plasma stayed in H-mode until the current ramp-down phase at the end of the shot, the expected increase in H_α intensity is obscured by other effects on the H_α signal and sometimes hardly visible, as in Figure 3.6. Other large effects on the H_α during the ramp-down phase include plasma shaping, making it limited and decreasing its cross-sectional area, and cooling leading to eventual recombination of the whole plasma.
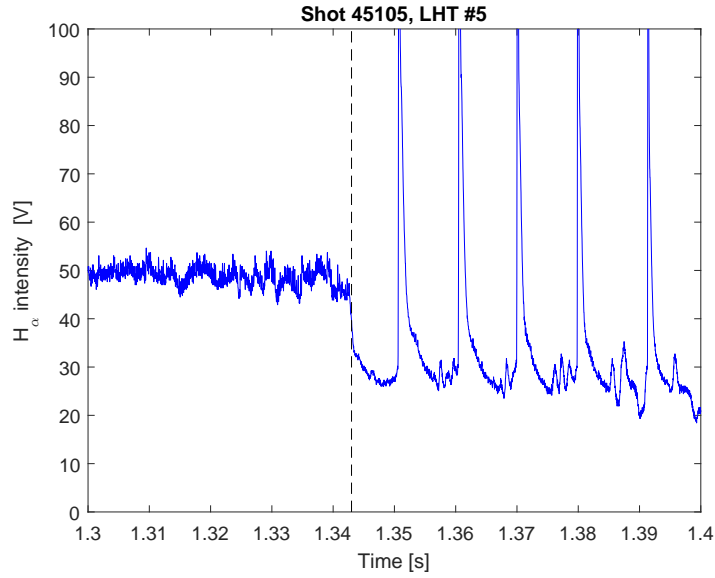
Figure 3.2: An ideal LHT H_α signature at $1.34\,\text{s}$ (black dashed line), showing a significant sudden drop from an approximately constant level. After the LHT, the signal relaxes relatively slowly to its new baseline, in about $5\,\text{ms}$.
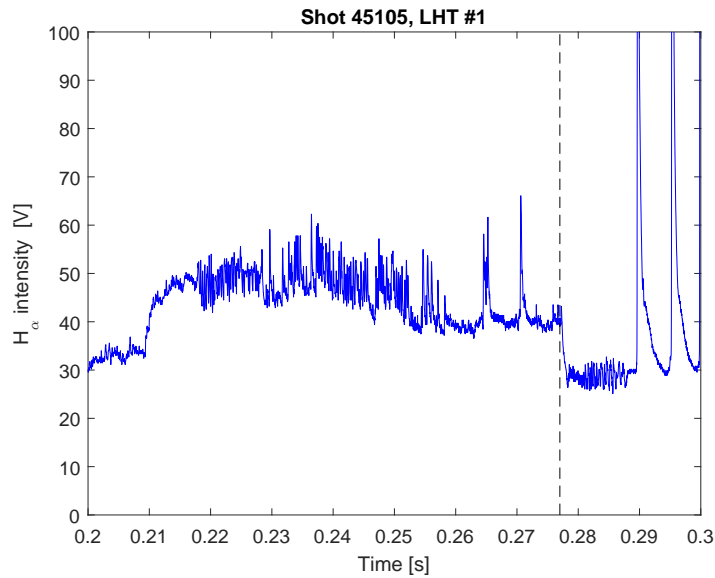

Figure 3.3: The H_α signature of an LHT at $0.28\,\text{s}$ preceded by two spikes at 0.26 and $0.27\,\text{s}$.
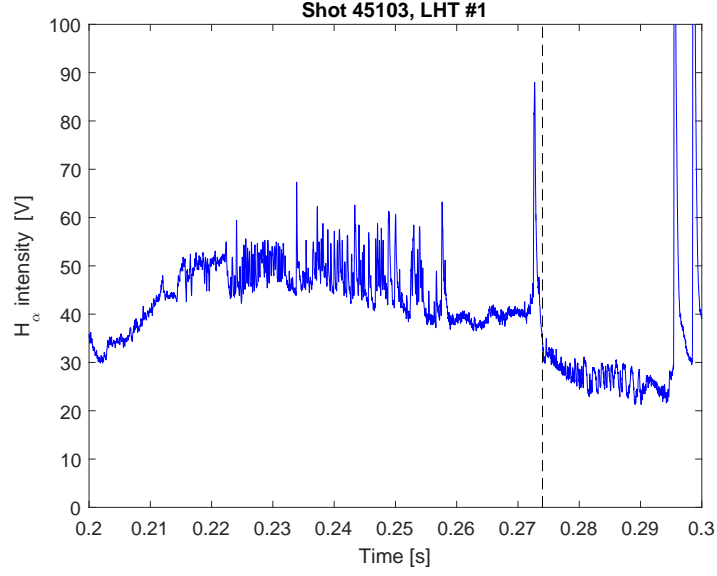
Figure 3.4: The H_α signature of an LHT at $0.27\,$s coinciding with an H_α spike.

Table 3.1: Classification of H_α signatures of L-H transitions in the training dataset on the basis of the presence of spikes and tails, and signal drop.

| Shot	First LHT Spike?	Tail?	Drop	Second LHT Spike?	Tail?	Drop	Third LHT Spike?	Tail?	Drop	Fourth LHT Spike?	Tail?	Drop
39872	No	Yes	60%									
43454	No	Yes	50%									
45103	Yes	No	25%	Yes	Yes	40%						
45104	Yes	No	10%	No	Yes	35%	No	Yes	40%			
45105	No	No	30%	No	Yes	35%	No	No	35%	No	Yes	40%
45106	Yes	Yes	35%	No	No	30%						
45107	Yes	Yes	20%	No	Yes	35%						
45108	Yes	No	5%	No	Yes	35%						
45109	Yes	No	25%	No	Yes	40%	No	No	35%			
47007	Yes	Yes	40%	Yes	Yes	50%						
47008	Yes	No	40%	Yes	No	50%						
47306	No	No	40%									
47307	No	No	30%									
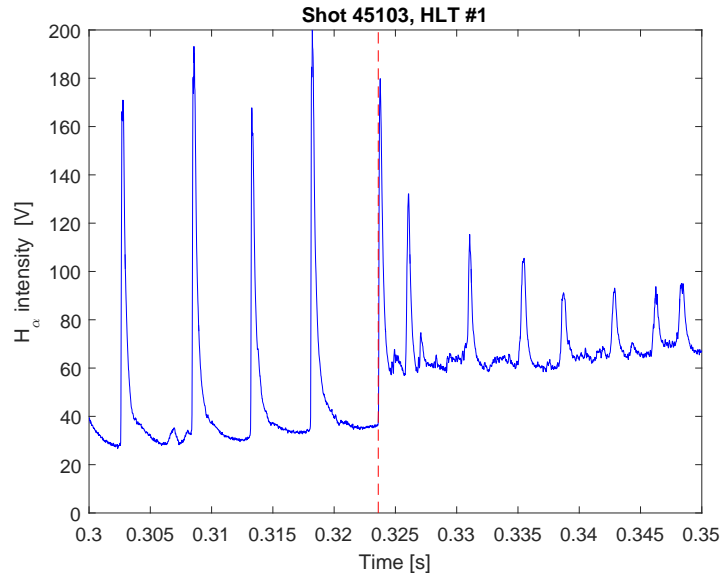
Figure 3.5: H_α time-trace of an HLT at $0.325\,$s (red dashed line), showing relaxation to a higher baseline after an ELM.
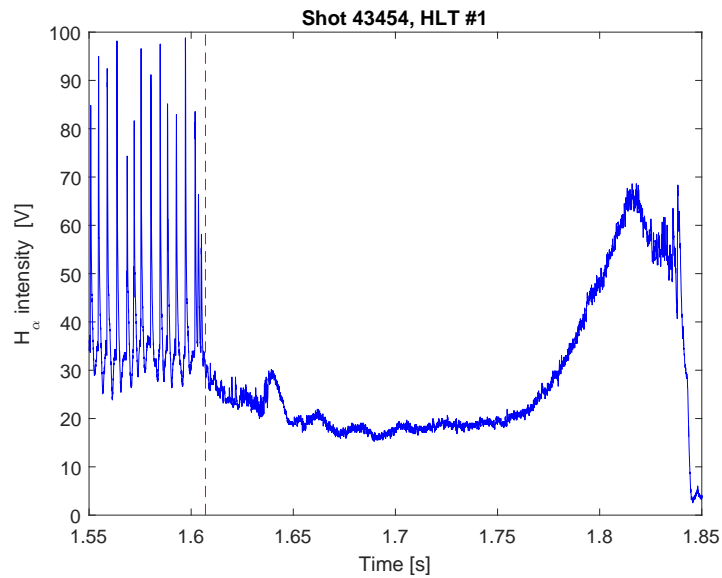

Figure 3.6: H_α time-trace of an HLT likely at $1.6\,$s. This HLT can only be seen as an end to the ELMs, but the expected increase in H_α intensity is not seen. The H_α increase is possibly obscured by other events in the current ramp-down phase.
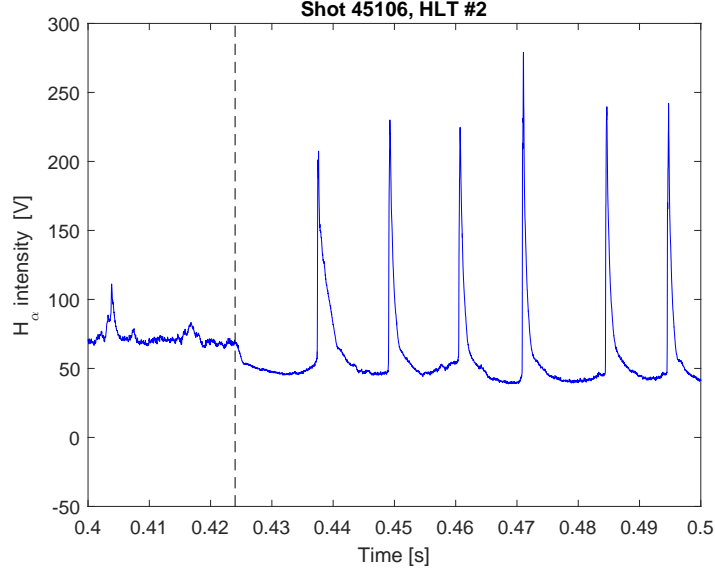
Figure 3.7: H_α time-trace of an LHT at $0.42\,\mathrm{s}$ and subsequent ELMs in shot 45106. During the first ELM, the H_α signal relaxes more slowly than during the other ELMs.

ELMs are visible on the H_α signal as well, as sudden peaks followed by a gradual relaxation to the same baseline value. For ELMs in the training dataset, this relaxation takes place on a time scale similar to that of an LHT. However, the first ELM in a H-mode period often shows a significantly slower signal relaxation, as seen in Figure 3.7 for shot 45106.

Summarizing, the H_α features which the model for the plasma should reproduce are changes in the baseline of the signal at LHTs and HLTs, where the baseline in H-mode is about 30% lower than in L-mode, and peaks at ELMs, where the signal approximately doubles within $0.3\,\mathrm{ms}$.

3.1.2   Features in the interferometer signal

LHTs and HLTs are also measurable as a sudden increase in the time-derivative of plasma density. In TCV, the line-integrated electron density $\int n_e dl$ is measured using an interferometer set-up, which measures optical pathlength differences caused by the refractive index of the plasma. For a wavelength in the sub-mm (i.e., far infra-red or FIR) range, the refractive index is a simple function of n_e. On TCV, a continuous-wave laser with a wavelength of $214\,\mu$m is used. The interferometer consists of 14 parallel laser chords in the vertical direction and a multi-element detector. The radial position of the 14 chords varies between $0.6910\,\mathrm{m}$ and $1.0620\,\mathrm{m}$, and is shown schematically in Figure 3.8. The sample rate for the FIR measurement is $20\,\mathrm{kHz}$, and the line-integrated density per fringe is $1.0360 \times 10^{19}\,\mathrm{m}^-2$ (amount of electrons per unit volume integrated over length).

Due to the improved confinement in H-mode, an LHT is seen as a sudden increase in the time-derivative of the signals of all FIR chords. To minimise noise, the average over all the FIR chords is used, except the outer two FIR chords on the high-field side, which sometimes suffer from fringe jumps. This average will be termed the FIR signal from here on out.

Figure 3.9 shows how different events are typically visible on the FIR signal in TCV discharge # 45108. In this shot, the plasma has become diverted at $0.2\,\mathrm{s}$, and at $0.25\,\mathrm{s}$ the density decreasing because the divertor is retaining a lot of impacting particles. In other words, the net recycling rate
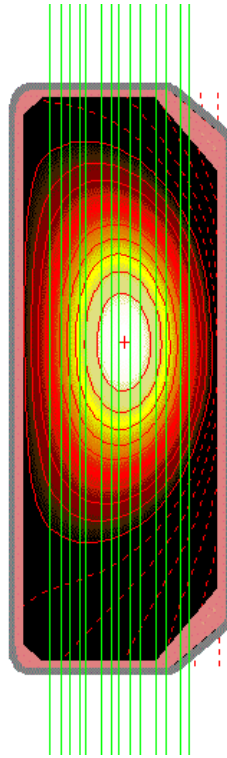
Figure 3.8: Poloidal view of the TCV vessel with the interferometer chords in green.

has decreased relative to the limiter configuration. At the LHT at 0.285 s, confinement suddenly improves and the density starts increasing. When the ELMs start, the density stops increasing on average. The ELMs themselves are visible as positive peaks. The increase is particularly large for those chords whose line of sight intersects the plasma boundary. This may be due to the release of retained particles from the divertor when it is hit by ELMs. Still, the net result of each ELM is a decrease in the FIR signal, and the density increases in between ELMs. At the HLT occuring at 0.32 s, confinement suddenly deteriorates, producing a steady decrease in density, visible as a negative time derivative of the FIR signal.

Especially before the LHT and after the HLT, it can be seen that the signal is disturbed by an oscillation with a frequency around 60 Hz. This oscillation is visible on all channels, and likely caused by a mechanical oscillation of the interferometry set-up, which is specific to the TCV.

Table 3.2 shows the amount by which the time derivative of the FIR signal changes during each transition, which is calculated using linear fits in the 10 ms just before the transitions and in the 10 ms just after. This change is of the order $5 \times 10^{20}\,\mathrm{m}^{-2}\mathrm{s}^{-1}$, though there is some variation between the different shots. In shots where the plasma stayed in H-mode until current ramp-down phase, the final HLT was not always recognisable from the density signals, in some cases leading to large outliers in the table.

Summarizing, the FIR signal features which the plasma model should reproduce are changes in the time-derivative at LHTs and HLTs, where the time-derivative is $5 \times 10^{20}\,\mathrm{m}^{-2}\mathrm{s}^{-1}$ higher in H-mode.
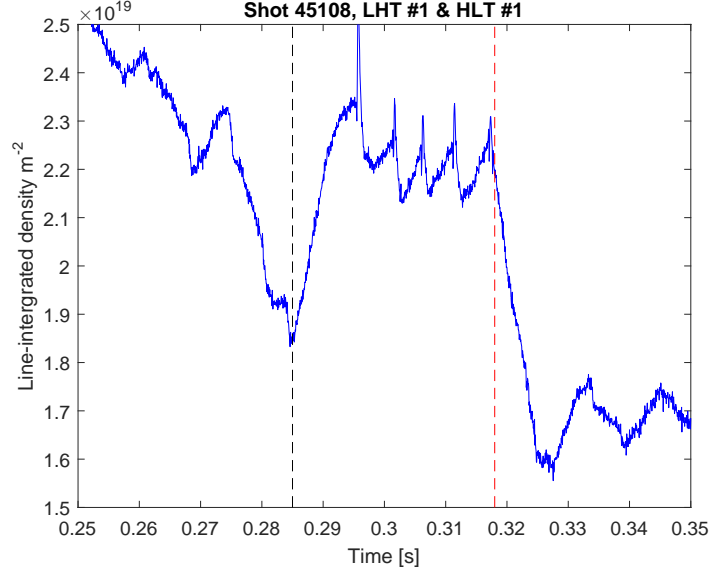
Figure 3.9: FIR time-trace of an LHT at 0.285 s, followed by a few ELMs and an HLT at 0.32 s. At the LHT, the time derivative suddenly increases, and at the HLT it suddenly decreases. ELMs cause the increase in signal after the LHT to stop, and also produce short spikes in the FIR signal.

Table 3.2: Increase (or decrease) of the time derivative of the density signal at the different LHTs and HLTs in the training dataset, from linear fits over time windows of 10 ms.

	1st LHT	1st HLT	2nd LHT	2nd HLT	3rd LHT	3rd HLT	4th LHT	4th HLT
Shot	Increase in density time derivative $[10^{20}\text{m}^{-2}\text{s}^{-1}]$							
39872	1.9	-2.7						
43454	2.3	2.3						
45103	1.5	-2.8	1.7	-3.7				
45104	8.3	-6.8	-1.4	-0.4	-0.4	-4.0		
45105	2.4	1.9	1.7	-5.0	4.6	-2.2	0.5	-2.3
45106	2.4	-4.1	5.2	1.0				
45107	4.9	-6.6	0.4	-1.7				
45108	9.4	-6.7	-0.9	3.7				
45109	1.1	-6.3	1.2	-1.4	3.6	0.1		
47007	5.3	0.5	2.1	1.5				
47008	12	-1.8	2.2	-4.0				
47306	5.2	4.6						
47307	8.6	82						

3.1.3   Features in the diamagnetic loop signal

The Diamagnetic Loop (DML) signal measures the toroidal diamagnetic flux produced by the plasma, which is sensitive to the (normalised) plasma pressure as explained in Section 2.3. The diamagnetic measurement signal comprises 4 magnetic loops, the location of which is shown in Figure 3.10. The diamagnetic loop D makes a single turn around the vacuum vessel and measures the total toroidal flux from the coils, the plasma, and vessel image currents. The compensation loop C consists of 80 turns and is used to compensate for the vacuum toroidal flux produced by the coils. Loop A makes a single turn tightly around the vacuum vessel and compensates for vessel currents. Finally, loop B also consists of 80 turns and is sensitive to the current diffusion in the toroidal coil [40]. The DML measurement runs at a sample rate of 10 kHz.

Similar to the FIR signal, the DML signal shows a sudden increase in its time derivative during an LHT and a sudden decrease during an HLT, as shown in Figure 3.11. The DML signal is also affected by the plasma current; as Figure 3.11 shows, the DML signal increases in the negative direction during the current ramp-up phase of the shot (because a negative current is ramped up). It goes back to 0 in the current ramp-down phase at the end of the shot (not shown). Moreover, ELMs affect the DML signal. The sudden outburst of energy and particles causes a sharp dip in the signal, which rises in between ELMs.

In Table 3.3, the amount by which the DML signal increases during LHTs and decreases during HLTs is tabulated for all shots in the test dataset. As with the FIR signal, this change is calculated using linear fits to the signal data in time windows of 10 ms just before and just after the transitions. The table shows that the DML signal's time derivative increases by typically 0.05 Wb/s during a transition. Summarizing, the DML signal features which the model should reproduce are changes in the time-derivative at LHTs and HLTs, where the time-derivative in H-mode is 0.05 Wb/s lower.
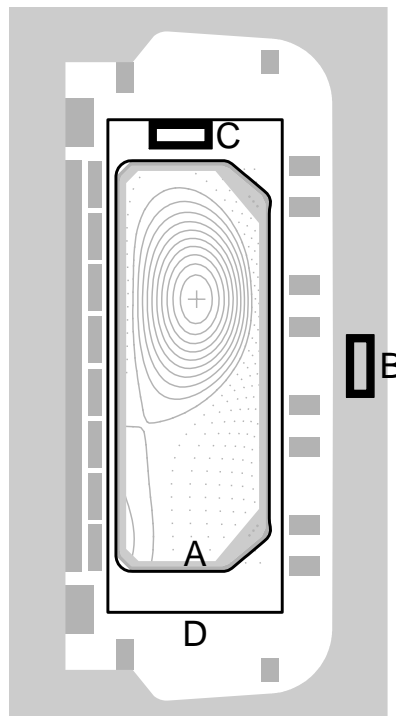

Figure 3.10: Poloidal cross-section of the TCV vessel showing the location of the 4 loops A-D which form the diamagnetic measurement system [40].
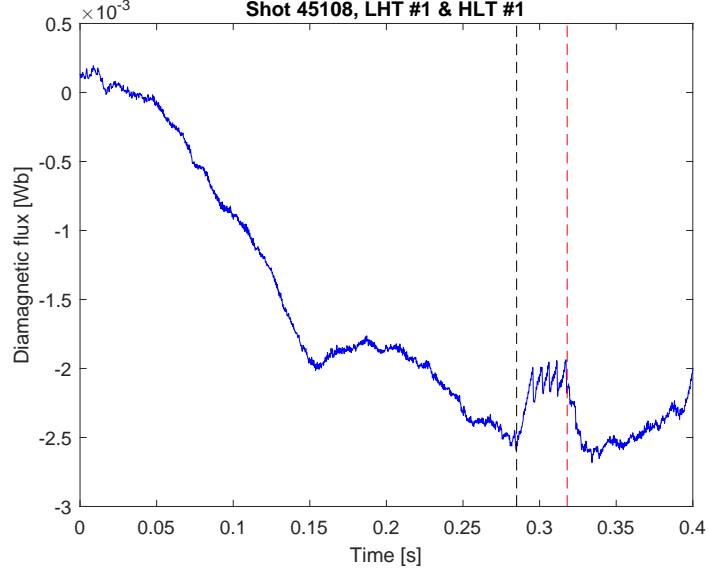
Figure 3.11: DML time-trace of an LHT at 0.28 s, followed a few ELMs and an HLT at 0.32 s. At the LHT, the time derivative suddenly increases, and at the HLT it suddenly decreases. Moreover, the DML signal increases in the negative direction during the first 0.2 s of the shot, while the current is being ramped.

Table 3.3: Increase (or decrease) of the time derivative of the diamagnetic loop signal at the different LHTs and HLTs in the training dataset, from linear fits over time windows of 10 ms.

	1st LHT	1st HLT	2nd LHT	2nd HLT	3rd LHT	3rd HLT	4th LHT	4th HLT
Shot	Increase in time derivative of the diagmagnetic flux[mWb/s]							
39872	10	-18						
43454	24	14						
45103	27	-50	36	-31				
45104	68	-55	51	-27	37	19		
45105	37	-59	43	-6	23	-3	22	-7
45106	23	-46	48	11				
45107	46	-56	48	3				
45108	59	-49	58	3				
45109	34	-38	56	3	9	8		
47007	2	2	-4	-18				
47008	-5	11	10	-20				
47306	18	14						
47307	14	-425						

3.1.4 Features in the X-ray electron temperature signal

The TCV XTe signal is an estimate of the central line-integrated electron temperature using the line intensity ratio between different X-ray measurements measured with different thickness Beryllium filters. Under the assumption of negligible high-Z impurities and a Maxwellian velocity distribution, this ratio can be related to the slope of the electron distribution in velocity space, from which the temperature can be calculated [45]. As the electron temperature rises following an LHT due to increased confinement, the XTe signal is expected to show a sudden increase (decrease) in its time-derivative at an LHT (HLT), similar to the FIR and DML signals.

In Figure 3.12, the XTe signal features during an LHT are compared with the DML signal features. The XTe signal shows some increase in time derivate, but this only happens about 50 ms after the visual detection of the transition (based on the H_α signal), compared to about 2 ms for the DML signal. This delay compared to the other signals is representative for all observed shots, and may be due to the fact that the central temperature is initially unaffected by the LHT, which starts as an increase in edge temperature that only spreads to the core on a time scale comparable to the energy confinement time. This long delay means that the LHT can already be detected on the basis of the other signals before the XTe signal shows any change. ELMs (starting at 0.54 s) are visible as a small decrease in the XTe signal, but are obscured by other sudden decreases in the signal. These other sudden decreases in central temperature are likely caused by the sawtooth instability.

The XTe signal responds too late to an LHT to be useful for the purpose of real-time LHT detection. Moreover, the XTe signal shows approximately the same behaviour for ELMs and sawtooth crashes, making it ineffective in detecting ELMs. For this reason, the XTe signal in not used in the detector.
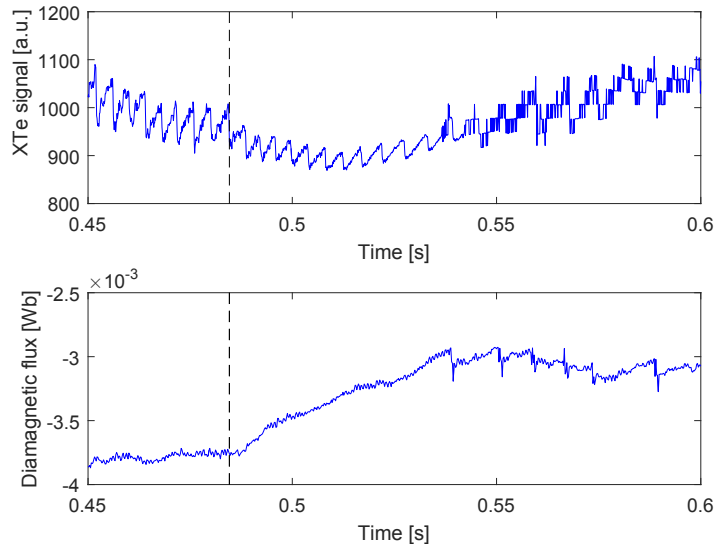

Figure 3.12: XTe (top) and DML (bottom) time-trace of an LHT at 0.485 s and subsequent ELMs from 0.54 s. Both signals show an increased time-derivative after the LHT, but the XTe signal responds much later to the LHT. ELMs cause a sudden decrease in both signals, but affect the DML signal more strongly than the XTe signal. Sawtooth crashes also cause sudden drops in the XTe signal.

3.2 Models for underlying plasma parameters

In this section, models are formulated for the dynamics of the plasma parameters that underlie the measured quantities discussed in Section 3.1. The modeling focuses on the dynamics during the LHT, HLT and ELM events and the differences in dynamics prior to and after these events occur. The combination of these models is a Markov Jump-Linear System (MJLS) of the plasma for the purpose of event detection. The transition probabilities for this MJLS are formulated in Section 3.2.1. A multi-model Kalman filter for the MJLS based on the theory of section 2.2.4 will be used in Chapter 4 to detect the events of interest.

It is assumed that the FIR signal is directly related to the average plasma density n_e, the DML signal directly measures the toroidal plasma flux ϕ_p, and the PD signal is proportional to the density n_{H^0} of neutral hydrogen in the edge region; $y = \begin{bmatrix} y_{H_\alpha} \\ y_{FIR} \\ y_{DML} \end{bmatrix} = \begin{bmatrix} c \cdot n_{H^0} \\ n_e \\ \phi_p \end{bmatrix}$. Here, $c \sim 1 \times 10^{-15} \, \mathrm{V/m^{-3}}$ is the estimated proportionality constant between the neutral density and the measured H_α intensity. The proportionality factor of 1[m] between the electron density and the FIR measurement is based on the assumption that the average length of the line of sight of te FIR chords through the plasma is approximately 1 m. However, the absolute value of both is actually unimportant to filter operation.

Now models for the evolution of n_e, n_{H^0}, and ϕ_p that reproduce the signals of interest discussed in the previous two sections are required. Specifically, the models should prescribe sudden changes in the baseline for n_{H^0} and the time derivative of n_e and ϕ_p at LHTs and HLTs, and sudden increases followed by relaxations in n_{H^0} during ELMs. In order to describe the dynamics around the relevant events, it is chosen to use three distinct modes in the modeling. The first is associated with the dynamics in L-mode, the second associated with the dynamics in H-mode and the third is associated with the dynamics during an ELM. More specifically, the latter focusses on the short time window during the ELM in which the H_α signal is increasing rapidly.

First, let us assume that the baseline for n_{H^0} and the time derivative of n_e and ϕ_p do not change except during transitions. These assumptions were shown to hold coarsely in the figures of Section 3.1, and will be relaxed later in this section. In L-mode, the modeled edge neutral density tends to some baseline value $\bar{n}_{H^0}$ on a time scale τ according to a first order (exponential) decay:

$$\frac{d}{dt} n_{H^0} = \frac{\bar{n}_{H^0} - n_{H^0}}{\tau} \quad \text{in L-mode.} \tag{3.1}$$

In H-mode, the neutral density tends to a baseline that is a factor $f_H \sim 70\%$ lower, which corresponds to the observed lowering of the baseline in Table 3.1.

$$\frac{d}{dt} n_{H^0} = \frac{f_H \bar{n}_{H^0} - n_{H^0}}{\tau} \quad \text{in H-mode.} \tag{3.2}$$

The model value τ should be taken as the typical time-scale of signal relaxations during LHTs, HLTs and ELMs, $\sim$1 ms.

At the start of an ELM (or an H_α spike in L-mode), the 'ELM rise' model should expect a fast increase in neutral density. Typically, the H_α signal during an ELM increases to double the baseline in $\tau_{rise} \sim 0.25$ ms, leading to the following model:

$$\frac{d}{dt} n_{H^0} = \frac{\bar{n}_{H^0}}{\tau_{rise}} \quad \text{during an ELM rise.} \tag{3.3}$$

The modeled electron density n_e is changing at a given rate $\dot{n}_e$ during L-mode:

$$\frac{d}{dt}n_e = \dot{n}_e \quad \text{in L-mode.} \tag{3.4}$$

Just after the transition to H-mode, the improved confinement causes an increase in the time derivative of n_e of about $\Delta \dot{n}_e \sim 5 \times 10^{20}\,\mathrm{m^{-3}s^{-1}}$ (see Table 3.9) compared to L-mode:

$$\frac{d}{dt}n_e = \dot{n}_e + \Delta \dot{n}_e \quad \text{in H-mode.} \tag{3.5}$$

The modeled diamagnetic flux ϕ_p is also changing at some given rate $\dot{\phi}_p$ in L-mode:

$$\frac{d}{dt}\phi_p = \dot{\phi}_p \quad \text{in L-mode.} \tag{3.6}$$

Again, just after the LHT, an increase in the time-derivative of around $\Delta \dot{\phi}_p \sim 0.05\,\mathrm{Wb/s}$ (see Table 3.11) is expected:

$$\frac{d}{dt}\phi_p = \dot{\phi}_p + \Delta \dot{\phi}_p \quad \text{in H-mode.} \tag{3.7}$$

During an ELM, the FIR and DML signals drop, following the sharp spike in H_α intensity. This is modeled using the L-mode dynamics for n_e and ϕ_p. This is imperfect, since the effects of ELMs on the FIR and DML signals occur later and on a slower time-scale than the H_α peaks. However, implementing these differences would require the addition of at least a fourth model, so instead, ELMs are detected on the basis of mostly the H_α signal.

The quantities $\bar{n}_{H^0}$, $\dot{n}_e$ and $\dot{\phi}_p$ may change significantly over the shot aside from the effect of transitions. However, they will do so on a relatively slow time-scale. The a priori prediction of the signal evolution on a transport time-scale requires extensive modeling and is beyond the scope of this work. Instead, this study focuses on modeling the large signal features that accompany the LHT, HLT and ELM events. The issue of slow signal evolution is tackled by introducing $\bar{n}_{H^0}$, $\dot{n}_e$, and $\dot{\phi}_p$ as state variables, which are modeled as constants:

$$\frac{d}{dt}\bar{n}_{H^0} = \frac{d}{dt}\dot{n}_e = \frac{d}{dt}\dot{\phi}_p = 0. \tag{3.8}$$

The total state hence becomes $x = \begin{bmatrix} \bar{n}_{H^0} & n_{H^0} & \dot{n}_e & n_e & \dot{\phi}_p & \phi_p \end{bmatrix}^T$. The measurements do not depend directly on the newly introduced state variables, and can be written as:

$$y = \begin{bmatrix} y_{H_\alpha} \\ y_{FIR} \\ y_{DML} \end{bmatrix} = \begin{bmatrix} 0 & c & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} x, \tag{3.9}$$

The multi-model observer discussed in Section 2.2.4 is used to estimate all state variables and their possible change over time in the presence of noise. In the component Kalman filters, process noise is introduced which accounts for model-reality mismatch, and for example allows changes in the estimated $\bar{n}_{H^0}$, $\dot{n}_e$, and $\dot{\phi}_p$, even though the models assume these to be constants. The complete dynamics of the state are as follows:

$$\frac{d}{dt}x = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{\tau} & -\frac{1}{\tau} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} x \quad \text{in L-mode.} \tag{3.10}$$
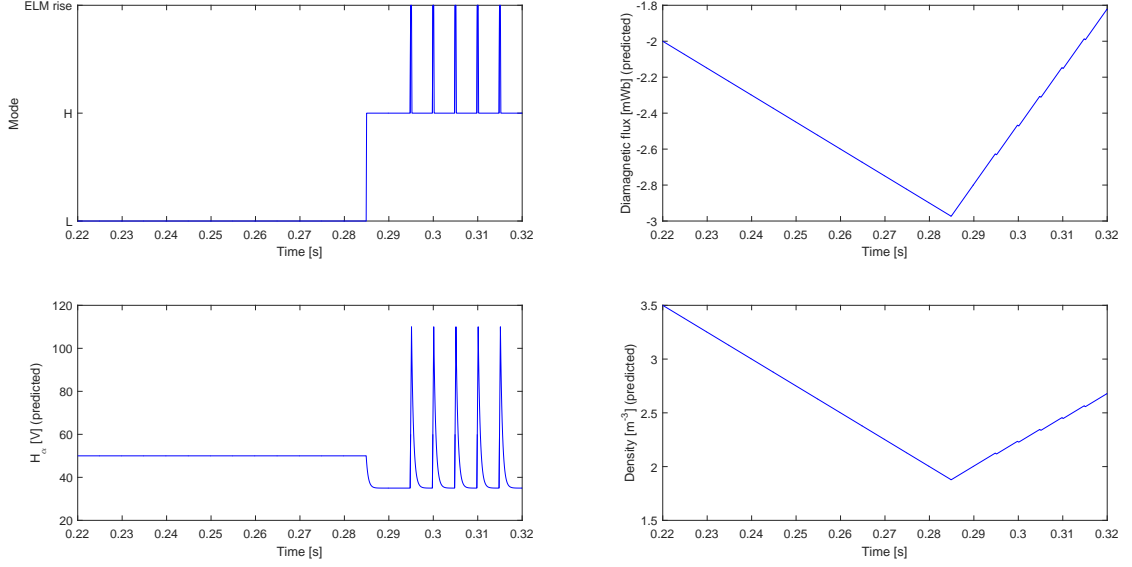
Figure 3.13: Model-only simulation of the H_α (bottom-left), FIR (bottom-right) and DML (top-right) signals. The chosen mode history (top-left) consists of an LHT at $0.285\,\text{s}$ and ELMs from 0.295 onwards.

$$\frac{d}{dt}x = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{f_H}{\tau} & -\frac{1}{\tau} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \Delta\dot{n}_e \\ 0 \\ \Delta\dot{\phi}_p \end{bmatrix} \quad \text{in H-mode.} \tag{3.11}$$

$$\frac{d}{dt}x = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{\tau_{rise}} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} x \quad \text{during an ELM rise.} \tag{3.12}$$

The role of each of these models is clarified in Figure 3.13, which shows the predicted signal evolution during typical features to be detected. The model history input is the subsequent occurence of two H_α spikes during L-mode, an LHT starting with an H_α spike, a number of ELMs and finally an HLT. This sequence of events translates to the mode history plotted in the top-left figure. The signal evolution for these events qualitatively agrees with that seen in Section 3.1.

As mentioned before, the TCV FIR signals suffer from a mechanical oscillation. The associated problems are tackled in Appendix B, which leads to the addition of two extra state variables, which reduce the effect these oscillations have on the detector's results.

3.2.1 Transition probabilities between modes

In order to design a multi-model estimator, the transition probability matrix $\Pi_k = [\pi_{ij,k}]$, which lists the probability of transitions between the different modes of the MJLS, has to be known. As noted in Section 2.2.4, Π_k governs the autonomous evolution of the mode probabilities as given by Equation 2.28. Also, if the Interacting Multiple Model (IMM) Hypothesis Reduction (HR) strategy is used, Π_k affects the way in which the estimates of different filters are mixed during each time step, as given by Equations 2.29 through 2.31.

For a data sampling rate of $10\,\mathrm{kHz}$, the base probability of a transition or back-transition during an individual time step is taken as (only) $p_0 \sim 0.1\%$, since transitions actually occur at only a few time steps out of thousands.

The chosen transition probabilities also depend on the heating power, through the explicit time-dependence of Π_k. Experimentally, an LHT becomes more likely at high heating power P compared to the threshold power, and an HLT becomes more likely at low (or no) heating power. This threshold power can be approximated by the scaling law value P_{sc}. The threshold power scaling by Martin [23] is used:

$$P_{sc} = 0.0488 n_{e20}.^{0.717} B_T^{0.803} S^{0.941}, \tag{3.13}$$

which is given in MW as a function of density in $1 \times 10^{20}\,\mathrm{m^{-3}}$, toroidal magnetic field in T, and plasma surface area in $\mathrm{m^2}$. Equilibrium reconstruction data from LIUQE contains contour coordinates of the last closed flux surface, which can be used to determine the plasma surface area.

The L-H transition probability π_{LH} should be 0 (or very small) at low heating power (for compactness, the time index k is dropped here). Depending on how efficiently the power is injected, more than twice the scaling law value of heating power may be necessary before the LHT occurs. Also, there may be a short delay between the application of a heating pulse and a transition. This means that π_{LH} should not increase limitlessly with increasing heating power. Furthermore, the H-L transition probability π_{HL} should be large for $P = 0$, decrease as the power increases, and saturate to some non-zero value for large P. If π_{HL} approached 0, this would force the detector to detect H-mode. Finally, ELMs and H_α spikes similarly increase in probability with increasing heating power but occur more often, so the transition probabilities π_{L-ELM} and π_{H-ELM} towards the ELM rise model is taken as $10 \times \pi_{LH}$.

The wanted behaviour is satisfied by the following expressions:

$$\pi_{LH} = p_0(\tanh(\frac{P}{w_\pi P_{sc}} - 1) + 1) \tag{3.14a}$$

$$\pi_{HL} = 2.5p_0 - \pi_{LH} \tag{3.14b}$$

$$\pi_{L-ELM} = \pi_{H-ELM} = 10\pi_{LH}, \tag{3.14c}$$

Which are shifted hyperbolic tangents. w_π determines the width of the hyperbolic tangent; the unshifted hyperbolic tangeant $\tanh(\frac{P}{w_\pi P_{sc}})$ reaches 76% of its maximum value at $\frac{P}{P_{sc}} = w_\pi$, and 96% of its maximum value at $\frac{P}{P_{sc}} = 2w_\pi$. A value of $w_\pi = \frac{1}{3}$ is chosen. The value of these transition probabilities as a function of input power is plotted in Figure 3.14.
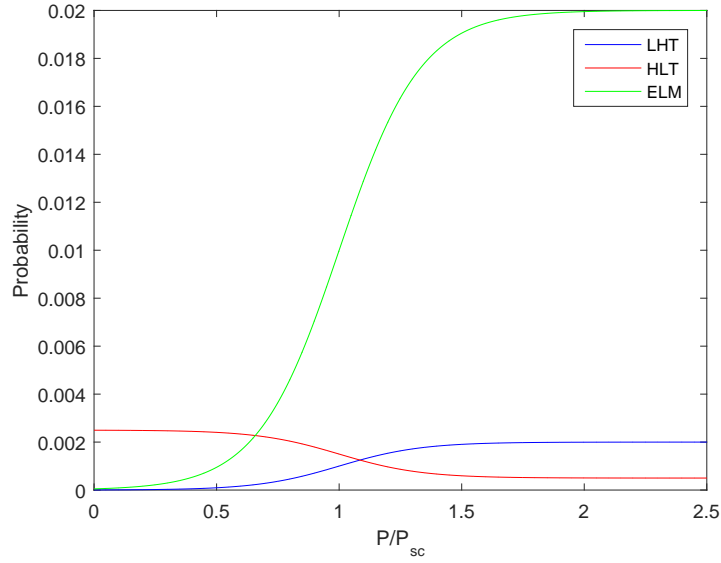
Figure 3.14: Scaling of different transition probabilities with the applied heating power. With increased heating power, the probability of an LHT (blue) or ELM (green) increases while the probability of an HLT (red) decreases.

To complete the transition probability matrix, a 20% probability of remaining in the ELM rise phase is chosen (since it is typically very short), and the remaining 80% is divided evenly over L- and H-mode. This means the detector relies on measurement data after an ELM peak to indicate the difference between H_α spikes and LHTs, or ELMs and HLTs. As such, the full transition probability matrix becomes:

$$\Pi = [\pi_{ij}] = \begin{bmatrix} 1 - \pi_{LH} - \pi_{ELM} & \pi_{HL} & 0.4 \\ \pi_{LH} & 1 - \pi_{HL} - \pi_{ELM} & 0.4 \\ \pi_{L-ELM} & \pi_{H-ELM} & 0.2 \end{bmatrix} \text{ where } \quad i, j \in \{\text{L}, \text{H}, \text{ELM}\}. \qquad (3.15)$$

Note that the sum of each column is 1. In other words, starting in any one particular mode, the sum of the probabilities of staying in that mode or transitioning to one of the other modes is 1.

Chapter 4

Event detection results

In this chapter, we use a Cooperating Multiple Model (CMM) filter to detect LHT, HLT, and ELM events in TCV diagnostic signals. This CMM filter is based on the Markov Jump-Linear System (MJLS) formulated in Chapter 3, with mode dynamics as given in Section 3.2 and transition probabilities between modes described in Section 3.2.1. The detector is a straightforwardly implemented filter for the MJLS according to the CMM theory described in Section 2.2.4. First, an overview of the training and test datasets in given in Section 4.1. Some different Hypothesis Reduction (HR) strategies are applied to data in the training dataset, and their differences in terms of detection performance are described in Section 4.2. Using the best performing HR strategy, the detector is applied to data from the test dataset in Section 4.3. In Section 4.4, the diamagnetic loop signal and the ELM timing specified by the detector are used to analyse the plasma energy losses during ELMs. Finally, the values of some influential filter parameters are objectively selected in Section 4.5, on the basis of a cost metric that penalises detection errors in both the training dataset and the test dataset.

4.1   Overview of the training and test datasets

An overview of all the TCV discharges used in this work is given in Table 4.1. In this table, the plasma shape is parametrised by κ_{95} and δ_{95}. κ is the vertical elongation of a flux surface $\kappa \equiv \frac{z_{max}-z_{min}}{R_{max}-R_{min}}$, defined by the maximum and minimum values of z and R on the flux surface. κ is the triangularity defined by $\delta \equiv 2\frac{(R_{max}+R_{min})-(R_{upper}+R_{lower})}{R_{max}-R_{min}}$ [46], where R_{upper} and R_{lower} are R at the upper-most and lower-most point on the flux surface. Both κ_{95} and δ_{95} are defined on the flux surface 95% of the way to the plasma edge. There, the shape paramaters are usually more well-defined than on the last closed flux surface. The snowflake divertor mentioned in the final column consists of a second-order magnetic null, which leads to the formation of many divertor legs [47].

Table 4.1: An overview of the training and test datasets used in this work. The "Dataset' column states whethr each shot belongs to the training or test dataset. 'Shot Goal' is the goal outlined in the TCV shot log, if any goal was stated. n_{LHT} is the number of L-H-L cycles occuring in each shot. $n_{e,19}$ is the achieved central electron density in 10^{19} m^{-3}, and β_p is the achieved normalised pressure (normalised by the poloidal magnetic pressure). κ_{95} is the elongation and δ_{95} the triangularity of the plasma. R_{axis} and z_{axis} the coordinates of the magnetic axis. The divertor geometry is also given.

Shot #	Dataset	Shot goal	n_{LHT}	$n_{e,19}$	β_p	κ_{95}	δ_{95}	R_{axis}	z_{axis}	Divertor
39872	Training	Snowflake	1	1.60	0.15	2.8	0.25	0.88	0.36	Snowflake
43454	Training	Snowflake	1	1.65	0.17	3.5	0.48	0.89	0.36	Snowflake
45103	Training		2	1.65	0.19	3.1	0.46	0.89	0.36	Snowflake
45104	Training		3	1.67	0.19	3.1	0.46	0.89	0.36	Snowflake
45105	Training		4	1.67	0.18	4.0	0.43	0.89	0.36	Snowflake
45106	Training		2	1.66	0.19	3.1	0.46	0.89	0.36	Snowflake
45107	Training		2	1.67	0.18	3.1	0.48	0.89	0.36	Snowflake
45108	Training		2	1.69	0.17	4.4	0.39	0.89	0.36	Snowflake
45109	Training		3	1.68	0.15	3.3	0.41	0.89	0.36	Snowflake
47007	Training	Snowflake	2	1.61	0.18	1.9	0.26	0.88	0.41	Snowflake
47008	Training	Snowflake	2	1.61	0.17	1.9	0.26	0.88	0.41	Snowflake
47306	Training		1	1.70	0.44	11.7	0.48	0.89	0.22	Single null
47307	Training		1	1.69	0.43	11.0	0.50	0.89	0.22	Single null
48005	Test	H-mode studies	1	1.56	0.31	3.8	0.29	0.90	0.07	Single null
48656	Test		1	1.56	0.31	3.8	0.29	0.90	0.07	Single null
48827	Test	Quasi-Stationary ELM-free H-mode	1	1.61	0.31	4.2	0.49	0.90	0.22	Single null
48854	Test		1	1.61	0.31	11.0	0.51	0.90	0.22	Single null
49330	Test		1	1.57	0.32	9.3	0.52	0.90	0.22	Single null

4.2 Event detection results for several hypothesis reduction strategies

In this section, the effect of using different HR strategies on the detection performance is investigated, with the goal of optimising the performance. The simple 1-Best HR strategy is used in Section 4.2.1, where between filtering steps, the state estimate of the most likely (i.e., 'best') model is used as the new state estimate for all filters. The more complex Interacting MM strategy is used in Section 4.2.2, for which each filter receives an individualised mixed state estimate, based on the probabilities of transitioning to (or staying in) the relevant mode. An HR strategy in which multiple time steps are explicitly taken into account is formulated and used in Section 4.2.3. Finally, a variation on the IMM strategy which emulates some advantages of the multi-step strategy and which is less computationally demanding is used in Section 4.2.4.

4.2.1 1-Best hypothesis reduction strategy

A simple HR strategy is the 1-Best algorithm, where the single best state estimate - as determined by the mode $B \in \{L, H, ELM\}$ which has the highest mode probability $\mu_{k|k}^{(i)}$ - is used by all filters in the next filtering step. In other words, $\hat{x}_{k-1|k-1}^{(B)}$ replaces $\bar{x}_{k-1|k-1}^{(i)}$ in Equation 2.32. The results for this HR strategy for TCV shot # 45103 are given in Figure 4.1. The wanted behavior of the 'best model' shown in the bottom-right panel is to move from L to H after each visually detected
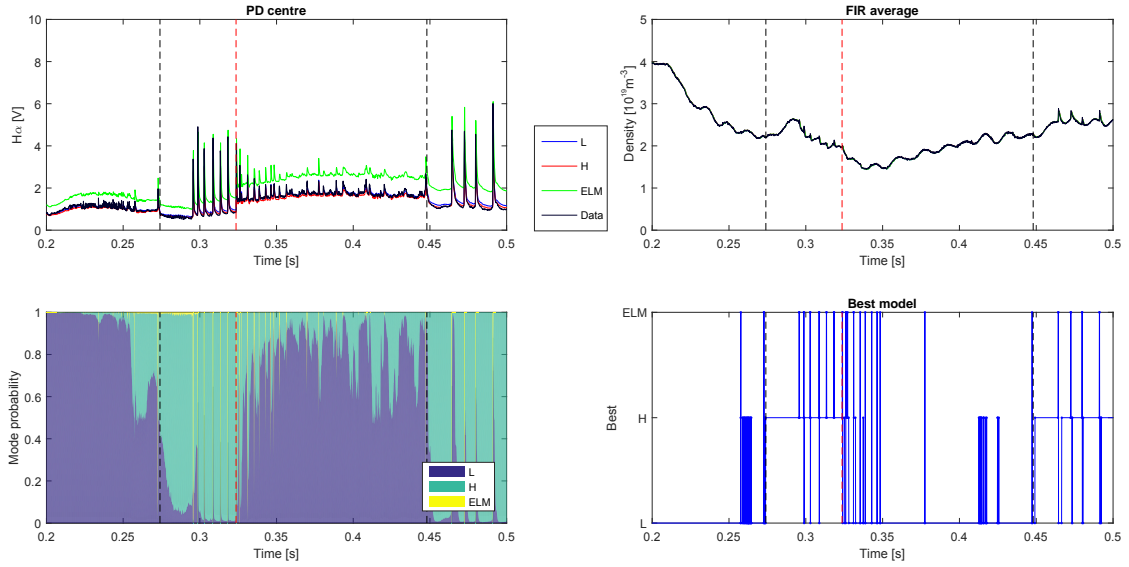
Figure 4.1: Detector results for TCV shot # 45103 when the 1-Best algorithm is used at each time step. Top: H_α (top-left) and FIR (top-right) measurements and expected (i.e., one-step-ahead predicted) measurements. Bottom-left: mode probabilities $\mu_{k|k}^{(i)}$. Bottom-right: chosen model (the one with the highest $\mu_{k|k}^{(i)}$). All as function of time. The vertical dashed lines represent the manually detected LHTs (black) and HLTs (red). Both LHTs and the HLT that occur within the selected time window are detected correctly. However, at some other times, short and 'deceptive' patterns in the signals are mistakenly detected as transitions.

LHT (black dashed vertical lines) and to move back from H to L after each visually detected HLT (red dashed vertical lines). Figure 4.1 shows that both LHTs and the HLT in the selected time interval are detected correctly.

However, some other signal features are wrongly detected as transitions. At 0.25 s, as the H_α signal relaxes back to its baseline after a small spike, this is wrongly interpreted as an LHT. Because of the relatively small sample time and slow time constant in the n_{H^0} model, there is only a small difference between the expected measurements between the two models, and even though the signal is not actually tending to a ~30% lower baseline, during a few time steps it looks like it is. The underlying problem is probably that H_α signal decay during all time steps cannot be modeled with a single time constant.

At 0.26 s, $\mu_{k|k}^{(L)}$ and $\mu_{k|k}^{(H)}$ are fluctuating around 50% and the ''best model' choice is constantly switching between L and H. This is because the real signal baseline is in between the baseline estimates of the by both models. Since both models start with the 'best' estimate at each time step, the alternation of the L and H models yields a signal which fluctuates around the real signal baseline only slightly. The underlying issue is caused by the HR strategy which makes hard decisions in the sense that only one mode hypothesis is selected as correct.

Around 0.3 s and 0.5 s, some ELMs are detected as the sequence ELM rise-L-H, where the L-mode model is chosen as the best for one or a few time steps. Moreover, around 0.33 s, some of the H_α peaks in L-mode are detected as ELM rise-H-L. This also happens because the decay time is sometimes overestimated by the models and sometimes underestimated.
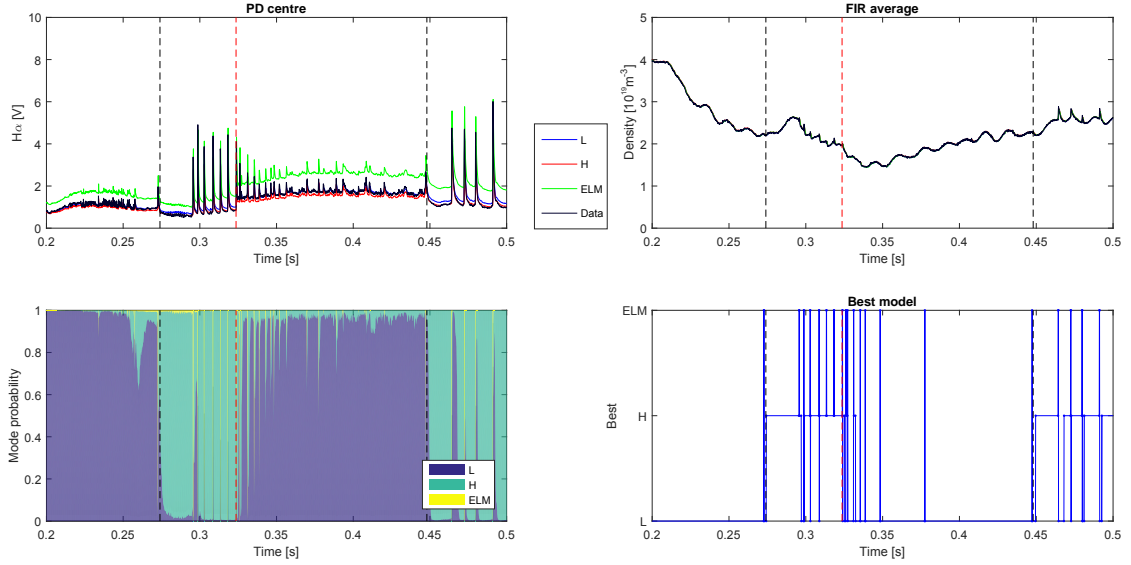
Figure 4.2: Detector results for TCV shot # 45103 when the IMM algorithm is used. Top-left: H_α. Top-right: FIR. Bottom-left: mode probabilities. Bottom-right: chosen model. The colored H_α and FIR signals are single-step and multi-step predicted measurements. Both LHTs and the HLT are detected. A few extra transitions are detected mistakenly.

4.2.2 Interacting multiple model hypothesis reduction strategy

An alternative HR strategy is the IMM algorithm, where each filter starts the next time step wih an individualised weighted average of all state estimates, where the weights are determined by the mode probabilities $\mu_{k|k}^{(i)}$ and the transition probabilities to the mode in question. See Section 2.2.4 and in particular Equations 2.29 through 2.31. For the Π-matrix given in Section 3.2.1 which says that mode transitions are rare ($\sim 0.1\%$ of time steps), this means that the 'best' model only dominates the mixed estimate for all models if its mode probability is by far the largest. If multiple models have a significant mode probability, their mixed estimates will be dominated by their own posterior estimates.

Figure 4.2 shows the results for TCV shot # 45103 when the IMM algorithm is used. The problem at $0.25\,\text{s}$ is solved, because as the L-mode probability increases after a few time steps of better measurement predictions, the mixed estimates move further apart and the L-mode predictions become worse. However, the problem with the ELMs and spikes is only mitigated, possibly because the mode probabilities just before an ELM are thrown away entirely when the 'ELM rise' model reaches $\sim 100\%$ mode probability.

4.2.3 Decisions based on multiple time steps

The detection of the modes following peaks in the H_α signal may be improved by explicitly taking into account multiple time steps. In this case, the different filters will first perform a time-update for N time steps and predict a series of N measurements based on some assumed N-step mode history, and only then perform a measurement update for the same time steps and compare their results to evaluate which filter produced the best prediction. The N-step mode histories that are used are L-L-...L and H-H-...H

However, such a multi-step decision strategy is not useful for each time window of N consecutive time steps. If, for example, time windows are included where an ELM occurs somewhere within the window, the detector will give unpredictable results unless $\sim N$ extra mode history hypotheses are included to allow for the possible occurence of one (or more) ELMs at any point in the time window, which would lead to large computational requirements.

Because multi-step decisions would be especially useful at only a handful of time instants (e.g., right after an H_α peak) and require many mode history hypotheses when used every N time steps, the single-step and multi-step strategies should be combined. By default, the detector then uses a single-step hypothesis reduction strategy, but only until it reaches a point at time k where a multi-step decision is required. The multi-step decision is initiated at time step k when at least one of the following two conditions is satisfied:

- Immediately following an H_α peak. In other words, when in the previous time step, the ELM model was chosen, but in the current time step, another model was chosen: $\mu_{k-1|k-1}^{(ELMrise)} \geq \mu_{k-1|k-1}^{(i)})\forall i \wedge \exists i : \mu_{k|k}^{(ELMrise)} < \mu_{k|k}^{(i)}$

- Immediately following an LHT or HLT as detected by the single-step decision strategy. In other words, when in the previous time step, L (H) mode was chosen, but in the current time step, H (L) mode is chosen: $(\mu_{k-1|k-1}^{(L)} \geq \mu_{k-1|k-1}^{(i)} \forall i \wedge \mu_{k|k}^{(H)} \geq \mu_{k|k}^{(i)} \forall i) \vee (\mu_{k-1|k-1}^{(H)} \geq \mu_{k-1|k-1}^{(i)} \forall i \wedge \mu_{k|k}^{(L)} \geq \mu_{k|k}^{(i)} \forall i)$

At that point, an N-step time update, measurement update, and mode probability are performed for two different model histories, N times L-mode and N times H-mode. The ELM rise prior mode probability is taken as 0 and the L-mode and H-mode are renormalised to a sum of 1. After this single iteration of the multi-step decision strategy, single-step iterations are used again.

The results for an algorithm switching between single-step and multi-step decision strategies are shown in Figure 4.3. In both cases, the 1-Best strategy is used for simplicity. The multi-step strategy uses a time window of 75 steps corresponding to 7.5 ms. The figure shows that this window is in some cases too long. Most notably, at 0.3 s, two complete ELMs fit into the decision time interval, which causes a faulty HLT detection.

A length of $N = 75$ time steps for multi-step decisions is apparently too long if the next ELM is included in the interval. However, in cases where the next ELM does not occur within the interval, such longer time intervals can be useful or even necessary. A possible improvement to the multi-step decision strategy is to first check whether the next ELM occurs within the interval, and if it does, to shorten the decision interval. This can be done using a simple single-step ELM detector to identify the occurence of the next ELM. This detector is based on two modes, 'ELM rise' and 'no ELM rise', where the latter has $\dot{x} = 0$ and the former has the same neutral density dynamics as the regular 'ELM rise' model, but no density or diamagnetic flux evolution, $\frac{d}{dt}n_e = \frac{d}{dt}\phi_p = 0$.

The general decision strategy then becomes as follows: the detector uses a single-step decision strategy until one of the aforementioned triggers indicate that a multi-step decision is required at time k. At that point, the single-step ELM detector runs for up to N_{max} time steps, or until it detects the next ELM rise at time $k + N$. A multi-step decision is then made for the interval between time k and $k + N$ (with $N = N_{max}$ if no next ELM was detected).

Figure 4.4 shows the results for this hybrid single-step/multi-step decision strategy with variable multi-step time interval. All features in the selected time window are detected correctly. In fact, now all features in the whole shot are detected correctly, with the exception of a few near the end.
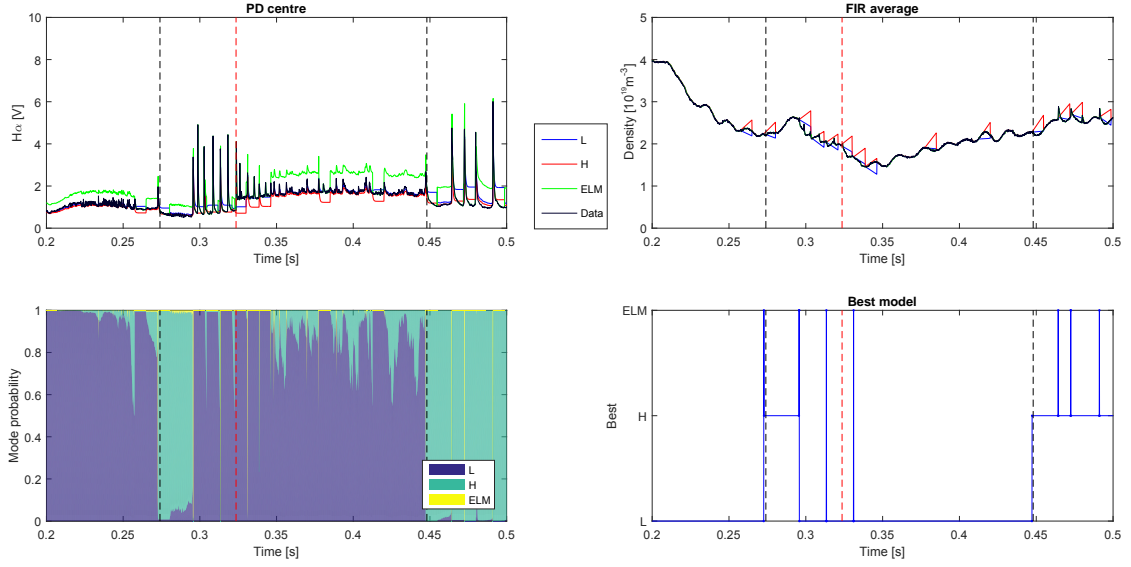
Figure 4.3: Detector results for TCV shot # 45103 when single-step or multi-step decisions are used, depending on the recent mode history. Top-left: H_α. Top-right: FIR. Bottom-left: mode probabilities. Bottom-right: chosen model. Both LHTs are detected, however two subsequent ELMs are mistaken for the HLT. Most ELMs after the mistaken HLT are also missed since multiple ones are combined into a single decision interval.

These exceptions are shown in Figure 4.5. The final HLT is detected correctly, but shortly afterwards, a decrease in the H_α signal related to other aspects of the current ramp-down phase is detected as an LHT. Without the explicit knowledge that the shot is nearing its end, it is difficult to distinguish this signal from a real LHT even visually, so it is perhaps unsurprising the the detector gives a false positive in this case.

In three cases, two ELMs follow each other very closely. In such a 'double ELM', when the second ELM starts, the H_α signal has not yet reached its baseline and the FIR signal is still strongly decreasing, which corresponds well to the L-mode model's prediction. This may be solved by rejecting multi-step decisions that are based on a time-interval that is too short.

During an ELMy H-mode, the minimum length of a decision interval should be related to the typical time interval between ELMs. To this end, an estimate of the inter-ELM interval Δt_{ELM} is tracked. If at time k, an ELM is detected a time Δt_k after the detection of the previous ELM and the previous estimate of the inter-ELM time was $\Delta t_{ELM,k-1}$, then the new observation gets a weight of 30% in the new estimate $\Delta t_{ELM,k}$. However, if there was a long pause after the previous ELM (at least twice as long as the expected interval), the estimate of the ELM interval remains unchanged:

$$\Delta t_{ELM,k} = \begin{cases} 0.7\Delta t_{ELM,k-1} + 0.3\Delta t_k & \text{for } \Delta t_k < 2\Delta t_{ELM,k-1}, \\ \Delta t_{ELM,k-1} & \text{for } \Delta t_k \geq 2\Delta t_{ELM,k-1}. \end{cases} \qquad (4.1)$$

During H-mode, a multi-step decision is deemed trustworthy if the decision interval is longer than 60% of the expected ELM interval. During L-mode, the minimum interval length is 40 time steps. This quite long minimum interval prevents false LHT detections in a few cases where the current ramp-up or plasma shaping causes a short dip in the H_α signal. Meanwhile, it leads to no false negatives since signal features in or shortly after L-mode are typically slower than in H-mode.
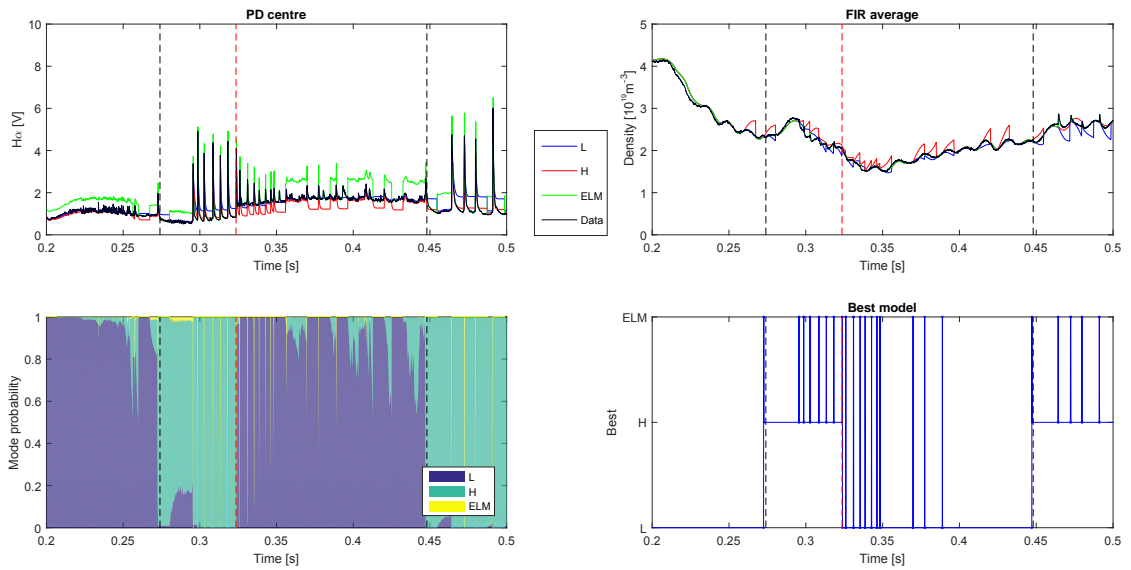
Figure 4.4: Detector results for TCV shot # 45103 when single-step or multi-step decisions of a time interval of varying lenth are used, depending on the recent mode history. Top-left: H_α. Top-right: FIR. Bottom-left: mode probabilities. Bottom-right: chosen model. All features are detected.
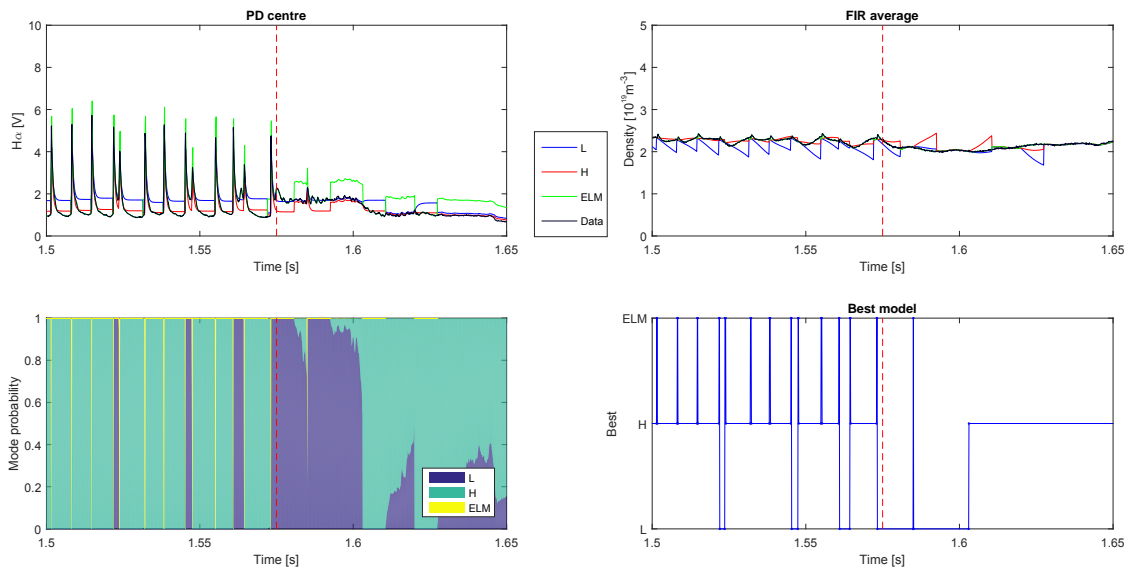

Figure 4.5: Detector results for TCV shot # 45103 when when single-step decisions or multi-step decisions for a time interval of varying lenth are used. Top-left: H_α. Top-right: FIR. Bottom-left: mode probabilities. Bottom-right: chosen model. Towards the end of the shot, three 'double' ELMs are mistaken for short L-modes. After the final HLT, signal features related to the ramp-down are mistaken for an LHT.
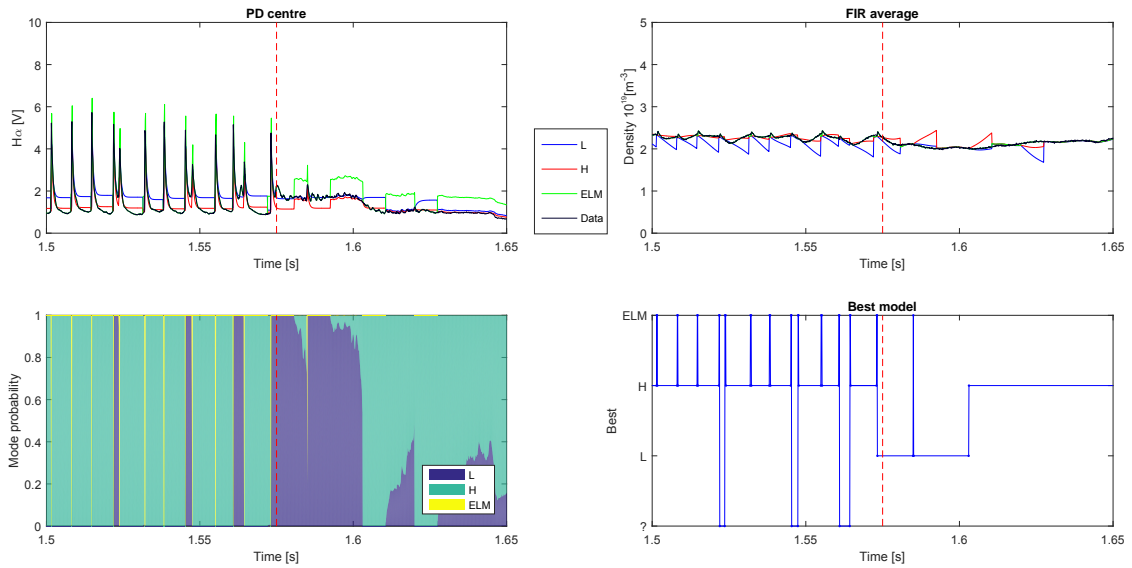
Figure 4.6: Detector results for TCV shot # 45103 when single-step decisions or multi-step decisions for a time interval of varying lenth are used, and short intervals are disregarded. Top-left: H_α. Top-right: FIR. Bottom-left: mode probabilities. Bottom-right: chosen model. Despite their high L-mode probability, the three 'double ELMs' now give a '?' as the best model, indicating that the decision interval was too short to be trustworthy and no decision could be made.

If at any time, a decision interval is too short to be trustworthy, the detector ignores the updated mode probabilities after the interval and chooses the best model based on the previous time step. In other words, the detector will stay in the same mode it was already in. If multiple subsequent short time intervals have a combined length longer than the minimum interval and all of them point at a transition, the combined interval is considered trustworthy and the transition is detected.

Figure 4.6 shows the detector results when this check on the length of each decision interval is implemented. It can be seen that the problem with double ELMs has been solved; the short intervals that were previously flagged as HLTs are now ignored.

4.2.4 A single-step strategy emulating multi-step behaviour

A disadvantage of the hybrid single-step/multi-step decision strategy described in the previous subsection is that it is difficult to implement in real-time in a natural way. When the need for a multi-step decision is encountered in real time, the detector needs to wait for all the necessary measurements to become available before it can calculate the innovations and hence the likelihoods with which to update the mode probabilities. Even the multi-step time-update is not without problems, as it is unknown on beforehand how many time steps are necessary. In the best-case scenario, these issues lead to a significantly varying computational load over the shot, so computational power has to be budgeted to accomodate the peak computational requirement and computational power may be wasted during other times. Also, the multi-step decisions require the somewhat ad hoc addition of the minimum interval length to prevent problems with the detection of double ELMs. This adds complexity, interferes with the MM filter's basic operation and adds more tuning parameters.

In this subsection, a hypothesis reduction strategy is formulated which similarly explicitly requires multiple time steps of new signal behaviour, but which can be executed time step per time step. This strategy is based on the IMM algorithm, since the IMM was the most robust of the single-step HR strategies which have been used.

The requirement of a minimum decision interval is implemented by differentiating between the model with the highest probability and the mode that is actually reported. Essentially, this is done by applying a conservative time delay to the reported mode. Conservative in the sense that the delay is applied only to an observed mode that deviates from the norm, i.e., the previously reported mode. This is done as follows: if the best model does not agree with the current mode, the 'conflicting mode interval counter' increases by 1. If the best model does agree with the current mode, the counter decreases by 1 (if it is not already 0). If the ELM model is chosen, it neither decreases nor increases the counter. Only if the current highest-probability model disagrees with the current mode, and the conflicting mode counter has reached a minimum, does the detector switch to reporting the new mode. This minimum interval length is 3 ms in L-mode and 70% of the typical inter-ELM time in H-mode. In this version of the detector, the inter-ELM time estimate is only updated if the new detected inter-ELM time is between 0.5 and 2 times the current estimate. Notably, this approach does not interfere with basic MM filter operation and is equivalent to minor post-processing of the 'best model' value.

During an ELM, the ELM mode probability is typically >99% for one or two time steps. As such, a few time steps after the ELM, the mode probability is quite evenly divided between L mode or H mode (as per the 40% to L, 40% to H, 20% to ELM transition probabilities away from the ELM model), regardless of how confident the detector was in both L and H models just prior to the ELM. Combined with the decision interval logic that favors the model with the highest probability regardless of the margin, this may cause problems. The chosen workaround is to adapt the transition probabilities as each ELM occurs. If at time step k, the ELM model has the highest probability while it did not have the highest probability at time step $k-1$, the transition probabilities from the ELM model to the L- and H-model are changed to the L and H-mode probabilities at time step $k-1$, renormalised to 0.8. The ELM→ELM probability remains 20%. This ensures that after each ELM (or H_α spike), the detector will tend to return to the same L and H mode probabilities as before the ELM.

The results for all shots for this version of the detector are given in Appendix A. These results are summarised in table 4.2, and show that the detector performs well for 11 out of 13 shots, and makes a minor mistake for the 2 other shots.

Table 4.2: The outcome of applying the detector to data in the training data-set. The 'Results' column reports how many errors (and what kind) the detector made compared to visual detection, the 'Notes' column briefly explains what signal features caused the wrong detections.

Shot #	Results	Notes
39872	Correct	
43454	Correct	
45103	Correct	
45104	Correct	
45105	Correct	
45106	Correct	
45107	Correct	
45108	Correct	
45109	1 false positives over 20 ms	The H_α signal relaxes to an unexpectedly high baseline during an ELM-free period
47007	Correct	
47008	1 false positive over 10 ms	The H_α signal relaxes to an unexpectedly low baseline after the final HLT
47306	Correct	
47307	Correct	

4.3   Results for the test data-set

The model values and hypothesis reduction strategy of the detector have been optimized to allow it to detect features in the training dataset. In this section, the detector is applied to data in the test data-set to verify its performance on shots that have not been used in its design.

Figures 4.7 through 4.13 show the detector results for the five shots in the test data-set. The detector detects most features of interest correctly. However, in some cases, irregular signals cause false positives. These results are summarized in Table 4.3.

Table 4.3: The outcome of applying the detector to data in the test data-set, with results and notes on what signal features were detected incorrectly.

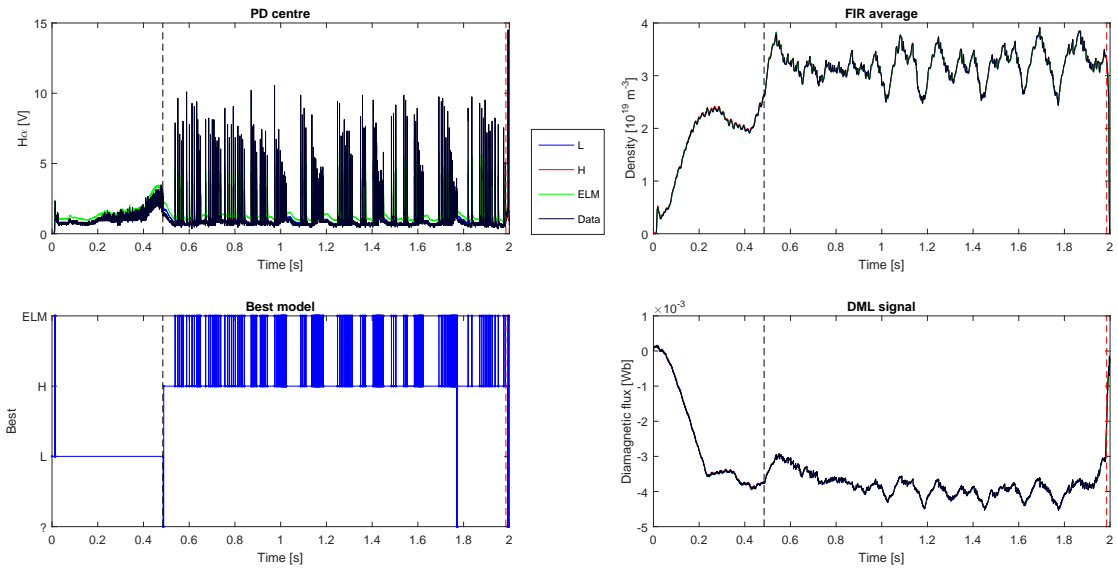| Shot #	Results	Notes
48005	Correct	
48656	Correct	
48827	1 false positive over 30 ms	H_α increase during ELM-free period
48854	Correct	
49330	2 false positives over 50 ms	Strong H_α fluctuations lead to unpredicatable outcomes

Figure 4.7: Detector results for shot # 48005. Top-left: H_α. Top-right: FIR. Bottom-left: chosen model. Bottom-right: DML. Colored measurement signals are the one-step-ahead measurement predictions. The vertical dashed lines indicate manually detected LHTs (black) and HLTs (red). All features are detected correctly.
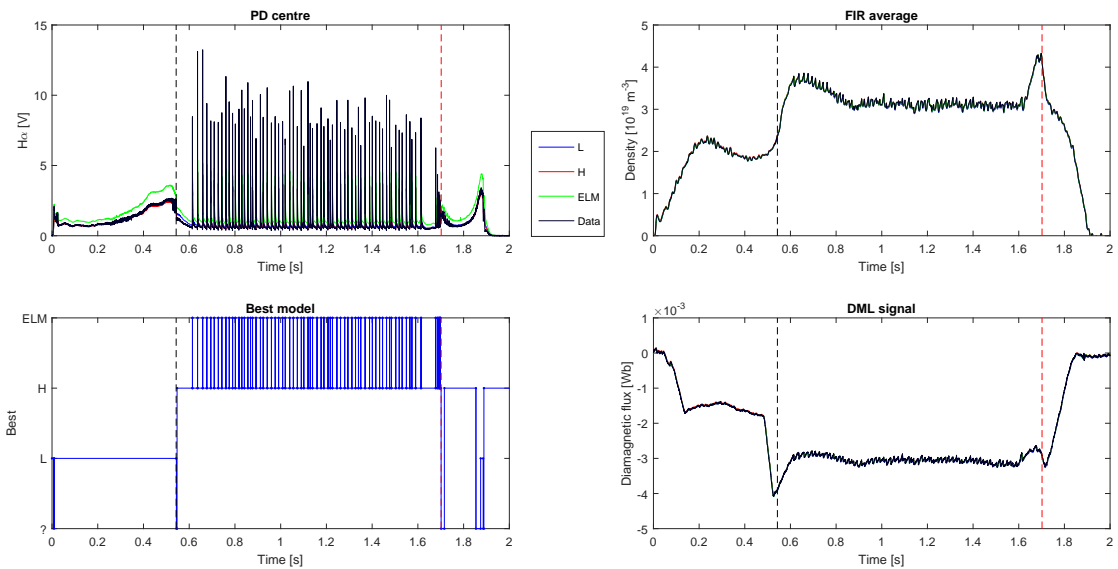

Figure 4.8: Detector results for shot # 48656. Top-left: H_α. Top-right: FIR. Bottom-left: chosen model. Bottom-right: DML. All features up until the current ramp-down are detected correctly.
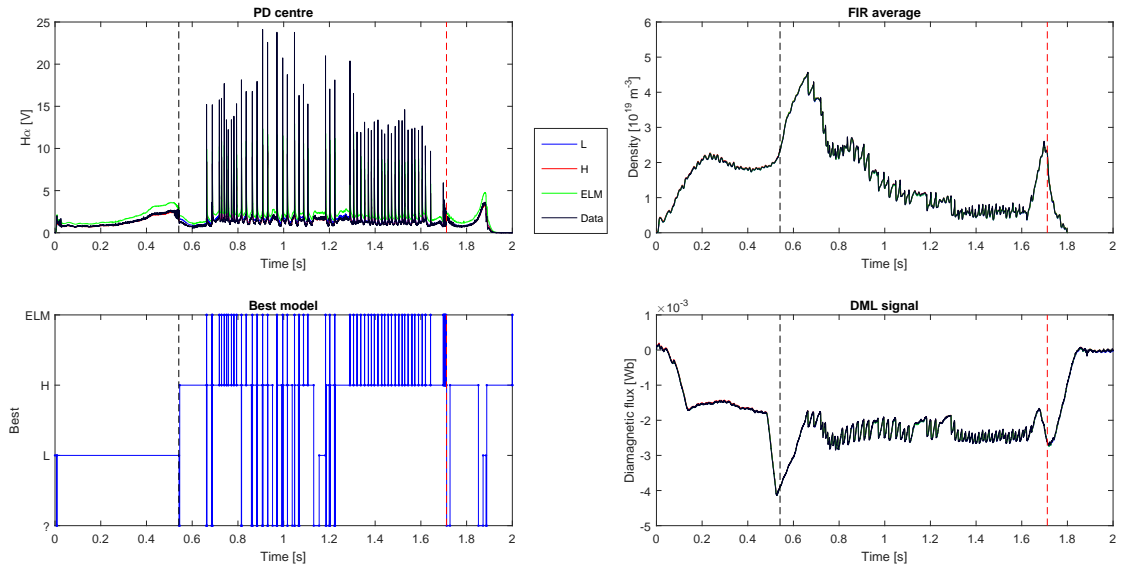
Real-time, Model-based Event Detection in Tokamaks.

Figure 4.9: Detector results for shot # 48827. Top-left: H_α. Top-right: FIR. Bottom-left: chosen model. Bottom-right: DML. An H_α signal feature at $1.15\,\mathrm{s}$ is shortly mistaken for L-mode.
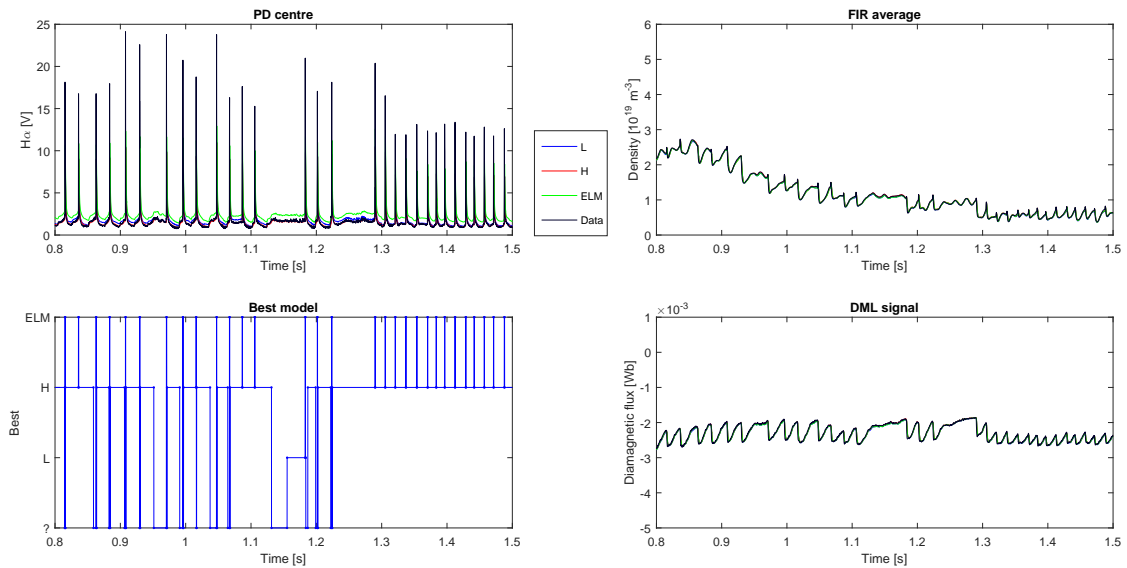

Figure 4.10: Zoomed-in detector results for shot # 48827. Top-left: H_α. Top-right: FIR. Bottom-left: chosen model. Bottom-right: DML. The colored signal curves are one-step measurement predictions. Around $1.13\,\mathrm{s}$, the H_α signal increases quite slowly during an ELM-free period. Despite the FIR and DML signals showing no decrease, the detector chooses L-mode until the ELMs resume at $1.18\,\mathrm{s}$.
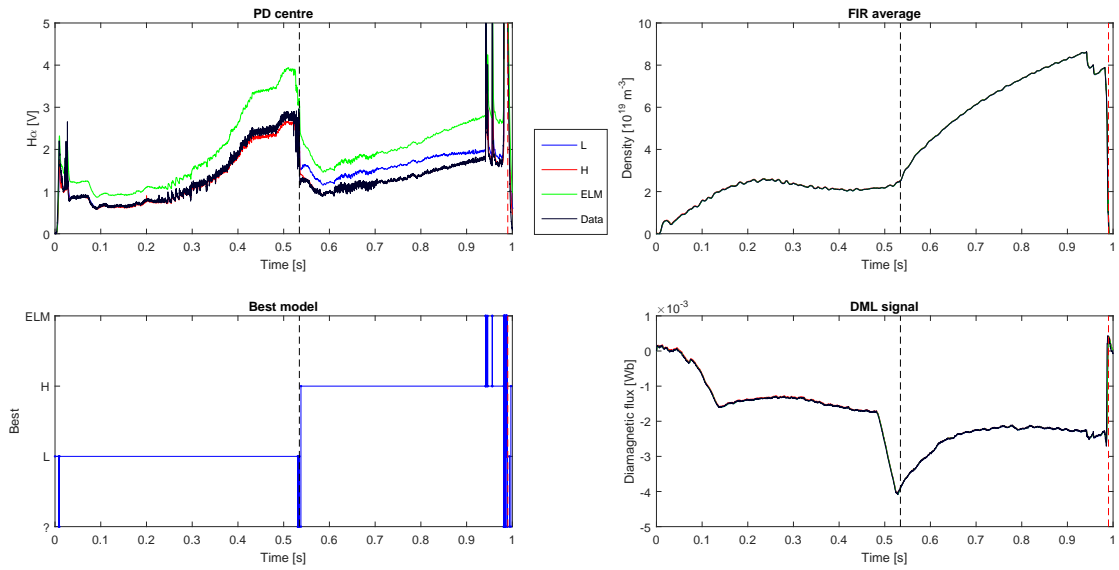
Figure 4.11: Detector results for shot # 48854. Top-left: H_α. Top-right: FIR. Bottom-left: chosen model. Bottom-right: DML. All features are detected correctly.
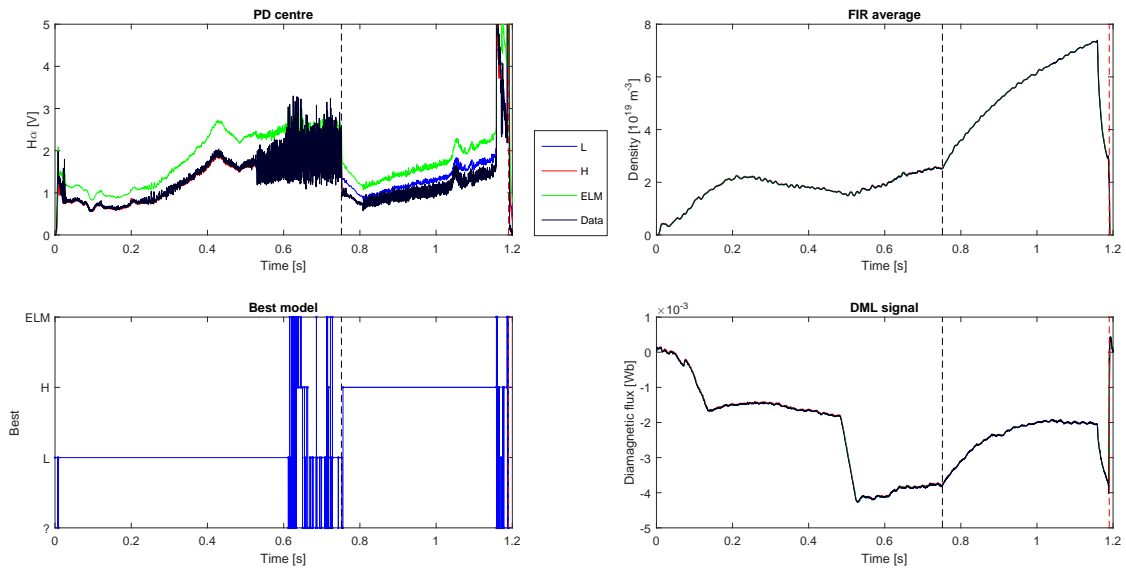

Figure 4.12: Detector results for shot # 49330. Top-left: H_α. Top-right: FIR. Bottom-left: chosen model. Bottom-right: DML. Before the LHT at 0.75 s, the H_α signal contains large dithers, possibly related to the Intermediate phase (see Section 2.1.4), which lead to faulty H-mode detections.
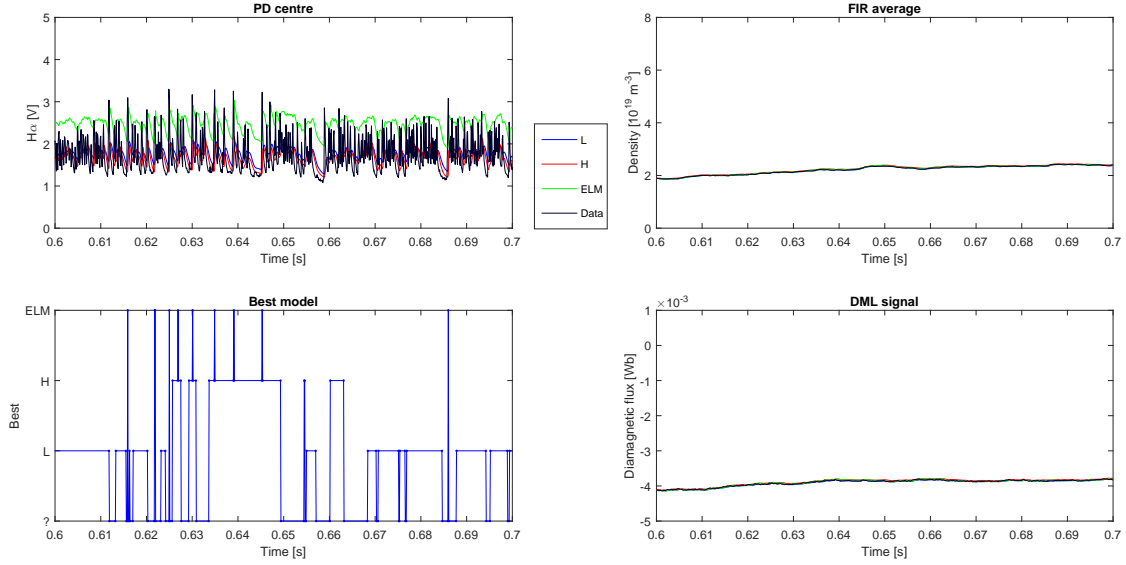
Figure 4.13: Zoomed-in detector results for shot # 49330. Top-left: H_α. Top-right: FIR. Bottom-left: chosen model. Bottom-right: DML. In the Intermediate phase just before the LHT, the H_α signal starts oscillating wildly. In some short time periods, a relaxation towards a lower baseline is seen. Many such short periods after each other trigger an H-mode detection twice.

4.4 Analysis of plasma energy losses during edge localised modes

In this section, the energy losses due to ELMs are analyzed using the ELM detection results of the adapated IMM filter described in Sections 4.2.4 and 4.3, and the relation between the plasma internal energy and the diamagnetic loop signal of Section 2.3. Moreover, the correlation of the ELM frequency and ELM energy losses is investigated for the detected ELMs in the dataset.

For TCV shot # 43454, the plasma shape term $s = \frac{1}{2}[s_1 + \delta_R s_2]$ (see Eqs. 2.53, 2.54 and 2.65) was calculated on the basis of the reconstructed poloidal magnetic flux Ψ for all time instants for which LIUQE equilibrium reconstruction data was available. In the TCV real-time control system, LIUQE runs at a constant sampling rate of $1\,\mathrm{kHz}$, but for archived shots, the sample frequency is lower and varies over the shot. For example, during the current ramp-up and ramp-down, the sample frequency is much lower.

Notably, the inclusion of plasma pressure in the weighting factor g (see Eq. 2.62) typically influenced the resulting plasma shape term (and hence β_p) by less than 1%-point, and the calculations converge to 1 part in 1000 in just 3 iterations. This validates the assumption of small influence of p on δ_R on which the iterative solution strategy was based.

At each time instant of a DML measurement, the most recent shape term calculation from LIUQE data was used to calculate β_p (i.e. a zero order hold was used). The calculated value of β_p throughout the shot is plotted along with the value calculated by LIUQE in Figure 4.14. As can be seen, there is significant agreement between the curves; at time instants for which LIUQE has calculated β_p, the calculated β_p is on average 2%-point higher while the rms difference between the two curves is 3%-points.

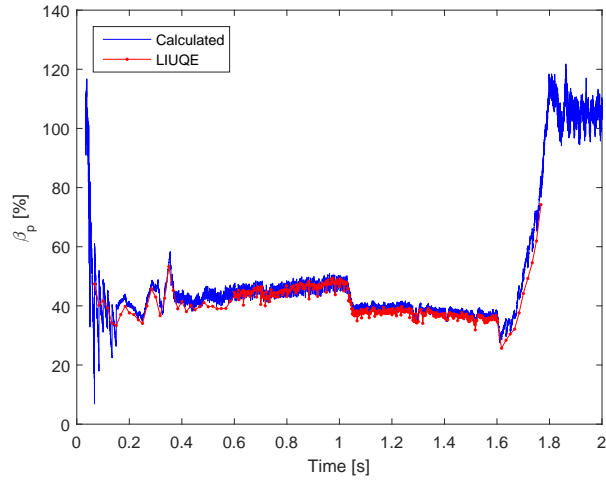Figure 4.14: β_p as calculated from the DML signal and equilibrium data from LIUQE (blue), and LIUQE's own calculated β_p (red) as a function of time. Throughout the shot, there is significant agreement between the two. The oscillations in the calculated β_p at the start of the shot, and its convergence to 100% towards the end, are incorrect results of a poor time resolution in LIUQE around those times, which appear because a zero order hold is applied to the plasma current in Eq. 2.61.

At the start of the shot, the calculated value of β_p is oscillating around the LIUQE values. This is because the archived time resolution of LIUQE is poor at this time. The plasma current I_p is taken as a zero-order hold from LIUQE, and the rapid increase in I_p during the current ramp-up leads to large errors in the diamagnetic flux term. Similarly, at the last time instant where LIUQE was run, the current is not yet ramped back to zero, so the expected diamagnetic flux is finite instead of 0, erroneously giving a calculated β_p of around 100%.

During the bulk of the shot, however, the time resolution of LIUQE reconstructions is good enough to allow accurate calculation of β_p in between the LIUQE time instants. As seen in Figure 4.15, this even allows the ELMs' effect on β_p to be visualised.

In order to speed up calculations, LIUQE's result for β_p can also be used directly, which saves the computational time needed to perform the Shafranov integrals. To this end, it is assumed that the plasma shape term s changes slowly with respect to the LIUQE sample frequency (i.e., LIUQE runs at an appropriate rate). Furthermore, it is assumed that s is directly measured by

$$y_s = \beta_{p,L} + \frac{\phi_p}{\phi_{0,L}}, \tag{4.2}$$

where $\beta_{p,L}$ and $\phi_{0,L}$ are LIUQE's result for β_p and ϕ_0 (ϕ_0 follows from LIUQE's reconstructed I_p), and ϕ_p is the measured diamagnetic flux. y_s contains noise, due to noise in the DML measurement, the magnetic probe measurements, the other measurements used to constrain LIUQE, and any approximations within LIUQE. In order to combat this noise, a simple Kalman filter is used to extract an estimate $\hat{s}$ out of the measurements y_s. The accurate modeling of plasma shape dynamics requires the complete modeling of the time-evolution of the MHD equilibrium. This is outside the scope of the present work, so s is simply modelled as a constant. Because of the model assumption of constant s, this filter works similarly to a simple low-pass filter. One important distinction is that due to the difference in time steps in offline LIUQE data, the process noise covariance is scaled proportional to the step time, leading to a larger Kalman gain at times of a lower sample frequency.
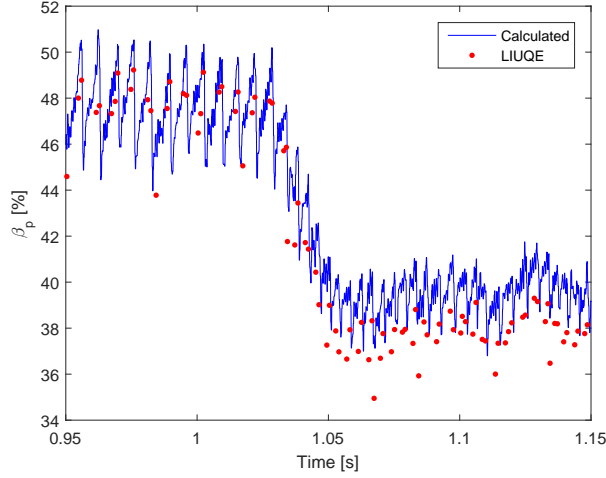
Figure 4.15: β_p calculated from the DML signal (blue line) and from LIUQE (red points), zoomed in in time. Energy losses due to ELMs are clearly visible in the first signal as a sharp decrease in β_p lasting typically just under 1 ms, while the LIUQE signal has too low time resolution to allow the visualisation of ELMS, and can only show trends in β_p.

Using such an estimate $\hat{s}$, β_p can be estimated in real-time at the DML time resolution of 10 kHz, when LIUQE's time resolution of 1 kHz is fast compared to the time scale of changes in the plasma current and shape (but optionally still slow compared to the time scale of e.g. ELMs). This estimate $\hat{\beta}_p$ is calculated similarly to Eq. 2.65 as:

$$\hat{\beta}_p = \hat{s} - \frac{\phi_p}{\phi_{0,L}}. \tag{4.3}$$

Using the information on ELM timing from the multi-model event detector, the percentage of plasma kinetic energy expelled during an ELM whose H_α signature starts at time k can be calculated using:

$$f_{E,loss} = \frac{\beta_{p,k} - \min(\{\beta_{p,k}, ..., \beta_{p,N_{hor}}\})}{\beta_{p,k}}, \tag{4.4}$$

where the minimum value of β_p within a horizon of the next N_{hor} time-steps is used. In principle N_{hor} is 20 time steps (2 ms), but is smaller in cases where the next ELM's H_α signature starts before 20 time steps . Note that for a typical type I ELM, the energy is expelled during 10 time steps or 1 ms for each ELM.

Using this formalism, it can be calculated that in shot # 43454, when the ELM frequency suddenly increases from $\sim$170 Hz to $\sim$225 Hz at $t = 1.03$ s (see Figure 4.15), the average expelled energy per ELM decreases from 9.8% to 4.8%, which may be useful information for real-time applications such as ELM monitoring and control.

In Figure 4.16, the relation between inter-ELM time and the fraction of expelled energy during an ELM is examined in 10 shots. To the eye, it seems clear that higher inter-ELM times typically correlate with higher fractions of expelled energy. However, a simple least-squares fit for a straight line does not agree with this result, instead being dominated by a few outliers with high inter-ELM time and low expelled energy. These outliers are likely mostly ELMs which were preceded by a short pause in ELMs. In order to reconcile the results with the intuitive reading of the plot, multiple linear fits excluding some different definitions of outliers have been taken, the results of which are summarised in Table 4.4. Pearson's r, the correlation coefficient mentioned in the table,
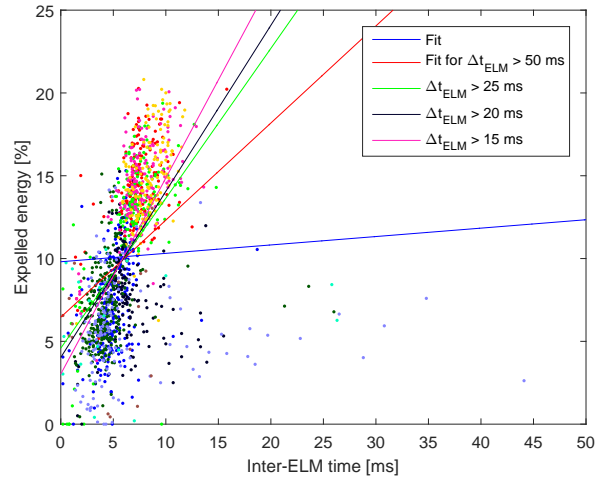
Figure 4.16: The inter-ELM time versus the fraction of energy expelled by the ELM. Different colour points indicate ELMs in different discharges. Higher inter-ELM times seem to correlate with larger fractions of expelled energy. A linear fit in which all points are included (blue line) is however dominated by some bottom-right outliers. Fits with increasingly strict upper bound on inter-ELM times (other lines) seem to visually fit the data increasingly better.

is a direct measure of the linear correlation between two variables. That is, $r = \pm 1$ if the fitted points lie perfectly on a straight line, with the $\pm$-sign determined by the sign of the fit's slope, and $r = 0$ if there is no (linear) correlation between the two variables. In all of these fits, 11 points with 0% expelled energy are excluded.

The figure shows that as the criterion for accepting points is strengthened, the fit seems to visually better represent the data, and the table shows that the calculated correlation coefficient increases in this case. Even for the strictest fit however, the slope seems to be too small and the offset too large to represent the main cluster of points well. Regardless, the results show some correlation between inter-ELM time and expelled energy, as expected.

Table 4.4: For the plots in Figure 4.16, the maximum allowed time between ELMs in order to exclude outliers, the amount of ELMs that were actually included in the fit, and Pearson's r of the fit are tabulated. As the allowed inter-ELM time decreases, the amount of included points changes only slightly, while the correlation coefficient of the fit improves significantly.

Maximum inter-ELM time [ms]	Amount of included points	Correlation coefficient (Pearson's r)
∞	1471	0.15
50	1455	0.41
25	1445	0.50
20	1442	0.53
15	1433	0.59

4.5 Objective selection of detector parameters

For a chosen decision strategy, many filter parameters can influence the behaviour of the multi-model filter. In Section 4.2, the filter values were broadly tuned to common-sense values for each HR strategy. This amounted to iteratively identifying signal features that were detected incorrectly but seemed easy enough to detect visually, and - using knowledge of how the filter values affect the result - adapting one or more of the filter values in steps of ∼30-100% until the feature was detected correctly.

In this section, a more objective approach to tuning the detector's behaviour is taken. The first step is to define a cost metric C for the detection quality, defined in terms of a timing error of L and H mode detections, and misses and false positives for ELMs. This step is taken in Section 4.5.1. Using this cost metric, several filter parameters are sequentially optimised in 1D parameter sweeps where all other parameters are kept constant at their common-sense values. The a priori transition probability p_0 and the process and measurement noise covariances are optimised to minimise the cost metric in Sections 4.5.2 and 4.5.3. Conclusions about the results of these parameter optimisations are drawn in Section 4.5.4. Moreover, Section 4.5.5 analyses whether both the kinetic measurements - the interferometer and diamagnetic loop - are ueful to the filter or whether one is enough.

This section's approach to optimising the filter parameters is still limited in the sense that not all relevant detector paramaters were optimised, the optimisation was not performed over the whole n-dimensional parameter space but only over a few 1D cross-sections, and it has only been performed for a set of 13 TCV shots selected from both the training and test datasets, rather than e.g. hundreds of shots. In this selection of shots, # 45106, 45107, 47008, 48305, and 48854 were excluded because the applied heating powers and resulting signal features were very similar to previous shots, to avoid giving too much weight to particular types of shots.

4.5.1 A cost metric for detection results

The cost metric C should increase by a set amount for every time step at which the detected mode does not correspond to the real mode in effect. Whether such a failure is caused by a delay in detection right after the transition, a transition that was missed entirely, or a faulty excursion into the wrong mode due to deceptive signal patterns does not matter.

Also, for every ELM that is missed (and every ELM that is faultily detected) the cost should increase. A weighting is chosen in which a missed ELM is equivalent to 100 time steps or 10 ms worth of wrong L or H-mode detections. This cost is calculated for all 13 shots which are not extremely similar to each other.

So for $N_{w,L/H}$ time steps of wrong L/H-mode detections and $N_{w,ELM}$ incorrect ELM detections across all 13 shots, the cost C is:

$$C = C_{L/H} + C_{ELM} = (N_{w,L/H} + 100N_{w,ELM})T_s. \tag{4.5}$$

Note the multiplication by the sample time T_s, which means that can be interpreted C as a time, which is 0 for perfect detection and rises as detection results deteriorate.

$N_{w,L/H}$ can be determined simply from the detected mode results and the visually found transition times. The true ELM times were determined as time steps at which a detector with common-sense parameters first chose the ELM model, while the actual mode was H-mode. These ELM times were visually checked and slightly corrected. The detected ELM times are the first time

steps at which a detector chooses the ELM model, having previously chosen the H-model. Faulty ELM detections during times when the wrong mode is detected are ignored, to avoid counting the mistake doubly. During times when H-mode is detected correctly, each detected ELM time must correspond within 0.2 ms to a true ELM time, and vice versa, else it is a faulty detection and contributes 1 to $N_{w,ELM}$.

The cost can be normalised by taking $\hat{C} = \frac{C}{N_{trans}}$ where N_{trans} is the total amount of transitions in all shots considered. In cases where the ELM cost contribution is low and the costs are dominated by timing errors, this makes the cost less dependent on how many shots are used or how many transitions occur in these shots. Note that the unit of the normalised cost is still ms; the cost is only normalised by the amount of transitions. When the two cost terms are analysed seperately, they can be naturally normalised to the average timing error per transition $\hat{C}_{L/H} = \frac{C_{w,L/H}}{N_{trans}}$ and the fraction of wrong ELM detections $\hat{C}_{ELM} = \frac{C_{w,ELM}}{N_{ELM}}$ with N_{ELM} the total number of actual ELMs.

4.5.2 Optimising the a priori event probability

The base transition probability p_0 determines, together with the normalised input power, the value of the off-diagonal elements in the transition probability matrix Π_k, see Section 3.2.1. p_0 represents the a priori probability that we assign to the occurence of an LHT or HLT. If p_0 is too low, the detector may miss transitions once a mode has established a high mode probability $\mu^{(i)}$. If p_0 is too high, unmodelled signal features which somewhat resemble a transition may be mistaken for a transition.

The common-sense value chosen for p_0 was 10^{-3}. In Figure 4.17, p_0 is varied between 5×10^{-7} and 0.02, and the cost of the resulting detection behaviour is plotted, split up into the average timing error per transition $\hat{C}_{L/H}$ in blue and the fraction of wrong ELM detections $\hat{C}_{ELM}$ in red. Typically, 5% of ELMs are missed. This is due to the 500 Hz low-pass pre-filter for the H_α signal, which makes the detector robust against faulty LHT and HLT detections due to fast and short signal features, but also filters out some high-frequent ELMs and some 'double ELMs'. Figure 4.17 shows that the ELM detection is quite unresponsive to changes in p_0. The cost associated with transition detection dominates the behaviour of the total cost, and is large for very high or very low values of p_0, as expected. Moreover, it has a minimum for $p_0 = 1 \times 10^{-5}$ and a plateau between 2×10^{-5} and 1×10^{-3}.

Upon inspection of the costs for individual shots, it appears that for most shots, the cost dependence on p_0 is quite mild for $1 \times 10^{-6} < p_0 < 1 \times 10^{-3}$. In fact, the cost behaviour of the dots in Figure 4.17 is dominated by the cost of only 4 discharges in the considered database, which are TCV discharges # 45104, 45105, 45109, and 47007. The two cost components for these 4 shots are also plotted in the figure as circles. Especially the timing error behaviour is well approximated by only these 4 discharges.

The cost behavior for these shots can be explained as follows. As p_0 is decreased, the detector becomes increasingly stricter in classifying features as transitions. This will lead to longer delays between a transition and its detection, since the new model needs to make more correct measurement predictions to counteract the lower transition probabilities from the previously correct model. For example, Figure 4.18 shows that if p_0 is lowered to 1×10^{-6}, a visually obvious LHT in shot 45104 is detected with a delay of 20 ms. If p_0 is decreased even further, events with mild signal features may be missed entirely. For example, Figure 4.19 shows that for $p_0 = 2 \times 10^{-6}$, an HLT with only a mild increase in H_α intensity is missed in shot 47007.
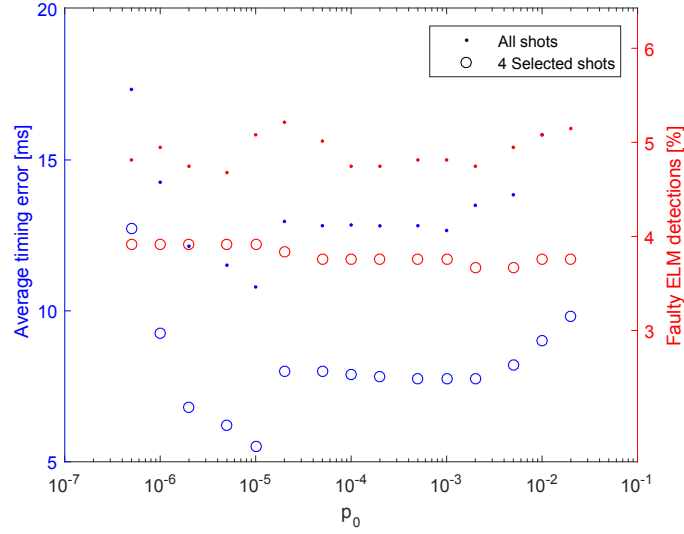
Figure 4.17: Normalised cost due to timing errors in LHT/HLT detections $\hat{C}_{L/H}$ (blue) and due to missed or false positive ELMs (red) $\hat{C}_{ELM}$, for all 13 shots (points) and 4 selected shots (circles), as a function of the base probability of a transition p_0. The y-axis limits are chosen such that the left and right vertical scale correspond exactly for the points (all shots), but not exactly for the circles (4 shots). The cost is smallest for $p_0 = 1 \times 10^{-5}$ and the cost behaviour for the 4 shots matches that for all shots.

On the other hand, if p_0 is increased too far, signal features which slightly resemble transitions may be wrongly detected as transitions, because the measurement predictions of the alternative model are better during a few time steps. For example, Figure 4.20 shows that for shot 45105, increasing p_0 to 0.02 leads to a false LHT detection. An H_α relaxation following a peak in L-mode is detected as an LHT, even though no significant change in the baseline is seen. When p_0 is decreased, the detector becomes more robust against such unmodeled signal features which temporarily resemble transitions. For example, Figure 4.21 shows that if p_0 is decreased to 1×10^{-5}, a signal feature shot 45109 which shortly shows a significant H_α increase but almost no change in the FIR or DML signals (and hence significantly resembles an HLT), is correctly rejected.

The results of this section show that the value $p_0 = 1 \times 10^{-5}$ is a good choice for the detector. Compared to the common-sense value of 1×10^{-3}, this may lead to slightly longer delays in transition detection for some shots, but in at least one case, allowed a misleading signal feature to be safely ignored, making the detector more robust. It should be noted that in the wider range $2 \times 10^{-6} \lesssim p_0 \lesssim 2 \times 10^{-3}$, the detector's cost is not very sensitive to p_0 for the investigated selection of discharges. Therefore, if the detector were applied to a wider range of shots in which previously unencountered signal features were present, some other value for p_0 in this wider range may prove to yield better detection results.
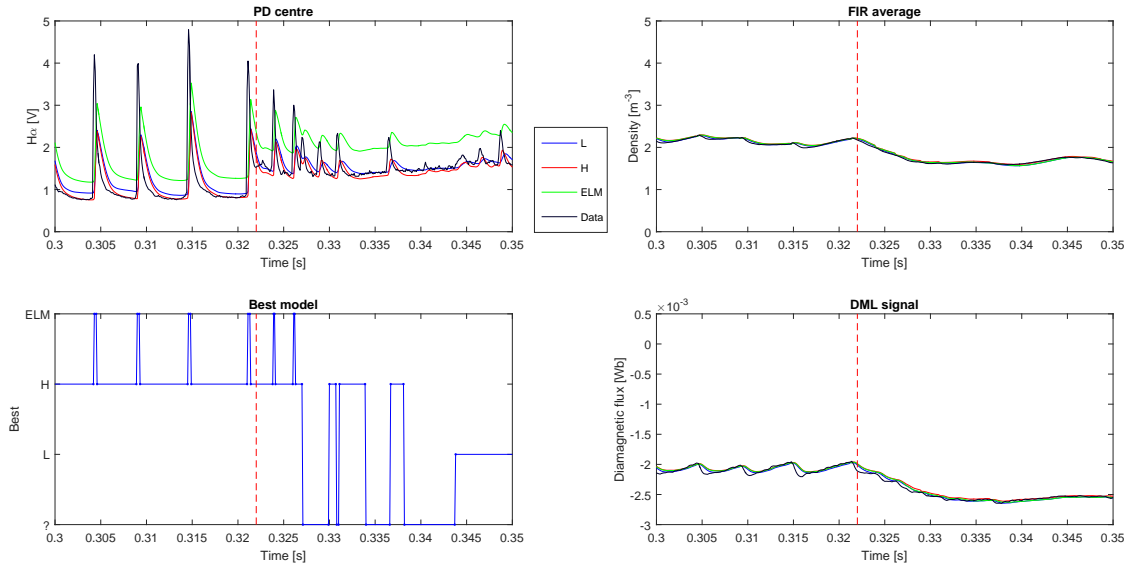
Figure 4.18: Detection result for shot 45104 with $p_0 = 1 \times 10^{-6}$. The LHT is not detected until 20 ms after the transition, due to the low transition probability.
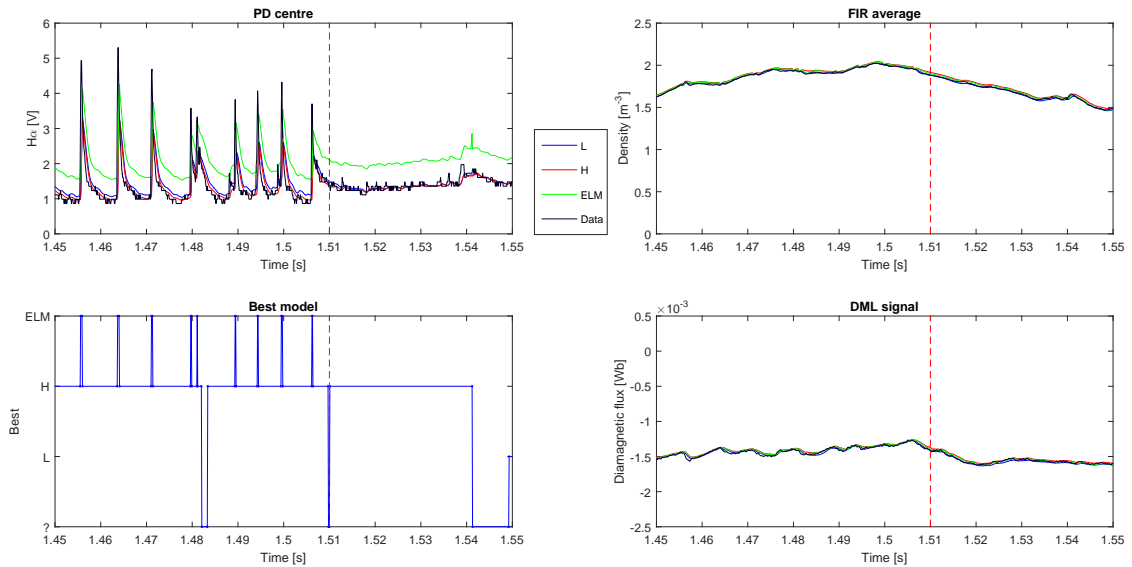

Figure 4.19: Detection result for shot 47007 with $p_0 = 2 \times 10^{-6}$. The final HLT for this shot shows only a modest increase in the H_α baseline. Due to the low transition probability, no HLT is detected. The detector stays in H-mode until the heating power is turned off at $1.54\,\text{s}$
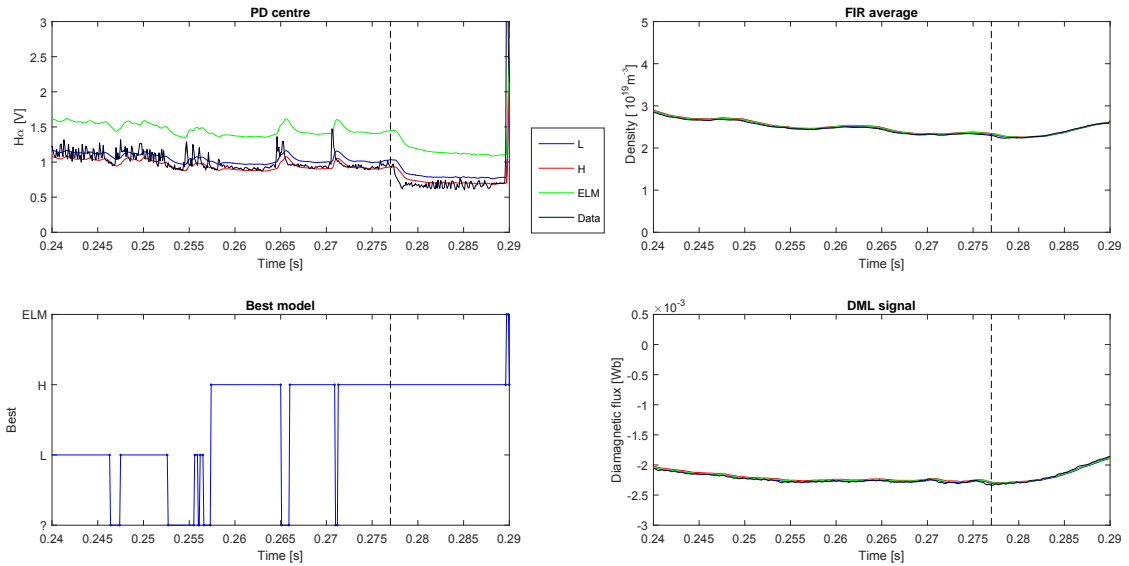
Figure 4.20: Detection result for shot 45105 with $p_0 = 0.02$. A signal relaxation after a small peak in the H_α signal, 20 ms before the real LHT, is wrongly detected as an LHT due to the high transition probability.
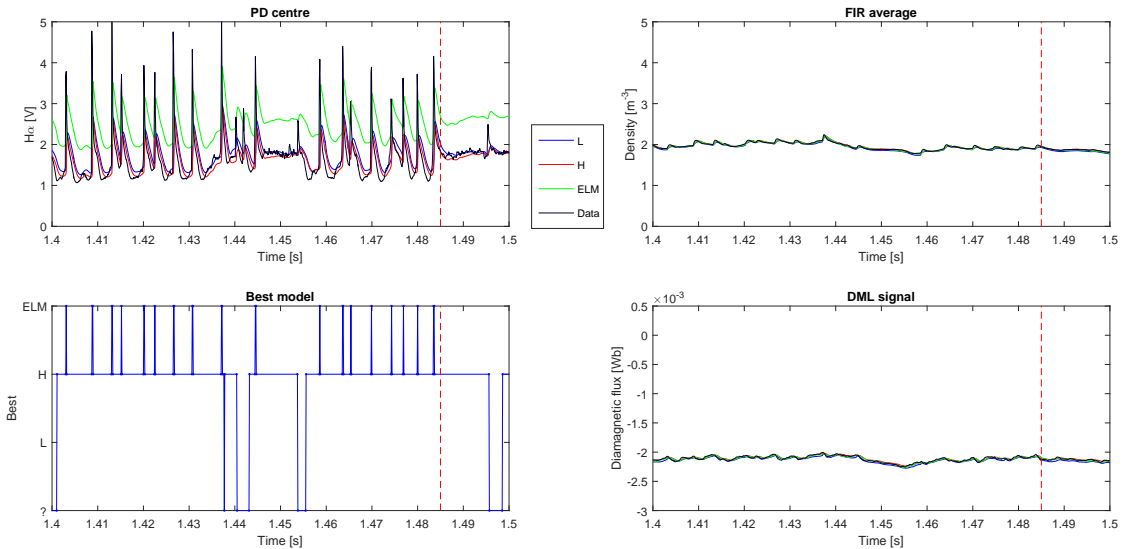

Figure 4.21: Detection result for shot 45109 with $p_0 = 1 \times 10^{-5}$. A misleading H_α relaxation towards a higher baseline is correctly rejected, thanks to the lack of corresponding signal features in the FIR and DML signals, the shortness of the feature, and the low transition probability.

4.5.3 Optimising the process noise and measurement noise covariances

The process noise covariance Q and the measurement noise covariance R influence many aspects of the detector's behaviour. For one, the ratio of Q to R determines the Kalman gain K, whose entries grow larger for larger Q. Because process noise adds to the estimation error covariance P every time step, while the measurement noise affects only the measurement on a single time step, the MM filter only performs well for $R \gg Q$. Just how much larger R is determines on what time scale the estimates converge to the measurements.

Moreover, when Q and R are both scaled by the same factor, the Kalman gain for the individual filters does not change, but the typical innovations covariance S is scaled by the same factor (aside from the effect of the initial estimation error covariance P_0). This means that the prior probability density distribution for the innovations z becomes broader (for higher S) or sharper (for lower S). Higher noise covariances and broader z distributions lead to smaller differences in the likelihoods L of different models, making the detector less susceptible to transitions and ELMs. Similar to the (inverse) effect of varying p_0, too large Q and R may cause the detector to miss transitions or ELMs, while too small Q and R may cause the detector to mistake unmodelled signal features for transitions or ELMs.

Lastly, because Q and R are diagonal and the system matrix F is a block diagonal, the dynamics of state variables (and estimation error covariances) underlying different measurements remain neatly separated throughout the detector operation. This means that by scaling all of those entries in Q and R relevant to a single measurement, the importance of that measurement when evaluating the model likelihoods L can be influenced without changing the Kalman gain for the associated states. Since $R \gg Q$, the 'weight' of each measurement signal is dominated by its entry in R compared to the typical scale of variations in the signal. If Q and R corresponding to one signal are decreased, the multi-variate gaussian distribution that is the prior distribution for z becomes sharper in the dimension corresponding to that signal, and the signal is given more weight in the probability update.

The common-sense values of Q and R are diagonal matrices

$$Q = \begin{bmatrix} Q_{\bar{n}_{H^0}} & & & & & \\ & Q_{n_{H^0}} & & & & \\ & & Q_{\dot{n}_e} & & & \\ & & & Q_{n_e} & & \\ & & & & Q_{\dot{\phi}_p} & \\ & & & & & Q_{\phi_p} \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} R_{H_\alpha} & & \\ & R_{FIR} & \\ & & R_{DML} \end{bmatrix}, \qquad (4.6)$$

with entries given in Table 4.5. In the table, c is the proportionality constant between the H_α measurement and the neutral density. I.e., the neutral density process noise $Q_{n_{H^0}}$ translates to $2 \times 10^{-3}\,\text{V}^2$ on the level of the measurements.

Optimising the scale of all noise covariances

In Figure 4.22, the ratio between each entry of R and Q is kept constant, but their common-sense values are multiplied by a constant R_0, which is varied between 0.6 and 1.5. The resulting normalised cost components $\hat{C}_{L/H}$ and $\hat{C}_{ELM}$ are plotted. The figure shows that $\hat{C}_{L/H}$ is minimum for $R_0 = 1.0$ (the common-sense noise covariances) and $\hat{C}_{ELM}$ is minimum for $R_0 = 1.05$. The total cost is smallest for $R_0 = 1$ by a small margin, and increases significantly for lower or higher R_0.

Table 4.5: Common-sense values for the entries of the process noise covariance Q and the measurement noise covariance R. $c \sim 1 \times 10^{-15}\,\text{V/m}^{-3}$ is the (arbitrary) proportionality constant between n_{H^0} and y_{H_α} defined in Section 3.2.

Symbol	Value	Unit
$Q_{\bar{n}_{H^0}}$	$1 \times 10^{-4}/c^2$	$(\text{m}^{-3})^2$
$Q_{n_{H^0}}$	$2 \times 10^{-3}/c^2$	$(\text{m}^{-3})^2$
$Q_{\dot{n}_e}$	1×10^{-1}	$(10^{19}\text{m}^{-3}/\text{s})^2$
Q_{n_e}	1×10^{-3}	$(10^{19}\text{m}^{-3})^2$
$Q_{\dot{\phi}_p}$	1×10^{-1}	$(\text{mWb/s})^2$
Q_{ϕ_p}	1×10^{-3}	$(\text{mWb})^2$
R_{H_α}	2.5×10^{-2}	V^2
R_{FIR}	5×10^{-2}	$(10^{19}\text{m}^{-2})^2$
R_{DML}	5×10^{-2}	$(\text{mWb})^2$
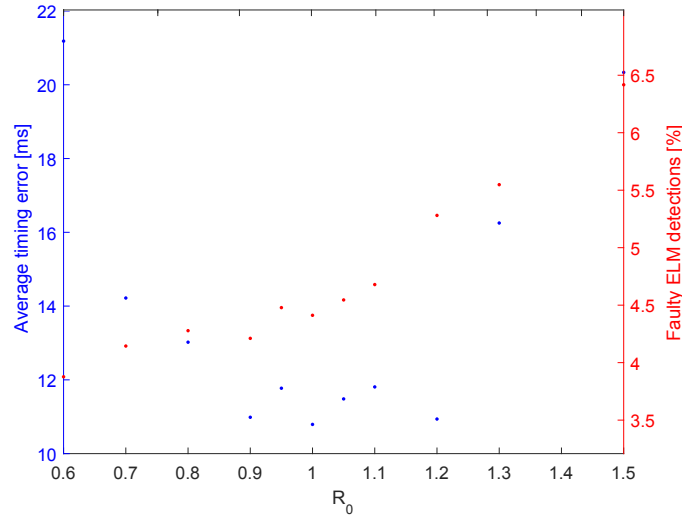
Figure 4.22: Normalised cost due to timing errors in LHT/HLT detections $\hat{Q}_{L/H}$(blue) and due to missed ELMs $\hat{Q}_{ELM}$ (red), for all 13 shots, as a function of the scaling factor for both noise covariances, R_0. The left and right vertical scale correspond exactly. The total cost is minimum at $R_0 = 0.9$.

For the individual shots, the detection results behave similarly to those for p_0. For example, the short unmodelled H_α signal increase in shot 45109 that caused a wrong L-mode detection for high p_0 is also wrongly detected for $R_0 \leq 0.8$. In Figure 4.23, the results for $R_0 = 0.8$ are shown. Also, the final HLT in shot 47007 featuring only a modest increase in the H_α baseline, which was missed for $p_0 \leq 5 \times 10^{-6}$, is also missed for $R_0 \geq 1.05$. The results for $R_0 = 1.05$ are shown in Figure 4.24. These results show that the scale of the noise covariances was already well-tuned, and does not need further improvement.
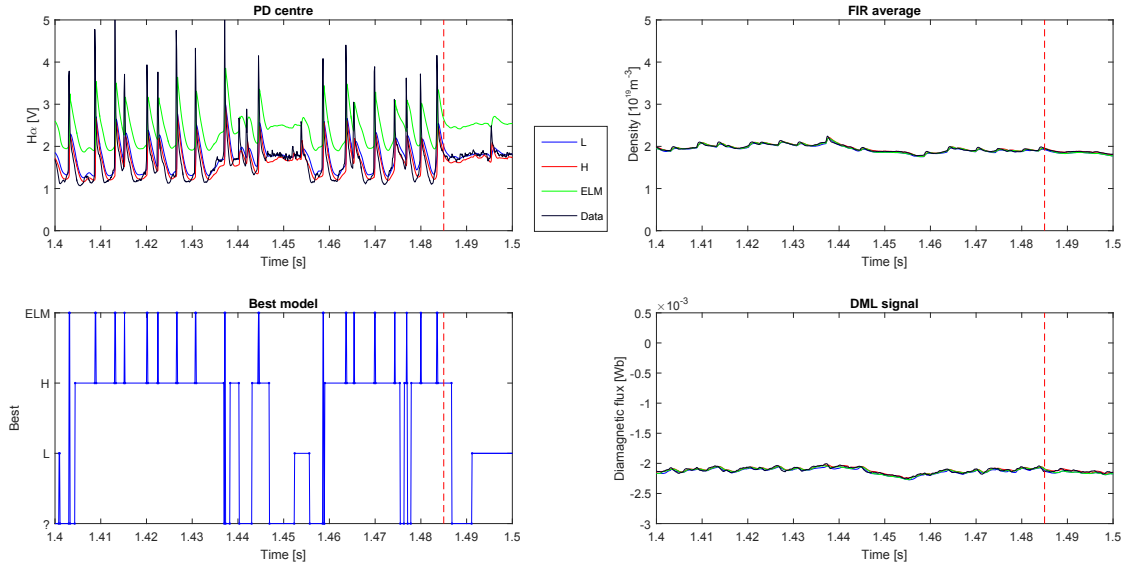
Figure 4.23: Results for shot 45109 with $R_0 = 0.8$. A misleading H_α signature is wrongly detected as H-mode, due to the low innovations covariance.
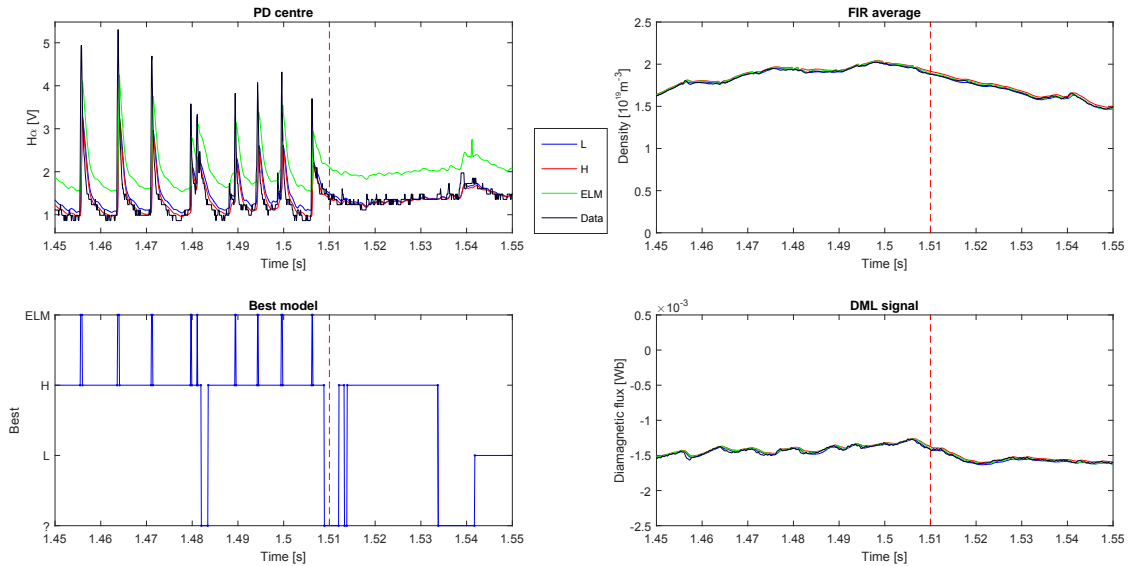
Figure 4.24: Results for shot 47007 with $R_0 = 1.05$. The final HLT is missed due to the high innovations covariance.
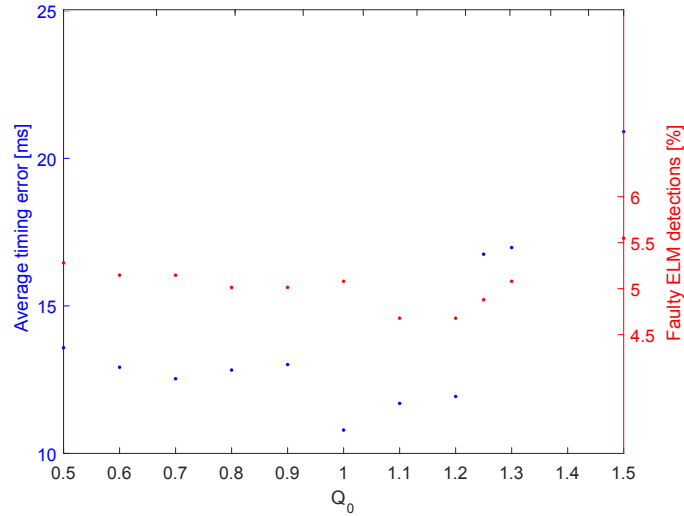
Figure 4.25: Normalised cost due to timing errors in LHT/HLT detections $\hat{Q}_{L/H}$(blue) and due to missed ELMs $\hat{Q}_{ELM}$ (red), for all 13 shots, as a function of the scaling factor for the process noise covariance, Q_0. The left and right vertical scale correspond exactly. The total cost is minimum at $Q_0 = 1.1$.

Optimising the Kalman gain

In Figure 4.25, the Kalman gain is varied by keeping R at its common-sense values (i.e., $R_0 = 1$) and multiplying Q by a scalar factor Q_0. The resulting normalised cost for all shots is plotted. Apparently, varying the Kalman gain has a smaller effect on the costs than varying both Q and R at the same time. Still, both costs increase for values of Q_0 far from 1. For Q_0 between 1 and 1.2, $\hat{C}_{L/H}$ decreases and $\hat{C}_{ELM}$ increases as a function of Q_0, but the total cost remains about the same. C is minimum for $Q_0 = 1.1$

For the individual shots, again similar behaviour is observed to the behaviour for varying p_0 or R_0. When Q (and hence the Kalman gain) is increased, it opens up the possibility of the dominant filter quickly adapting its estimates to new signal features, before the alternative filter that better describes the new measurements has reached a high mode probability. For example, Figure 4.26 shows that for $R_0 = 1.25$, the HLT in shot 45105 is missed because of the high Kalman gain. Conversely, a low Q and hence a low Kalman gain can cause the dominant filter to adapt only slowly to small unmodeled disturbances, allowing an alternative filter to grow a large mode probability. For example, Figure 4.27 shows that for $R_0 = 0.9$, the short increase in H_α baseline in shot 45109 is wrongly detected as an HLT. These results show that the Kalman gain was already well-tuned, and does not need to be adapted.

Optimising the weight of different measurements

In Figure 4.28, the weight of the H_α measurement is varied compared to the two kinetic measurements, by scaling the FIR and DML measurement noise covariances R_{FIR} and R_{DML} and the relevant process noise covariances by the same factor, and keeping R_{H_α} and the associated process noise covariance constant. This ensures the Kalman gain does not change. As can be seen from Table 4.5, the common-sense values for $\frac{R_{DML}}{R_{H_\alpha}}$ and $\frac{R_{FIR}}{R_{H_\alpha}}$ are 2. In Figure 4.28, this ratio is varied between 0.6 and 6, and the resulting normalised cost is plotted for 13 shots.
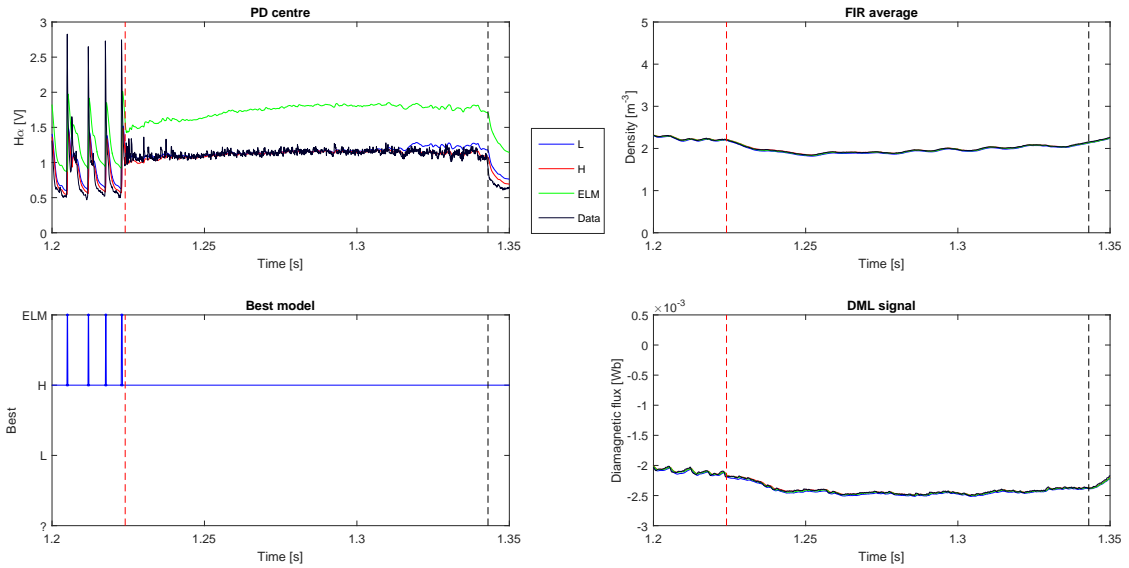
Figure 4.26: Results for shot 45105 with $R_0 = 1.25$. Due to the high Kalman gain, the HLT is missed because the H-mode filter can adapt its estimate quickly compared to how long it takes for the L-mode probability to grow, given that its expected measurements better match the real measurements.
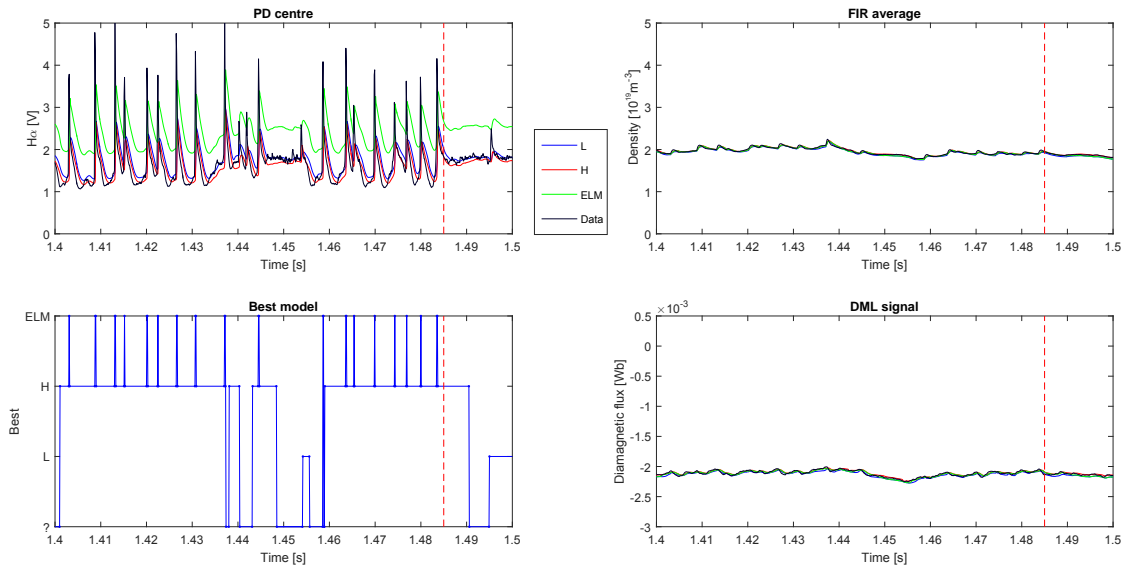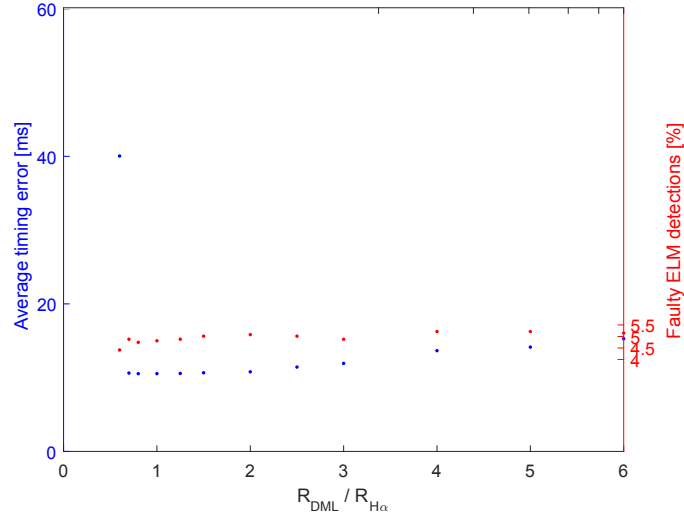

Figure 4.27: Results for shot 45109 with $R_0 = 0.9$. Due to the low Kalman gain, the unmodelled increase in H_α is detected as an HLT, because the L-mode probabilty grows faster than the H-mode filter can adapt its estimate to the measurements.

Figure 4.28: Normalised cost due to timing errors in LHT/HLT detections $\hat{Q}_{L/H}$ (blue) and due to missed ELMs $\hat{Q}_{ELM}$ (red), for all 13 shots, as a function of the ratio between different measurement noise covariances $\frac{R_{DML}}{R_{H_\alpha}} = \frac{R_{FIR}}{R_{H_\alpha}}$, where R_{H_α} and the Kalman gain are kept constant. The left and right vertical scale correspond exactly. The total cost is minimum at $\frac{R_{DML}}{R_{H_\alpha}} = 0.8$.

Figure 4.28 shows that the cost decreases with decreasing $\frac{R_{DML}}{R_{H_\alpha}}$ for $\frac{R_{DML}}{R_{H_\alpha}} \geq 0.8$, but increases dramatically at $\frac{R_{DML}}{R_{H_\alpha}} = 0.6$. For many shots, decreasing $\frac{R_{DML}}{R_{H_\alpha}}$ and hence increasing the weight of the kinetic signals allows the detector to detect transitions with a smaller delay. However, when the ratio becomes too low, some major errors occur. For example, Figure 4.29 shows that for shot 43454, the high weight on the kinetic signals causes ELMs to be detected as an HLT. Figure 4.30 shows that for shot 43454, an HLT is missed because its effect on the kinetic signals is smaller than expected. These results show that more emphasis can be put on the kinetic measurements. Therefore, the ratio $\frac{R_{DML}}{R_{H_\alpha}}$ is lowered to 1.
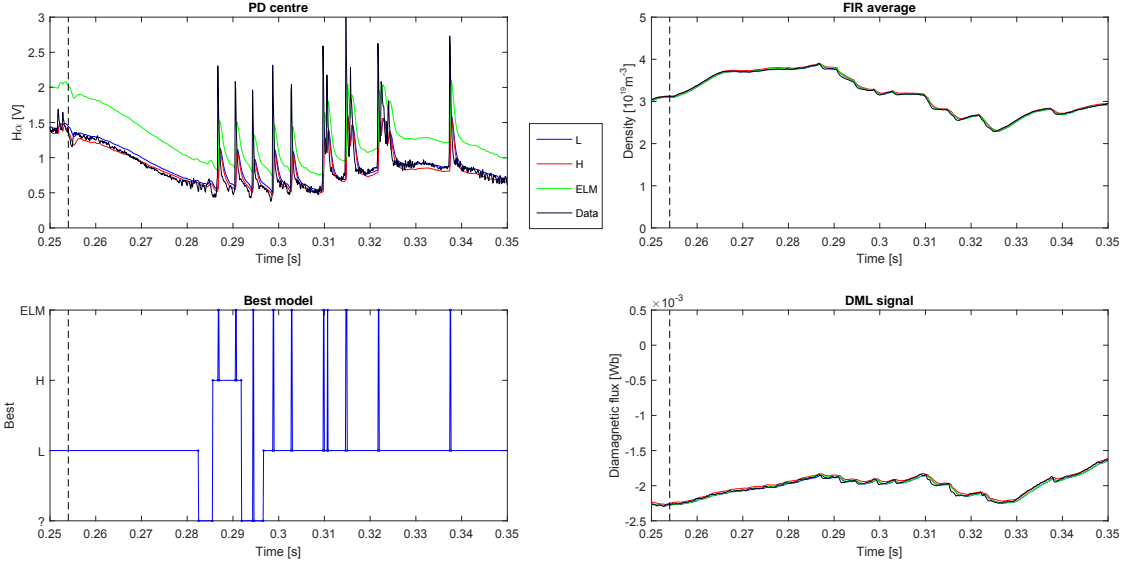
Figure 4.29: Results for shot 43454 with $\frac{R_{DML}}{R_{H_\alpha}} = 0.6$. Due to the high weights on the kinetic signals, the decreases in the DML and (especially) FIR signals due to ELMs are detected as an HLT, even though the H_α signal is described better by the H-mode model during the first few ELMs. This causes the detector to miss more than a second of H-mode.
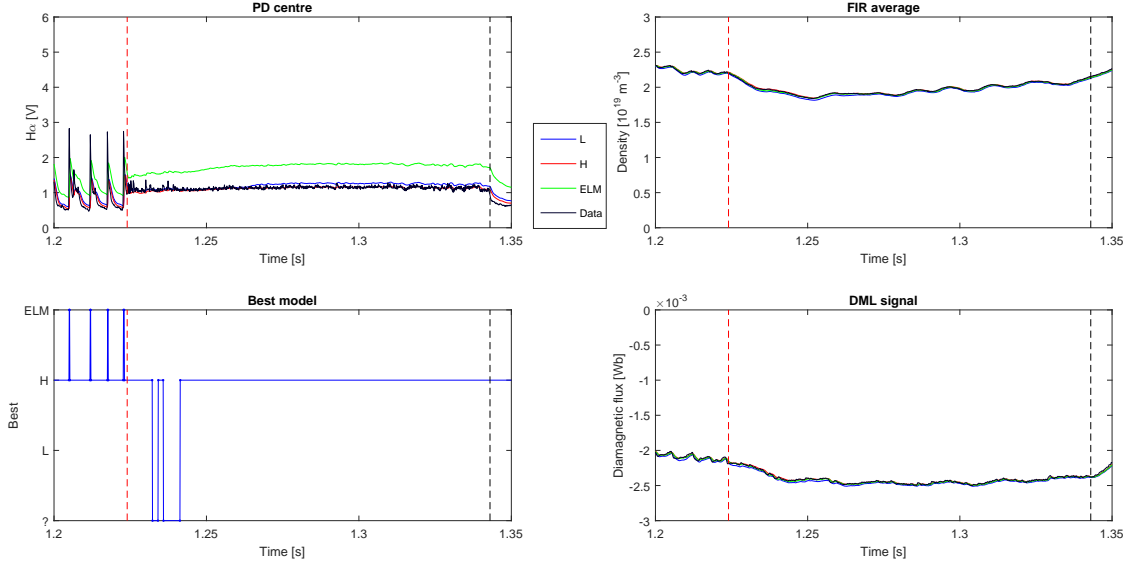
Figure 4.30: Results for shot 45105 with $\frac{R_{DML}}{R_{H_\alpha}} = 0.6$. The HLT is missed because the weights on the kinetic signals are high, and the decreases in these signals are smaller than the L-mode model expects. The derivatives decrease by about $2 \times 10^{20} \, \mathrm{m^{-3}/s}$ and $20 \, \mathrm{mWb/s}$ compared to the expected $5 \times 10^{20} \, \mathrm{m^{-3}/s}$ and $50 \, \mathrm{mWb/s}$.

4.5.4 Conclusions on parameter optimisation

Using an objective cost function, the performance of the multi-model filter was verified and some filter parameters were optimised to give the best performance on a series of 13 TCV discharges. For many of the filter paramaters, this optimisation constituted a trade-off between minimising the detection delay after a transition and rejecting deceptive signal features vaguely similar to the features to be detected. As a result of the optimisation, the a priori transition probability p_0 was decreased from 1×10^{-3} to 1×10^{-5} to make the detector more robust against such deceptive features. Moreover, all noise covariances relevant to the DML and FIR measurements were decreased by a factor 2 to give more weight to these measurements, to compensate for deceptive patterns in the H_α measurements. The scale of all covariances and the Kalman gain were shown to be already well-tuned.

4.5.5 The use of multiple kinetic measurements

The cost metric introduced in this section can also be used to assess whether it is useful to include both FIR and DML measurements in the detector. After all, these contain mostly the same information and may be redundant. In Figure 4.31, the resulting cost is plotted as function of whether the FIR and/or DML signals are used. The H_α signal is used in all cases. When neither the FIR nor the DML signal is used, the detector performs poorly, similar to the poor performance for very high ratio $\frac{R_{DML}}{R_{H_\alpha}}$ seen in the previous subsection. When at least one kinetic signal is used, the performance improves, but it does not improve further when both signals are used. The small differences between the 3 sets of results is likely within the range of minor parameter tuning, but this will not be adressed further.
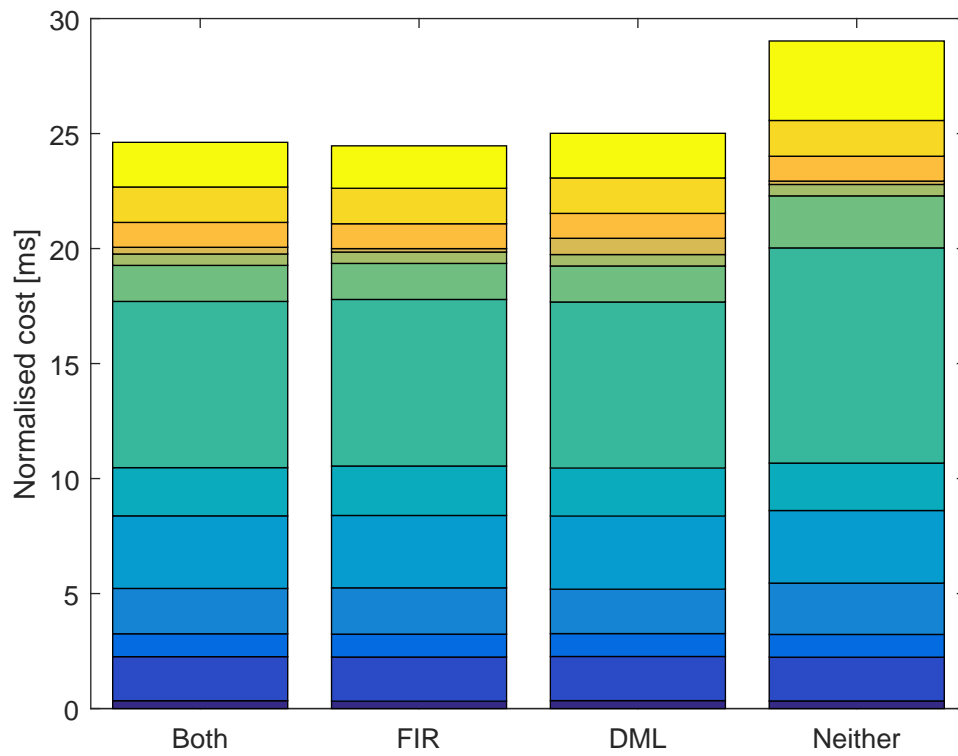
Figure 4.31: The per-shot cost contribitions to the total normalised cost when the FIR and/or DML signals are used in the estimator, additional to the H_α signal. The different color segments represent different shots. The detector performs significantly worse when neither the FIR nor the DML signal is used, but performance does not improve significantly by using both signals compared to only one.

Chapter 5

Conclusions and Outlook

The research question for this work was: *"can specific events in the TCV tokamak, particularly, L-H transitions and ELMs, be detected in real-time using a model-based approach?"*

The model-based detector developed in this work can detect L-H transitions and H-L transitions with an average timing error of 10 ms per transition, on a dataset of 18 discharges on the TCV tokamak with varying conditions. Each shot had at least one and up to four L-H-L cycle(s), and 14 out of 18 shots had one or more ELMy periods. The timing error includes time delay between the actual transition and its detection, but also false positives where particular signal features were wrongly identified as transitions. Moreover, 95% of ELMs were detected correctly, where mistakes were caused by low-pass pre-filtering of the H_α measurement signal. In fact this pre-filtering constitutes a trade-off between the robustness of transition detection against short but sharp changes in the H_α signal, and the detection of high-frequency ELMs. The loss of plasma internal energy during ELMs was also analysed by using the ELM timing from the detector and the plasma diamagnetic flux measurement. The diamagnetic flux is linked to the plasma internal energy using LIUQE equilibrium reconstruction data, in a way that is accurate as long as LIUQE's sample frequency is much faster than changes in the plasma's flux surface geometry (technically, in the outcome of the Shafranov integrals) and plasma current. At a 1 kHz sample frequency, these conditions are readily fulfilled.

The first sub-question was *"which diagnostic signals have to be taken into account to detect the events?"*. It has been shown that the H_α signal and the ratio of applied heating power to the L-H threshold power from a scaling law are sufficient to detect most L-H transitions and back-transitions in the considered dataset of 18 discharges. The addition of either the DiaMagnetic Loop (DML) signal or a combination of Far InfraRed interferometer (FIR) signals can make the detection more robust to transitions with an unclear H_α signature and to unmodeled signal features, but the addition of both the FIR and DML signals does not further improve detection results for the considered dataset.

The second sub-question was *"how should the evolution of the underlying plasma parameters producing the measured signals be modelled during different modes of operation?"*. The detection results have been obtained by using a Cooperating Multiple Model estimator, consisting of simple models which directly produce the signal behaviour that is actually observed during transitions. For the H_α signal, this means the signal will suddenly tend to a lower/higher baseline (for L-H/H-L transitions) and that it will peak during ELMs. For the FIR and DML signals, this means a sudden increase/decrease in the signal's time-derivative.

The third sub-question was *"which hypothesis reduction strategy should be used to combine the state estimates based on different models?"* Of the different hypothesis reduction strategies that were used to share and combine the state estimates of the different filters, the hybrid decision strategy which combined continous single-step decisions with occasional multi-step decisions (both with the 1-best algorithm) performed best. However, the time-varying computational requirements make such an algorithm unsuitable for real-time implementation, and it was shown that a slightly modified single-step IMM hypothesis reduction strategy gave similarly good results. These minor modifications consisted of a conservative time delay in the reported mode and changes to the ELM→L and ELM→H transition probabilities at each ELM event.

Futhermore, the detector performance has been objectively assessed using a cost function that contains terms for both wrong L or H-mode detections and wrong ELM detections. This analysis showed that the rough, common-sense optimisation of many parameters already gave near-optimal costs in single-paramater sweeps. However, it was found that decreasing the prior probability of a transition and giving more weight to the FIR and DML measurements could lead to a significant increase in robustness against unmodeled signal disturbances.

A number of possibilities to improve and use the detector were not adressed in this work. As an outlook toward possible future research, a few will now be listed.

Low-pass pre-filtering of the H_α measurement signal allows a trade-off between having robust transition detection and complete ELM detection. However, ideally, detection of L-H and H-L transitions should be unhindered by fast signal dynamics resembling a transition only on a short scale, while fast ELMs should still be detected. To this end, two versions of the H_α signal could be used, where one is pre-filtered using a low-pass filter and the other is used as-is.

The detector may be able to make better use of the sequence in which signal features occur. If a signal feature which resembles an L-H transition is observed, and shortly afterwards a series of high-amplitude, constant-frequency ELMs starts, this should increase the probability that the original signal frequency was indeed an L-H transition. Similarly, if after a possible H-L transition, small and high-frequent H_α blips appear which quickly decrease in frequency and size before disappearing, this should validate the H-L transition detection.

Some H-L transitions which occur near the end of the shot, during the current ramp-down phase, are missed because their signal features are not very clear. This is likely due to other signal features in the current ramp-down phase which obscure the HLT features. Perhaps the detection of these transitions can be improved if the detector directly uses measurements (or a reference signal) of the plasma current to more easily allow HLT detections when the plasma current drops significantly below its established flat-top value.

Some signal features which were difficult to detect correctly were likely related to the Intermediate phase [31]. In some shots, such a phase occurs right before the plasma enters H-mode, and causes oscillations in the turbulence levels, which for example cause large noise in the H_α signal. If signal features relating to the Intermediate phase are shown to behave reproducibly, these can perhaps be described by simple models similar to the L/H/ELM models, and added to the filter.

The detector parameter optimisation presented in this work is incomplete both in the sense that not all relevant parameters were tuned, and that all parameters were tuned independently rather than simultaneously across the n-dimensional parameter space. In the future, some more extensive parameter optimisation could be considered.

Finally, the detector should be implemented on the real-time control system for the TCV tokamak. This will allow monitoring of the ELM frequency and amplitude, which can be used for feedback control of the transient heat load due to ELMs when actuators such as non-axisymmetric resonant magnetic perturations or fuel pellet injection are used to trigger small ELMs. Such work will be a good preparation for future devices like ITER, where large ELM heat loads constitute a major threat to the life-time of the divertor and the trade-off between good confinement and low (transient) divertor heat loads is more critical than ever.

Real-time implementation of the detector also allows monitoring of transitions between L and H-mode. An attractive mode of operation for future devices like ITER is to trigger a transition to H-mode and then lower the heating power to below the L-H power threshold, staying just above the H-L power threshold so that the transport bifurcation allows the plasma to stay in H-mode. In this case, an underestimate of the H-L power threshold will lead to an undesired back-transition into L-mode. If such an accidental H-L transition is detected by the detector, extra heating power can be applied until the plasma again transitions to H-mode. Preparatory experiments on TCV for such a scenario would also be valuable for ITER.

Bibliography

[1] Mitsuru Kikuchi, Karl Lackner, and Minh Quang. *Fusion physics*. IAEA, 2012. ISBN 978-92-0-130410-0. 1

[2] T Hamacher and a M Bradshaw. Fusion as a future power source: recent achievements and prospects. *Proceedings of the World Energy Congress*, (October):1–17, 2001. URL http://fire.pppl.gov/energy{_}eu{_}wec01.pdf. 1

[3] N. Cruz, J-M. Moret, S. Coda, et al. An optimal real-time controller for vertical plasma stabilization. 2014. doi: 10.1109/TNS.2015.2478965. URL http://arxiv.org/abs/1406.4436http://dx.doi.org/10.1109/TNS.2015.2478965. 1

[4] J Stober, L Barrera, K Behler, et al. Feedback-controlled NTM stabilization on AS-DEX Upgrade. *European Physical Journal Conference*, 7:1–6, 2015. doi: 10.1051/epjconf/20158702017. 1

[5] T C Blanken, F Felici, M R De Baar, et al. Modeling , observer design and robust control of the particle density profile in tokamak plasmas. (Cdc):7628–7635, 2015. ISSN 07431546. doi: 10.1109/CDC.2015.7403425. 1

[6] Justin E Barton, William P Wehner, Eugenio Schuster, et al. Simultaneous Closed-loop Control of the Current Profile and the Electron Temperature Profile in the TCV Tokamak. *Acc 2015*, pages 3316–3321, 2015. 1

[7] D Kalupin, M Z Tokar, B Unterberg, et al. Predictive modelling of L and H confinement modes and edge pedestal characteristics. *Nuclear Fusion*, 45(6):468, 2005. ISSN 0029-5515. doi: 10.1088/0029-5515/45/6/008. URL http://stacks.iop.org/0029-5515/45/i=6/a=008. 3

[8] W. Weymiens, S. Paquay, H. J. de Blank, et al. Comparison of bifurcation dynamics of turbulent transport models for the L-H transition. *Physics of Plasmas*, 21(5):052302, 2014. ISSN 1070-664X. doi: 10.1063/1.4871856. URL http://scitation.aip.org/content/aip/journal/pop/21/5/10.1063/1.4871856. 3

[9] Jan Weiland. Simulations of the L-H transition on experimental advanced superconducting Tokamak. *Physics of Plasmas*, 21(12):122501, 2014. ISSN 1070-664X. doi: 10.1063/1.4901597. URL http://scitation.aip.org/content/aip/journal/pop/21/12/10.1063/1.4901597. 3

[10] F Wagner. A quarter-century of H-mode studies. *Plasma Physics and Controlled Fusion*, 49 (12B):B1–B33, 2007. ISSN 0741-3335. doi: 10.1088/0741-3335/49/12B/S01. 3, 8

[11] S González, J Vega, a Murari, et al. Automatic location of L/H transition times for physical studies with a large statistical basis. *Plasma Physics and Controlled Fusion*, 54(6):065009, 2012. ISSN 0741-3335. doi: 10.1088/0741-3335/54/6/065009. URL http://stacks.iop.org/0741-3335/54/i=6/a=065009?key=crossref.fc9308ceada9e3d53001c8d5dc01fa92. 3, 7

[12] S Pamela, T Eich, L Frassinetti, et al. Non-linear MHD simulations of ELMs in JET and quantitative comparisons to experiments. *Plasma Physics and Controlled Fusion*, 58(1):014026, 2016. ISSN 0741-3335. doi: 10.1088/0741-3335/58/1/014026. URL http://stacks.iop.org/0741-3335/58/i=1/a=014026?key=crossref.817ac800c3246c7874dcb00f9e79bbb9. 3

[13] M Keilhacker, A Gibson, C Gormezano, et al. High fusion performance from deuterium-tritium plasmas in JET. *Nuclear Fusion*, 39:209–234, 1999. ISSN 0029-5515. doi: 10.1088/0029-5515/39/2/306. 5

[14] K. Ikeda. Progress in the ITER Physics Basis. *Nuclear Fusion*, 47(6):9–11, 2007. ISSN 0029-5515. doi: 10.1088/0029-5515/47/6/E01. 5

[15] John Wesson. *Tokamaks*. Clarendon Press, Oxford, 3rd edition, 2004. 5, 6

[16] F Wagner, F Ryter, A R Field, et al. Recent Results of {H}-Mode Studies on {ASDEX}. *Plasma Physics and Controlled Fusion Research (Proc. {13^{th}} Int. Conf., Washington, U.S.A.)*, 1:277, 1991. 6

[17] Y R Martin and TCV Team. ELMing H-mode Accessibility in Shaped TCV Plasmas. In *Proc. 18th IAEA Fusion Energy Conference*, 2000. URL http://www-pub.iaea.org/mtcd/publications/pdf/csp{_}008c/pdf/exp5{_}30.pdf. 6

[18] J. Freidberg. *Plasma Physics and Fusion Energy*. Cambridge University Press, Cambridge, 1 edition, 2007. ISBN 9780521851077. 6

[19] A Pitzschke. Pedestal Characteristics and MHD Stability of H-Mode Plasmas in TCV. 4917, 2011. 7

[20] Y R Martin and TCV Team. H-mode threshold power in TCV Ohmic plasmas. *Plasma Physics and Controlled Fusion*, 44:A143–A150, 2002. 7

[21] Y R Martin and TCV Team. Synchronization of L-mode to H-mode transitions on the sawtooth cycle in Ohmic TCV plasmas. *Plasma Physics and Controlled Fusion*, 46(5A):A77–A85, 2004. ISSN 0741-3335. doi: 10.1088/0741-3335/46/5A/008. URL http://stacks.iop.org/0741-3335/46/i=5A/a=008?key=crossref.86d2294a2fc0a1388623f76c15b0bc39. 7

[22] ITER physics basis. Chapter 2: Plasma confinement and transport. *Nuclear Fusion*, 39(12):2175–2249, 2002. ISSN 0029-5515. doi: 10.1088/0029-5515/39/12/302. 7

[23] Y.R. Martin, T. Takizuka, and ITPA CDBM H-mode Threshold Database Working Group. Power requirement for accessing the H-mode in ITER. *Journal of Physics: Conference Series*, 012033(123), 2008. doi: 10.1088/1742-6596/123/1/012033. 7, 42

[24] Jw Connor and Hr Wilson. A review of theories of the L H transition. *Plasma physics and controlled fusion*, 1, 2000. URL http://iopscience.iop.org/0741-3335/42/1/201. 8

[25] H. Biglari, P. H. Diamond, and P. W. Terry. Influence of sheared poloidal rotation on edge turbulence. *Physics of Fluids B: Plasma Physics*, 2(1):1, 1990. ISSN 08998221. doi: 10.1063/1.859529. URL http://scitation.aip.org/content/aip/journal/pofb/2/1/10.1063/1.859529. 8

[26] P H Diamond, S-I Itoh, K Itoh, et al. Zonal flows in plasmaa review. *Plasma Physics and Controlled Fusion*, 47(5):R35–R161, 2005. ISSN 0741-3335. doi: 10.1088/0741-3335/47/5/R01. 8

[27] P. De Vries. *Magnetic Islands in Tokamak Plasmas*. PhD thesis, Universiteit Utrecht, ISBN 90-3931731-3, 1997. 8

[28] P Gohil, G R McKee, D Schlossberg, et al. Dependence of the H-mode power threshold on toroidal plasma rotation in the DIII-D tokamak. *Journal of Physics: Conference Series*, 123: 012017, 2008. ISSN 1742-6596. doi: 10.1088/1742-6596/123/1/012017. 8

[29] Eun-jin Kim and P. H. Diamond. Zonal Flows and Transient Dynamics of the L - H Transition. *Physical Review Letters*, 90(18):185006, 2003. ISSN 0031-9007. doi: 10.1103/PhysRevLett. 90.185006. 8

[30] A H Nielsen, G S Xu, J Madsen, et al. Simulation of transition dynamics to high confinement in fusion plasmas. pages 1–13, 2014. ISSN 03759601. doi: 10.1016/j.physleta.2015.10.004. 9

[31] G S Xu, H Q Wang, M Xu, et al. Dynamics of L H transition and I-phase in EAST. 54, 2014. doi: 10.1088/0029-5515/54/10/103002. 9, 76

[32] G S Xu, B N Wan, H Q Wang, et al. First Evidence of the Role of Zonal Flows for the L - H Transition at Marginal Input Power in the EAST Tokamak. *Physical Review Letters*, 125001 (September):1–5, 2011. doi: 10.1103/PhysRevLett.107.125001. 9

[33] G. D. Conway, C. Angioni, F. Ryter, et al. Mean and oscillating plasma flows and turbulence interactions across the L-H confinement transition. *Physical Review Letters*, 106(6):1–4, 2011. ISSN 00319007. doi: 10.1103/PhysRevLett.106.065001.

[34] T. Estrada, C. Hidalgo, T. Happel, et al. Spatiotemporal structure of the interaction between turbulence and flows at the L-H transition in a toroidal plasma. *Physical Review Letters*, 107 (24):3–7, 2011. ISSN 00319007. doi: 10.1103/PhysRevLett.107.245004.

[35] L. Schmitz, L. Zeng, T.L. Rhodes, et al. The role of zonal flows and predatorprey oscillations in triggering the formation of edge and core transport barriers. *Nucl. Fusion*, 54(7):073012, 2014. ISSN 0029-5515. doi: 10.1088/0029-5515/54/7/073012. URL http://stacks.iop.org/0029-5515/54/i=7/a=073012?key=crossref.568a89d1219d521817eb4ff189d82222. 9

[36] Dan Simon. *Optimal State Estimation*. John Wiley & Sons, Hoboken, New Jersey, 1st edition, 2006. 13, 14, 15

[37] Brian D. O. Anderson, John B Moore, and Mansour Eslami. *Optimal Filtering*. Prentice Hall, Englewood Cliffs, New Jersey, 1st edition, 1982. 16

[38] A. Doucet and C. Andrieu. Iterative algorithms for state estimation of jump Markov linear systems. *IEEE Transactions on Signal Processing*, 49(6):1216–1227, 2001. ISSN 1053587X. doi: 10.1109/78.923304. URL http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=923304. 16

[39] X.R. Li and V.P. Jilkov. Survey of maneuvering target tracking. part v: multiple-model methods. *IEEE Transactions on Aerospace and Electronic Systems*, 41(4):1255–1321, 2005. ISSN 0018-9251. doi: 10.1109/TAES.2005.1561886. URL http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1561886. 16, 17, 19

[40] J. M. Moret, F. Buhlmann, and G. Tonetti. Fast single loop diamagnetic measurements on the TCV tokamak. *Review of Scientific Instruments*, 74(11):4634–4643, 2003. ISSN 00346748. doi: 10.1063/1.1614856. 21, 36

[41] Loïc Curchod. *High Density Plasma Heating in the Tokamak à Configuration Variable*. PhD thesis, EPFL PhD thesis 5012, 2011. 21, 22, 23, 24

[42] V D Shafranov. Determination of the parameters β_I and l_i in a Tokamak for arbitrary shape of plasma pinch cross-section. *Plasma Physics*, 13(9):757, 1971. ISSN 0032-1028. doi: 10.1088/0032-1028/13/9/006. 22

[43] B.~P. Duval, A Bortolon, A Karpushov, et al. Bulk plasma rotation in the TCV tokamak in the absence of external momentum input. *Plasma Physics and Controlled Fusion*, 49: 195, 2007. ISSN 0741-3335. doi: 10.1088/0741-3335/49/12B/S18. URL http://cdsads. u-strasbg.fr/abs/2007PPCF...49..195D. 23

[44] J. M. Moret, B. P. Duval, H. B. Le, et al. Tokamak equilibrium reconstruction code LIUQE and its real time implementation. *Fusion Engineering and Design*, 91(2015):1–15, 2015. ISSN 09203796. doi: 10.1016/j.fusengdes.2014.09.019. 24, 25

[45] T P Donaldson, Abramovich M (Eds), Stegun I A, et al. Theory of foil-absorption techniques for plasma X-ray continuum measurements. *Plasma Physics*, 20(12):1279–1289, dec 1978. ISSN 0032-1028. doi: 10.1088/0032-1028/20/12/005. URL http://stacks.iop.org/ 0032-1028/20/i=12/a=005?key=crossref.bf33a8f325f0ea5be6a523b476af403d. 38

[46] T C Luce. An analytic functional form for characterization and generation of axisymmetric plasma boundaries. *Plasma Physics and Controlled Fusion*, 55(9):095009, sep 2013. ISSN 0741-3335. doi: 10.1088/0741-3335/55/9/095009. URL http://stacks.iop.org/ 0741-3335/55/i=9/a=095009?key=crossref.9ed4aa4893102add5119acd80d784590. 44

[47] F Piras, S Coda, I Furno, et al. Snowflake divertor plasmas on TCV. *Plasma Physics and Controlled Fusion*, 51(5):055009, 2009. ISSN 0741-3335. doi: 10.1088/0741-3335/51/ 5/055009. URL http://stacks.iop.org/0741-3335/51/i=5/a=055009?key=crossref. a281e046d21242eba91cedfc60587859. 44

Appendix A

Event detection results for the training dataset

In this appendix, the detection results of the discharges in the training dataset are presented, for the best performing hypothesis reduction strategy. The hypothesis reduction strategy is described in Section 4.2.4, and its results are summarised in Table 4.2. The correct detection of LHTs and HLTs can be seen by observering that shortly after each black vertical dashed line (visually identified LHT), the detector chooses H-mode, and shortly after each red dashed line (visually identified HLT), the detector chooses L-mode.
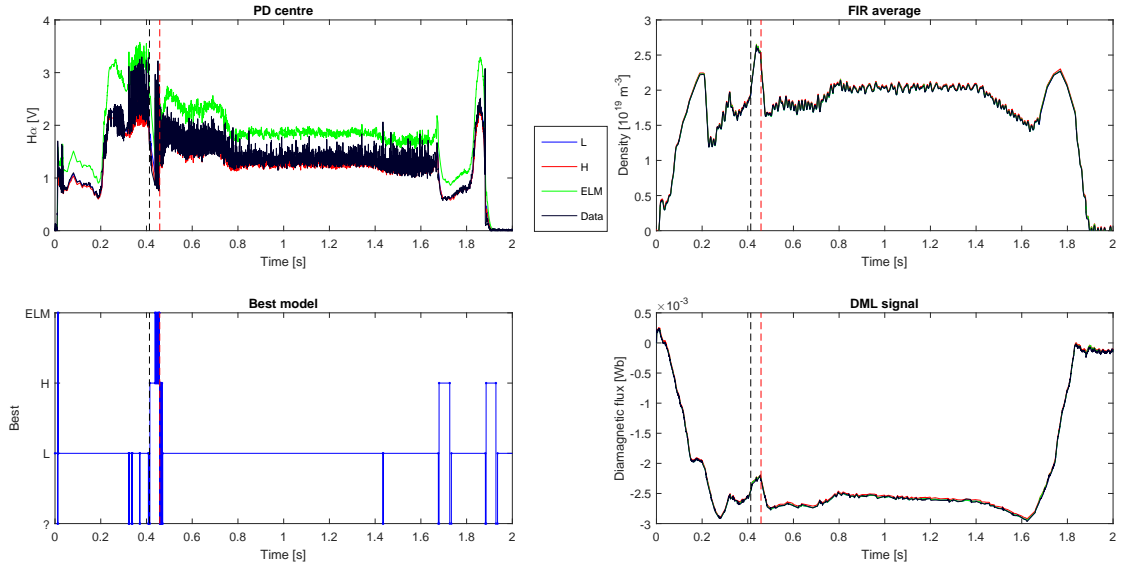
Figure A.1: Detector results for shot 39872. All features up until the current ramp-down are detected correctly.
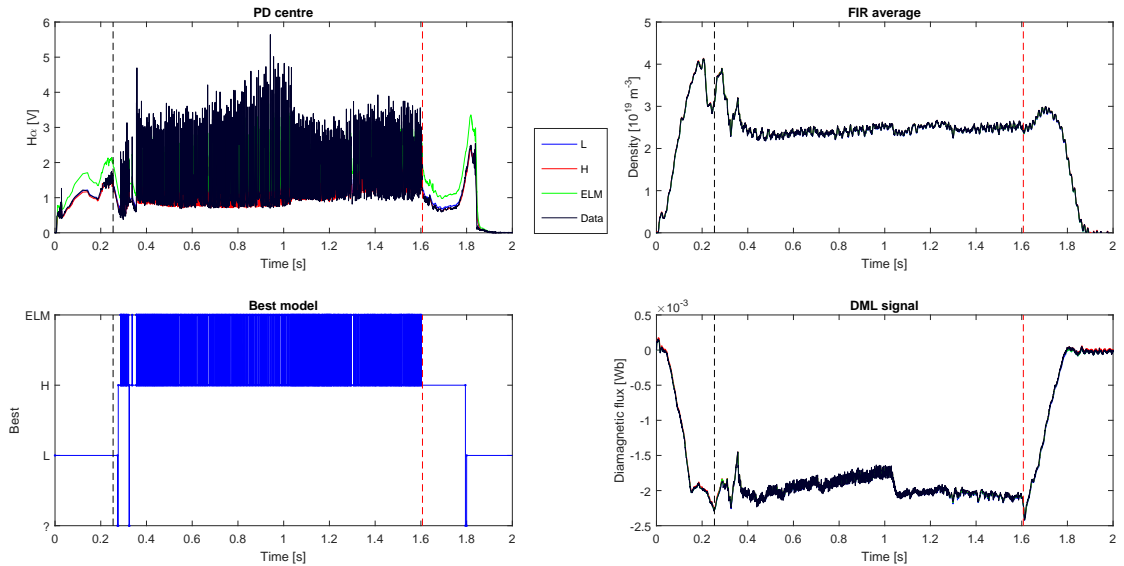
Figure A.2: Detector results for shot 43454. All features up until the current ramp-down are detected correctly.
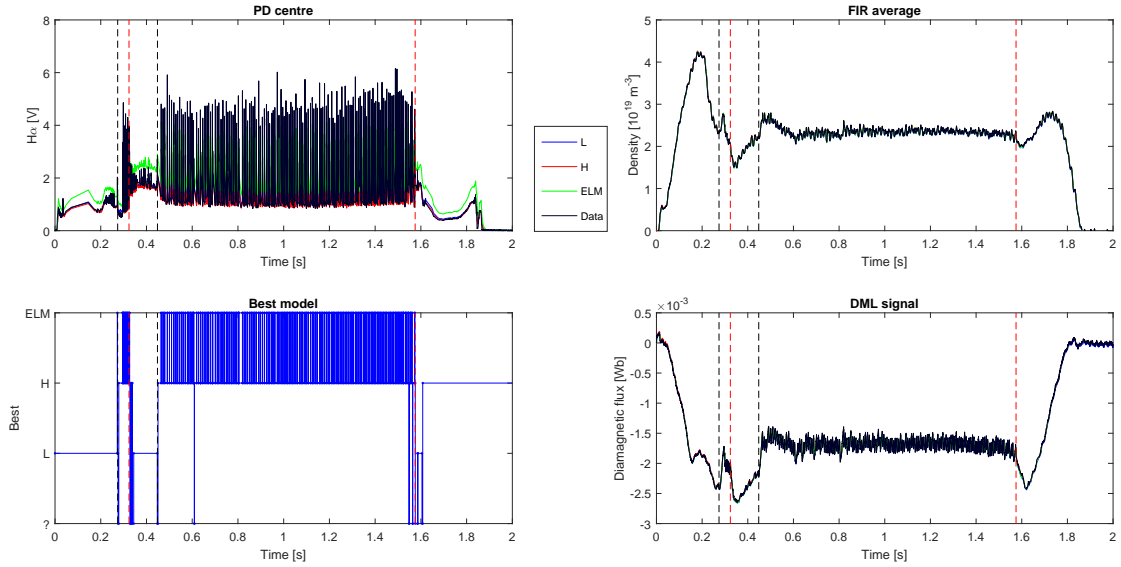
Figure A.3: Detector results for shot 45103. All features up until the current ramp-down are detected correctly.
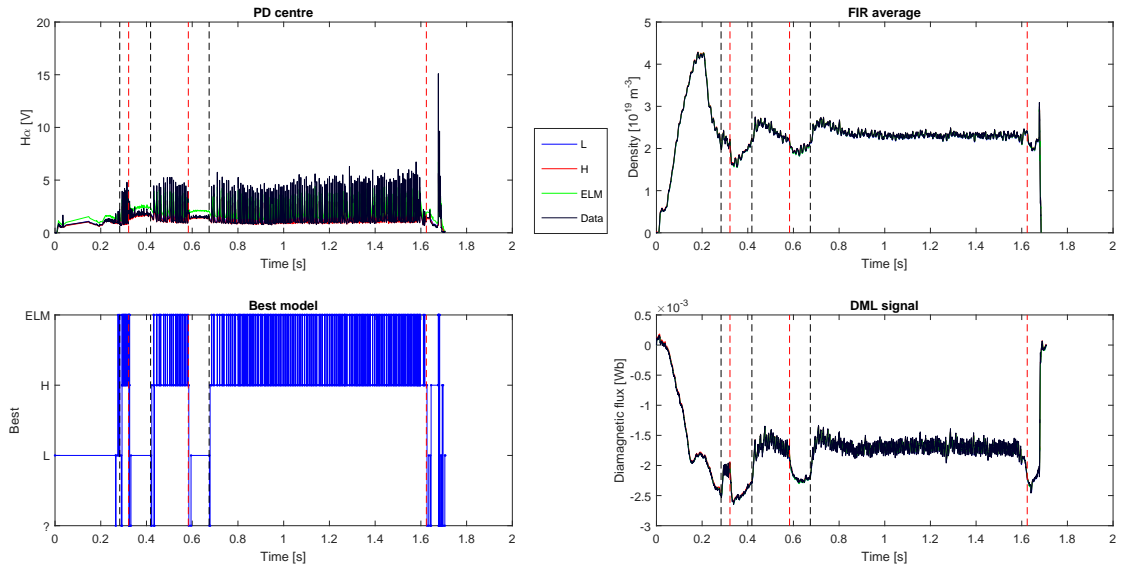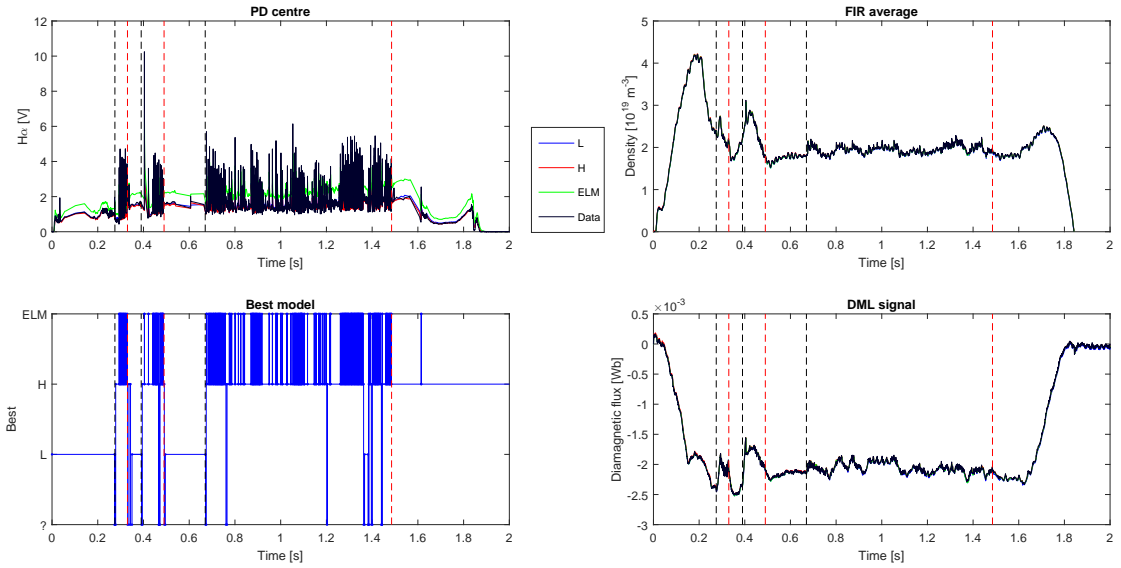

Figure A.4: Detector results for shot 45104. All features up until the current ramp-down are detected correctly.
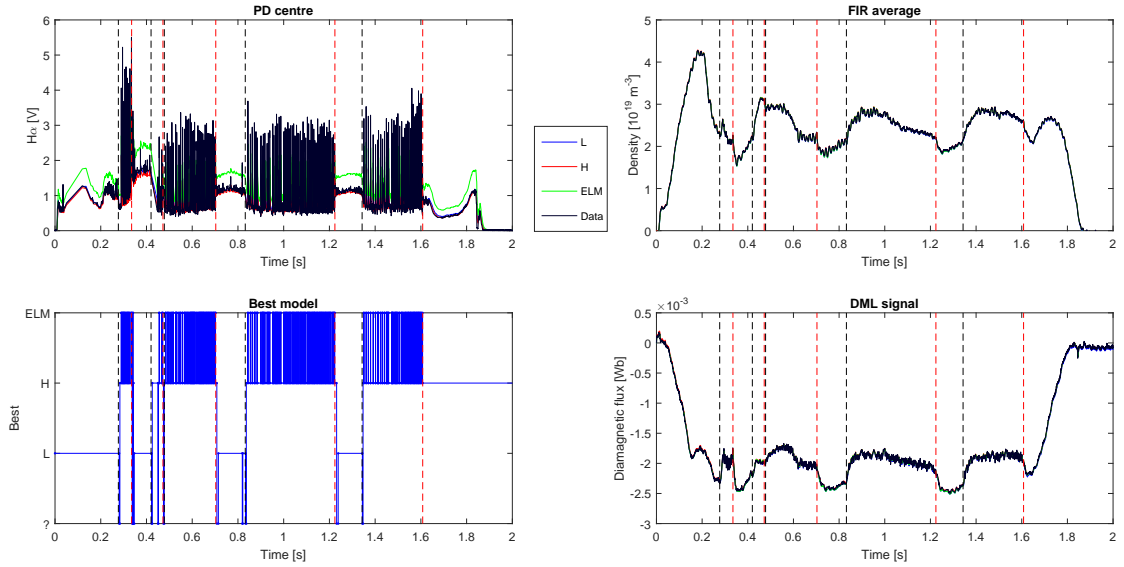
Figure A.5: Detector results for shot 45105. All features up until the current ramp-down are detected correctly.
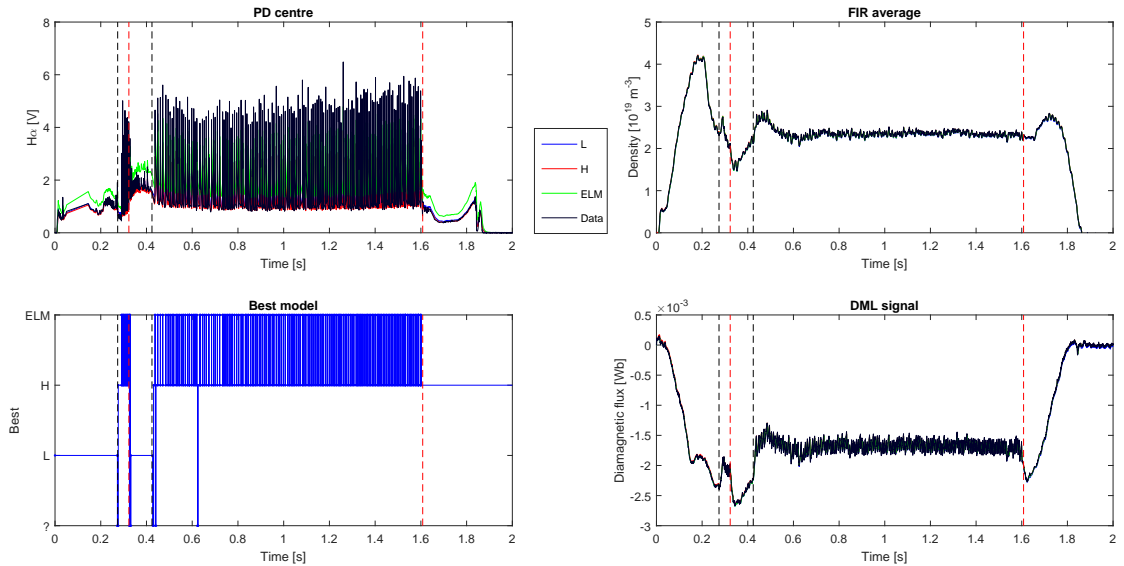
Figure A.6: Detector results for shot 45106. All features up until the current ramp-down are detected correctly.
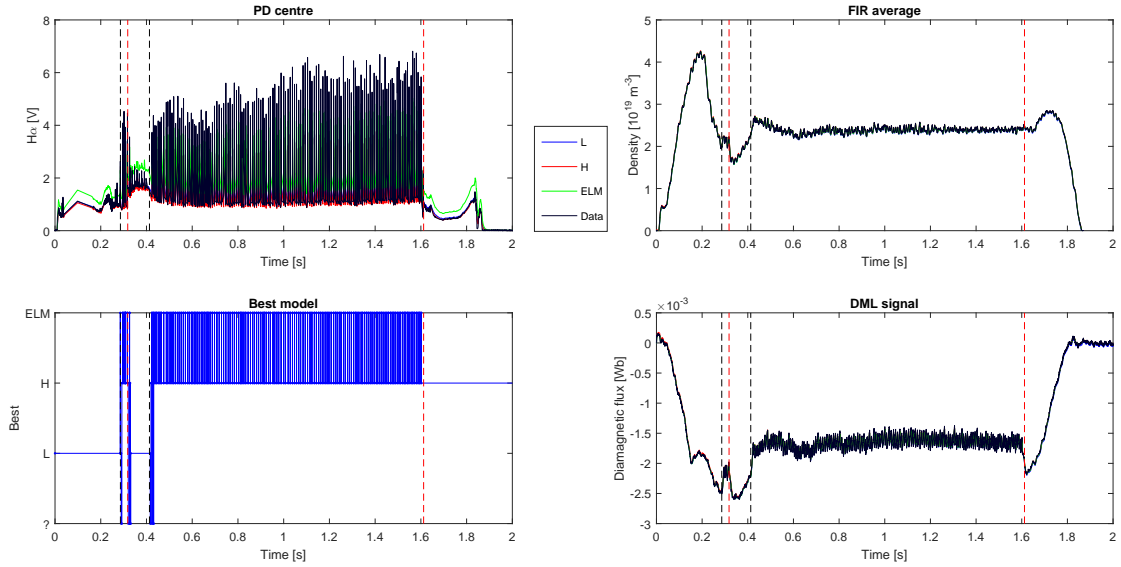
Figure A.7: Detector results for shot 45107. All features up until the current ramp-down are detected correctly.
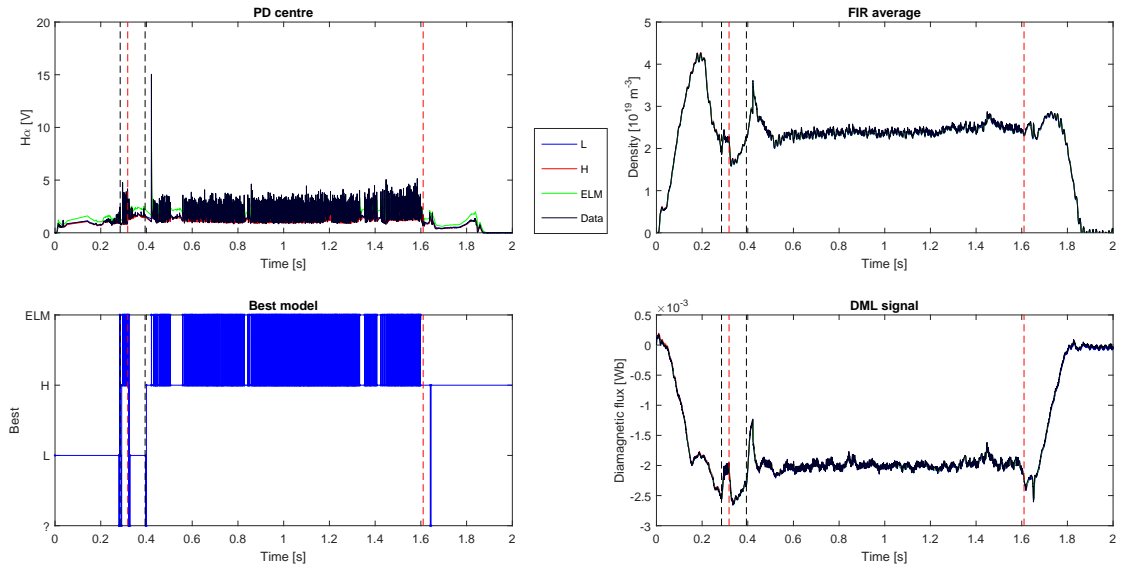
Figure A.8: Detector results for shot 45108. All features up until the current ramp-down are detected correctly.

Figure A.9: Detector results for shot 45109. At 1.37 s, an irregular H_α signal causes a faulty L-mode detection.
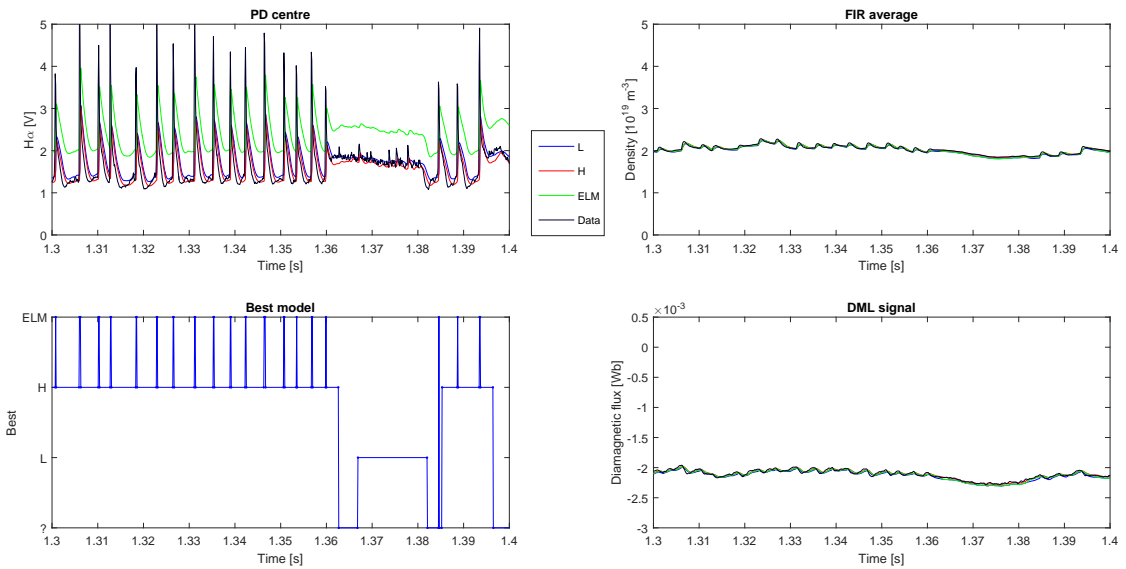
Figure A.10: Zoom-in of the results for shot 45109. The H_α signal relaxes to a higher baseline and is affected by small dithers (suggestive of an HLT), while the FIR and DML signals are almost unaffected (suggestive of no HLT). The heating power is actually increasing at this moment, so an HLT seems unlikely. The detector choses L-mode and thus is probably incorrect at this moment.
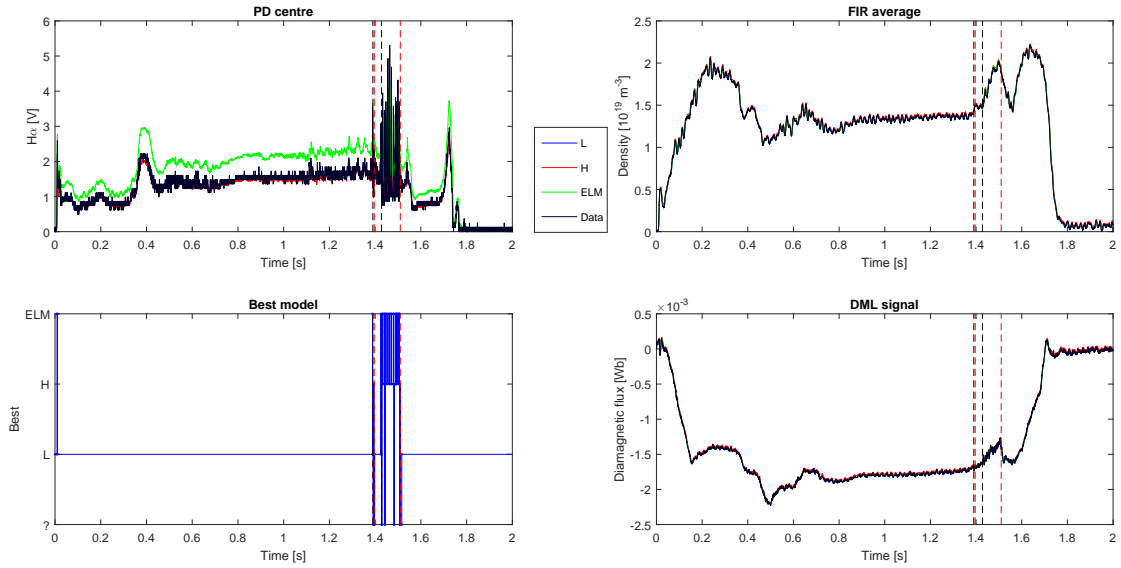
Figure A.11: Detector results for shot 47007. All features are detected correctly.
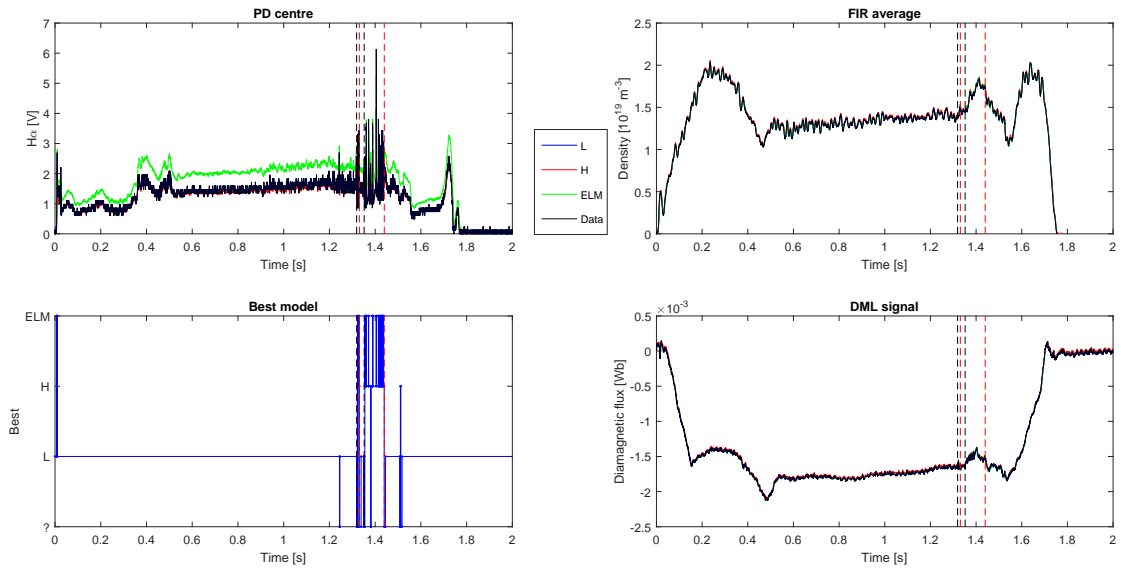
Figure A.12: Detector results for shot 47008. An irregular H_α signal at $1.51\,\mathrm{s}$ is wrongly detected as H-mode.
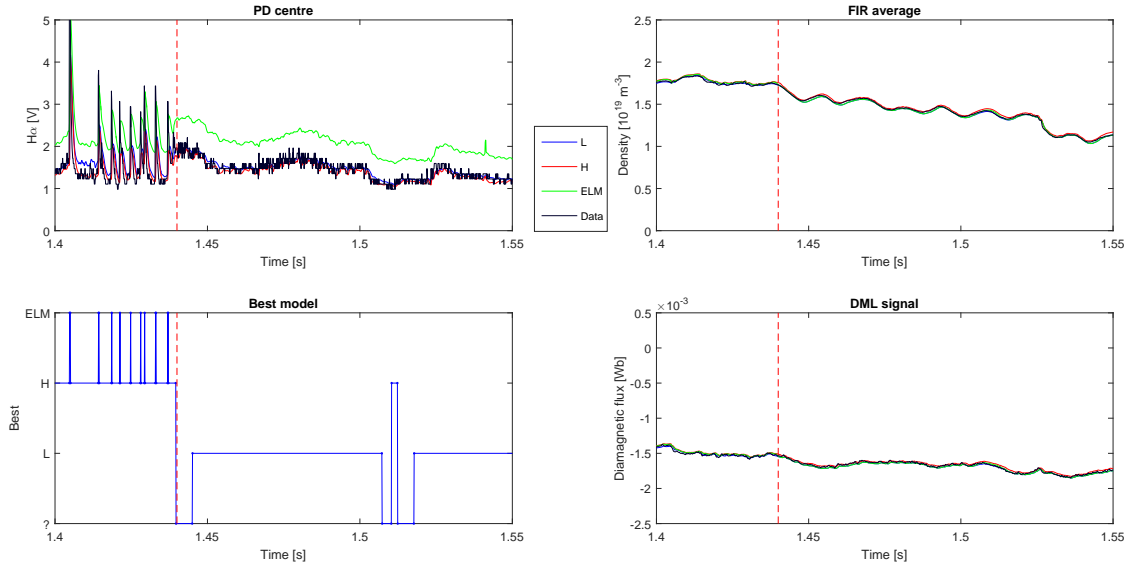
Figure A.13: Zoom-in of the results for shot 47008. After the HLT, a mild decrease in the H_α signal which is possibly related to the current ramp-down is mistaken for an LHT.
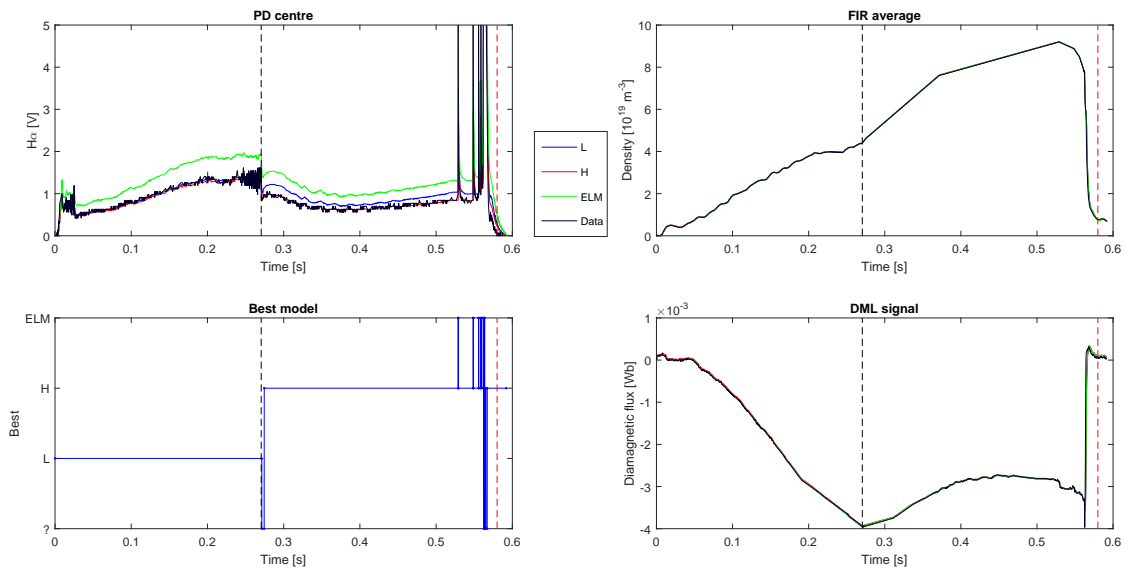

Figure A.14: Detector results for shot 47306. All signal features up until the current ramp-down are detected correctly.
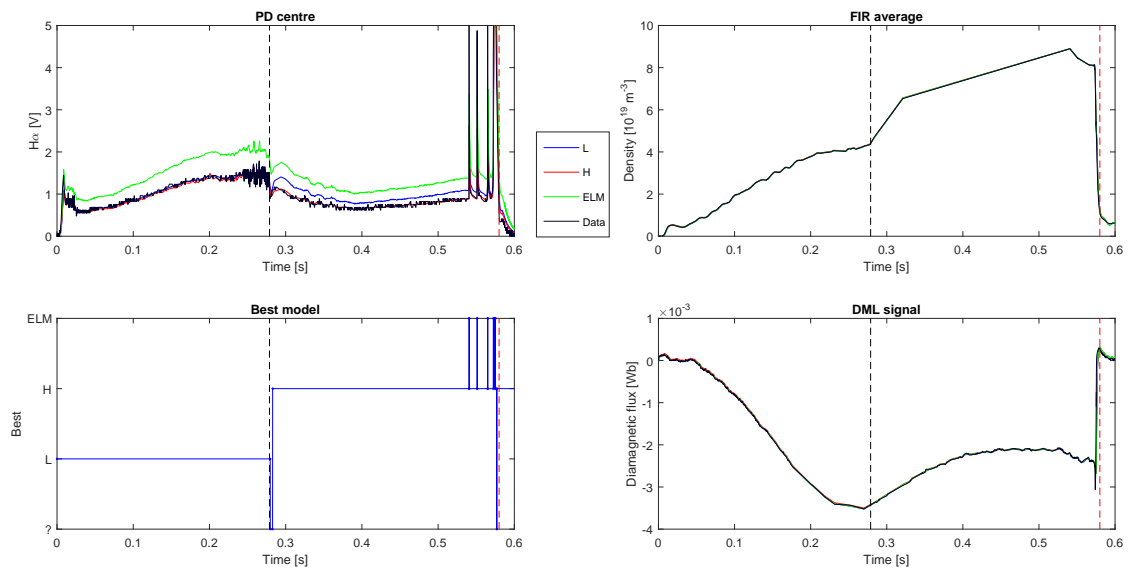
Figure A.15: Detector results for shot 47307. All signal features up until the current ramp-down are detected correctly.

Appendix B

Disturbance rejection in the interferometer signal

The TCV Far Infra-Red interferometry (FIR) signal is known to contain disturbances due to a 60-70Hz mechanical oscillation in the FIR diagnostic mounting. This disturbance is present in all 14 interferometer channels available on TCV and is of course specific to the TCV tokamak, as it depends on the way the interferometry system is mounted. This oscillation in the signal is not caused by fluctuations of the plasma electron density, and poses a significant disturbance to estimation of the plasma density on TCV. In order to estimate the plasma density, the mechanical oscillation must be filtered out of the signal. In this appendix, a linear oscillator model is formulated and added to the system of model equations on which the Multi-model Kalman filter is based. Ideally, the disturbances frequency, amplitude and phase should all be known so the effect can be filtered out.

As seen in Eq. 3.10, the model for the density evolution and measurements (not including mode transitions) is:

$$x = \begin{bmatrix} \dot{n}_e \\ n_e \end{bmatrix}, \quad \frac{d}{dt} x = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} x, \quad y = \begin{bmatrix} 0 & 1 \end{bmatrix} x, \tag{B.1}$$

Where simply a linear increase or decrease in the measured density is assumed. To compensate for the oscillation, two oscillating state variables are added, one 'visible' (i.e., directly measured), x_v, and the other invisible (i.e., not directly measured), x_i. Also, an angular frequency state variable ω is added. In an analogy with a simple mechanical harmonic oscillator, x_v is the position (which is typically measured) while ωx_i is the velocity (not typically measured). The angular frequency is modeled as constant: $\frac{d}{dt}\omega = 0$. The oscillation variables follow the dynamics:

$$\frac{d}{dt} \begin{bmatrix} x_v \\ x_i \end{bmatrix} = \frac{d}{dt} \begin{bmatrix} \omega x_i \\ -\omega x_v \end{bmatrix} \tag{B.2}$$

Combining the two lines of Eq. B.2 yields $\frac{d^2}{dt^2} x_v = i\omega^2 x_v$; as intended, x_v and x_i form a harmonic oscillator at frequency ω. This particular choice of oscillating state variables and dynamics was made over alternatives including $\frac{d}{dt} \begin{bmatrix} x_v \\ x_i \end{bmatrix} = \frac{d}{dt} \begin{bmatrix} x_i \\ -\omega^2 x_v \end{bmatrix}$, which lead to the same combined dynamics for $\frac{d^2}{dt^2} x_v$. The present choice of state variables has the advantage that x_v and x_i have the same amplitude and unit. The most natural way to assign process noise to the two oscillating states is in such a way that the oscillation of the amplitude is equally affected by

process noise in the two states. For the present choice of state variables, this can be easily done by choosing the fourth diagonal entry of the process noise matrix equal to the third. Note that no state variable is needed for the amplitude. At each point in time, $\sqrt{x_v^2 + x_i^2}$ gives the amplitude of the oscillation and the initial amplitude is set by the root-sum-square of the initial conditions for x_v and x_i. An extra amplitude state would be redundant, which goes against the state-space formalism.

The total dynamics are described by:

$$x = \begin{bmatrix} \dot{n}_e \\ n_e \\ x_v \\ x_i \\ \omega \end{bmatrix}, \quad \frac{d}{dt}x = \begin{bmatrix} 0 \\ \dot{n}_e \\ \omega x_i \\ -\omega x_v \\ 0 \end{bmatrix}, \quad y = n_e + x_v = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 \end{bmatrix} x. \tag{B.3}$$

Now, the system evolution is non-linear, because of the appearance of the products ωx_v and $-\omega x_i$ in Eq. B.3. Hence, the system can be observed using an extended Kalman filter (EKF), which linearises the non-linear system dynamics around the previous posterior state estimate $\hat{x}_{k-1|k-1}$ to obtain the system matrix F which is used in the time-update of the estimation error covariance. The time-update of the state itself is performed with the non-linear dynamics (Eq. B.3).

An extended Kalman filter based on the dynamics described above was tested on a pure sine wave (ideal disturbance, no relevant signal) with frequency 60 Hz, amplitude $1.5 \times 10^{18} \, \mathrm{m}^{-3}$, and initial phase 22.5°, while the filter started with initial conditions $\omega_0 = 2\pi \cdot 70 \, \mathrm{rad/s}$, $x_{v,0} = x_{i,0} = 0$. A measurement noise covariance $R = 1 \times 10^3$ is used (same as in the multi-model case), and the process noise covariance

$$Q = \begin{bmatrix} \sigma_{\dot{n}_e}^2 & & & & \\ & \sigma_{n_e}^2 & & & \\ & & \sigma_{x_v}^2 & & \\ & & & \sigma_{x_i}^2 & \\ & & & & \sigma_\omega^2 \end{bmatrix} = \begin{bmatrix} 0.1 & & & & \\ & 2 \times 10^{-4} & & & \\ & & 5 \times 10^{-5} & & \\ & & & 5 \times 10^{-5} & \\ & & & & 30 \end{bmatrix}, \tag{B.4}$$

where $\sigma_{\dot{n}_e}^2$ and $\sigma_{n_e}^2$ also correspond to the multi-model case.

Figure B.1 shows that within four oscillations, the filter's estimate has converged closely to the real disturbance. In reality, the disturbance will not look like a perfect sine wave and especially its amplitude may change over time. In figure B.2, the filter is applied to data from shot 45108, since in this shot the effect of the disturbance was large compared to most other shots. Specifically, the FIR data from the last 0.14 s of the shot is used. At this time, the shot itself is over and the real plasma density is 0, such that the measured signal consists of only the disturbance. Also in this case, after about 4 oscillations, the disturbance estimate has converged quite closely to the measurements. However, the varying amplitude causes some tracking error of the disturbance itself. This likely cannot be eliminated without sacrificing sensitivity to 'real' signal changes in the frequency range close to the frequency of the oscillation, e.g., those due to mode transitions.

Next, it is investigated how the filter behaves when applied to real FIR data. Figure B.3 shows the results for the first 0.5 s of TCV shot 45108. Some changes had to be made for the filter to work on this data. Specifically, the frequency process noise covariance σ_ω^2 was reduced from 30 to 0.1. In the figure, it can be seen that the frequency estimate does not change strongly due to the frequency of the disturbance, but is affected significantly by 'real' changes in the signal, such as the drop in density after divertor formation at 0.2 s. If the covariance is kept at 30 (in principle large enough for fast convergence to the real frequency), the frequency estimate is driven unstable
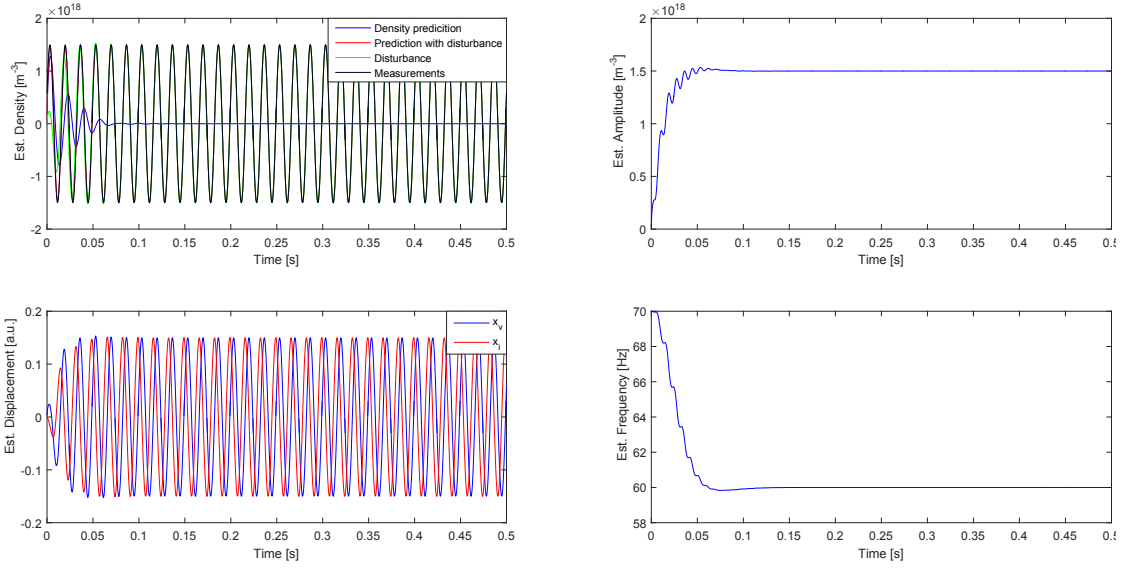
Figure B.1: Disturbance rejection for the FIR signal, tested on a pure 60 Hz sine wave with initial phase 22.5°. Plotted are the estimated density measurement (top-left), oscillation amplitude (top-right), oscillation frequency (bottom-right), and the visible (blue) and invisible (red) displacement states (bottom-left). The filter starts at 70 Hz with an amplitude of 0. The estimated amplitude converges sufficiently close to 0 within 0.1 s.
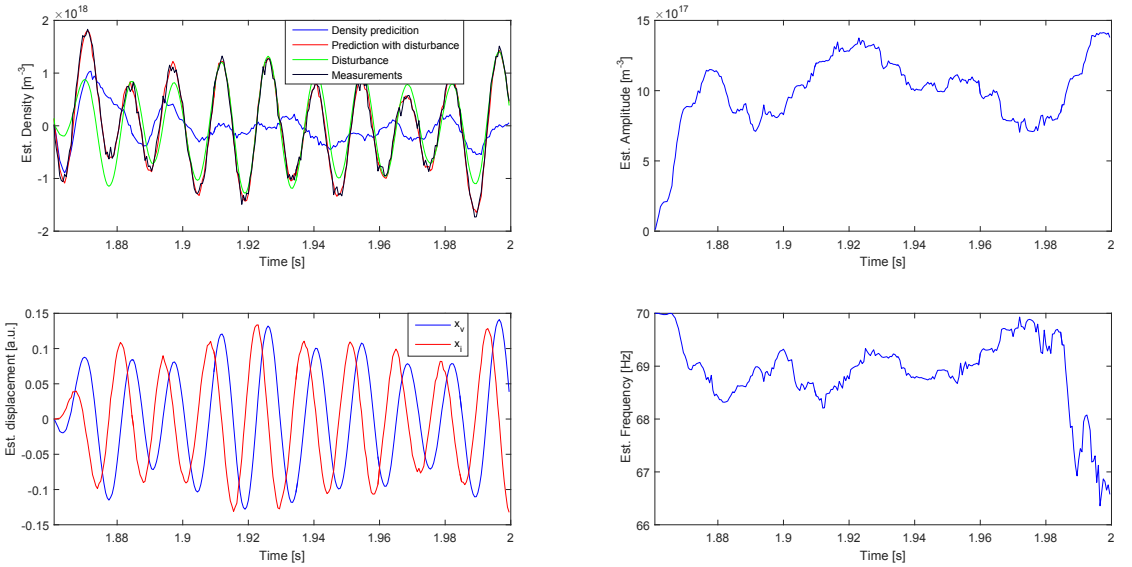

Figure B.2: Disturbance rejection for the FIR signal, tested on a post-shot disturbance to zero signal. The filter starts at 70 Hz and 0 amplitude. Within 0.05 s, the density prediciton has converged to 0 quite closely. The tracking error that remains is largely due to the varying amplitude of the oscillation, which is in principle indistinguishable from 'real' signal changes.
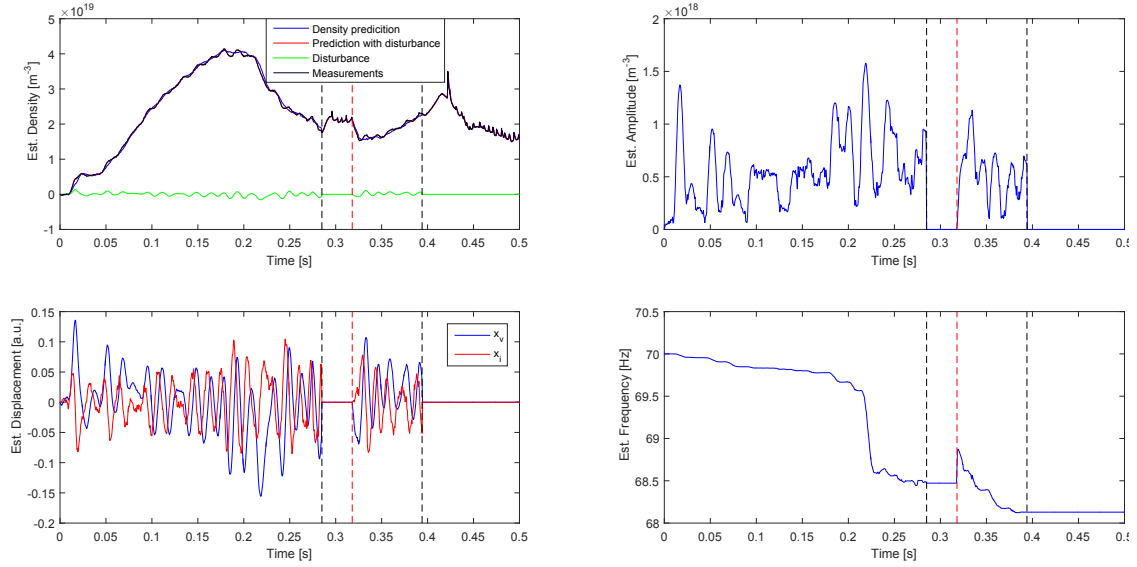
Figure B.3: Disturbance rejection for the FIR signal, tested on shot 45108. The filter starts at 70 Hz and 0 amplitude. The oscillation is effectively filtered out of the density signal, but the disturbance model had to be disabled during H-mode periods and the frequency estimate responds more strongly to signal caused by fluctuations of the plasma density (i.e., at 0.2 s) than to the disturbance.

and even becomes negative. This problem where 'real' signal changes have a much larger effect on $\hat{\omega}$ than the disturbance occurs for all values of σ_{ω}^2 between 0.01 and 30 that were used. Changing σ_{ω}^2 changes the effect of both signal components on $\hat{\omega}$ but not the ratio of their influences.

On a related note, the disturbance estimate is in principle affected by ELMs, which cause an approximately periodic (though not harmonic) disturbance of the FIR signal. For $\sigma_{\omega}^2 = 0.1$, the frequency estimate would increase to 110 Hz within the long H-mode period starting at 0.4 s (the actual ELM frequency is ~250 Hz). One way to prevent this is to take a lower σ_{ω}^2 in H-mode. However, taking $Q = 0$ in H-mode only limits the frequency state increase to 80 HZ because during the time in L-mode, an estimation error covariance is established which gives a nonzero $K_{k,\omega}$ even after Q becomes 0. This $K_{k,\omega}$ is the fifth entry of the Kalman gain K_k, i.e., the entry that allows the measurements to affect $\hat{\omega}$. In figure B.3, the problem is tackled by disabling the disturbance model during H-mode. This is done by taking $x_v = x_i = 0$, $y = n_e$, and neglecting the part of K that affects the state variables belonging to the disturbance.

Despite the apparently inaccurate frequency estimation, the disturbance is filtered from the FIR signal to a satisfactory degree in figure B.3. This introduces the possibility of neglecting possible changes in the disturbance frequency and filtering using a model assuming a constant frequency. This leads to the following system dynamics:

$$x = \begin{bmatrix} \dot{n}_e \\ n_e \\ x_v \\ x_i \end{bmatrix}, \quad \frac{d}{dt}x = \begin{bmatrix} 1 & & & \\ & & & \omega \\ & & -\omega & \end{bmatrix} x, \quad y = \begin{bmatrix} 0 & 1 & 1 & 0 \end{bmatrix} x. \tag{B.5}$$

Now, the frequency ω is taken as a constant 70 Hz throughout the shot. The FIR data can be filtered using a linear Kalman filter based on the above dynamics. The result is shown in figure
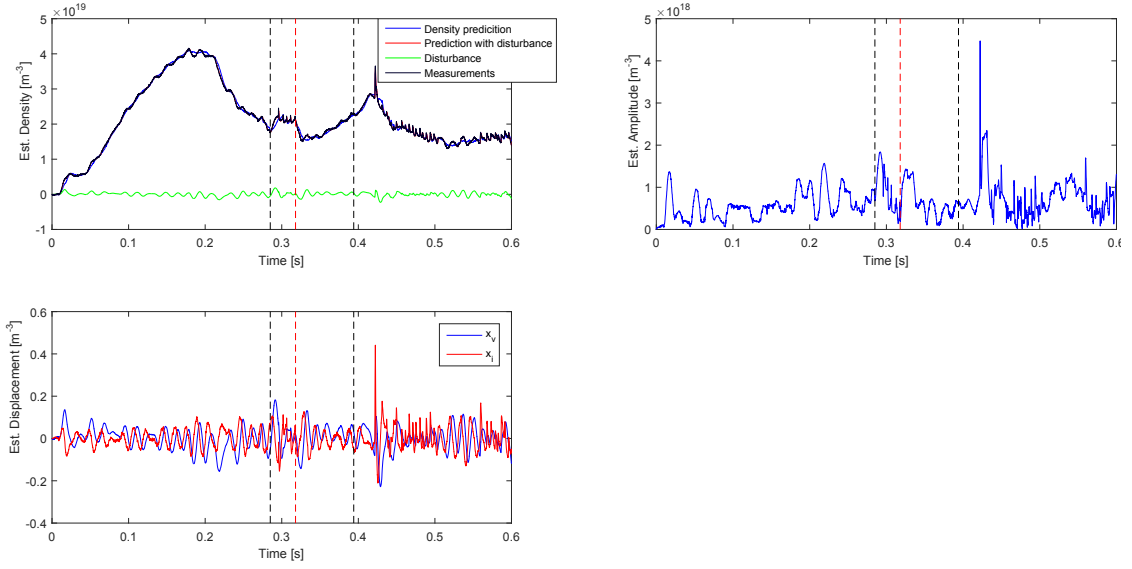
Figure B.4: Disturbance rejection for the FIR signal using a linear model, tested on shot 45108. The performance is on par with the non-linear filter. Just after 0.5 s, no ELMs affect the FIR signal. Also at that time, the oscillating disturbance is filtered out well.

B.4. The signal is filtered equally well as for the non-linear Kalman filter. Therefore, the density model with two extra states that describe an oscillation at a known frequency is the one which is used in the rest of this work.

The linear disturbance filter is also tested on data from shot 39872. This shot has only a short H-mode and therefore longer periods of L-mode where the disturbance is easily visible. Figure B.5 shows that for this shot, the disturbance is also filtered out quite well, although changes in the disturbance behaviour are still visible in the filtered signal.
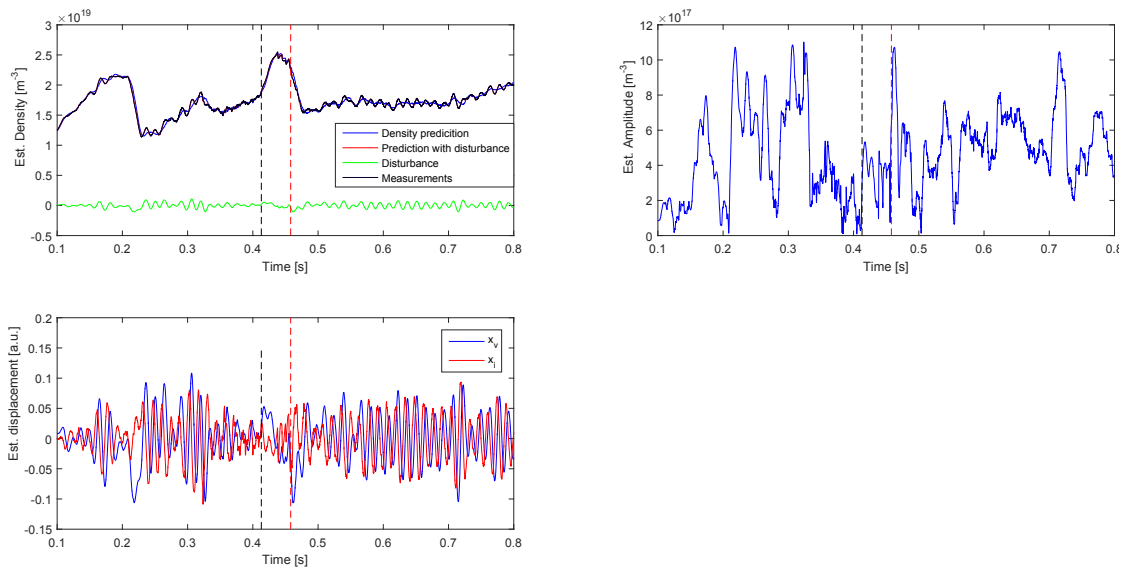
Figure B.5: Disturbance rejection for the FIR signal using a linear model, tested on shot 39872. In this shot, the disturbance behaves less predictably at some points. For example, its amplitude peaks strongly at 0.7 s. Apart from such transient effects, the disturbance is filtered out effectively.