Analysis of Plasma Filaments with Fast Visible Imaging in the Mega Ampère Spherical Tokamak

by

Thomas Paul Middleton Farley

The University of Liverpool

Department of Electrical Engineering and Electronics

Liverpool supervisor: Prof. James Bradley
CCFE supervisor: Dr Fulvio Militello

Thesis submitted in accordance with the requirements of the University of Liverpool for the degree of Doctor in Philosophy by Thomas Paul Middleton Farley

July 2019

Declaration

I hereby declare that this thesis is my own work and no further sources of information have been used other than the references cited. Neither this thesis nor any part of it have been submitted to any other university or institution for the application of another degree or qualification.

Signed: 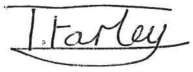 (T. Farley)

Date: 10-07-2019

Abstract

The cross-field transport of particles in the scrape-off layer (SOL) of magnetic fusion devices is dominated by the convection of coherent filamentary plasma structures. In this thesis, we present a new technique for the analysis of filaments in fast visible camera data. The new technique operates by inverting the background subtracted emission in the camera images onto a basis set of uniformly emitting field line images, constructed using information from magnetic equilibrium reconstructions. The output of the inversion is a 2D mapping of emission parametrising the average intensity of field lines in the SOL by the coordinates of their intersection with the mid-plane. Filaments manifest in the inverted emission profile as blobs of raised emission. A filament detection technique has been developed to identify these regions of increased emission and fit them with 2D Gaussians. This yields the positions, widths and amplitudes of the filaments. A tracking algorithm is then applied to calculate the filaments velocities and lifetimes.

Data from a synthetic camera diagnostic is used to assess the capabilities and limitations of the new technique and quantify its errors. This exercise shows it can detect $\sim$36% of all filaments in the analysis region, corresponding to $\sim$74% of filaments above the targeted amplitude threshold. This sensitivity is achieved with a true positive detection rate of 98.8%. Standard errors in the radial and toroidal positions of the filaments are estimated to be $\sim$2 mm, while errors in the toroidal and radial widths are around $\sim$3 mm and $\sim$7 mm respectively. The shapes of the probability density functions (PDFs) of the filament parameters are also qualitatively recovered and the effect of filament overlap on filament amplitude measurements is investigated. Valuable insight is gained into effects from the non-orthogonality of the field line basis functions and the resulting spatial dependence of the measurement errors.

Finally, the technique is applied to MAST data and compared to Langmuir probe measurements. Good agreement is found between the two diagnostics, including exponential waiting times and symmetrical conditionally averaged waveforms. Measurements of the PDFs of filament properties provide valuable inputs for analytic models of SOL transport and show filament lifetimes to be exponentially distributed. The depth of field of the technique enables measurement of the toroidal filament spacing, with results supporting the assertion that filaments are generated uniformly and independently, and are thus described by Poisson statistics underpinning several analytic models.

"With magic, you can turn a frog into a prince. With science, you can turn a frog into a Ph.D and you still have the frog you started with."

- Terry Pratchett, *The Science of Discworld*

Contents

List of Figures

List of Tables

Publications

Research performed during the period of this study has, so far, contributed to the following publications:

[1] **T. Farley**, N. R. Walden, F. Militello, M. Sanna *et al.*, "Filament identification in wide-angle high speed imaging of the Mega Amp Spherical Tokamak",*Review of Scientific Instruments*, **90**(9), 093502 (2019)

[2] B. Cannas, S. Carcangiu, A. Fanni, **T. Farley** *et al.*, "Towards an automatic filament detector with a Faster R-CNN on MAST-U", *Fusion Engineering and Design*, **146**, 374âĂŞ377 (2019)

[3] F. Militello, **T. Farley**, K. Mukhi, N. Walkden and J. T. Omotani, "A two-dimensional statistical framework connecting thermodynamic profiles with filaments in the scrape off layer and application to experiments" *Physics of Plasmas*, **25**(5), (2018)

[4] N. R. Walkden, F. Militello, J. Harrison, **T. Farley** *et al.*, "Identification of intermittent transport in the scrape-off layer of MAST through high speed imaging" *Nuclear Materials and Energy*, **12**, 175âĂŞ180. (2017)

[5] F. Militello, N. R. Walkden, **T. Farley**, W. A. Gracias *et al.*, "Multi-code analysis of scrape-off layer filament dynamics in MAST" *Plasma Physics and Controlled Fusion*, **58**(10), 105002 (2016)

Chapters 4 and 5 of this thesis include a significant amount of content from publication [1], while Chapter 6 includes content published in [3] and [5] and Chapter 7 briefly discusses results from publication [2].

Prior to the analysis of filaments, this PhD project was dedicated to the development of negative ion sources for neutral beam injectors, before funding problems with the associated SNIF experiment led to the change in project focus. While work from the original project is not discussed in this thesis, it contributed towards the following publications:

[1] D. Kogut, R. Moussaoui, N. Ning, J-B. Faure, J-M. Layet, **T. Farley**, J. Achard, A. Gicquel, G. Cartry, "Impact of positive ion energy on carbon-surface production of negative ions in deuterium plasmas", under review in *Journal of Physics D: Applied Physics* (2019)

[2] G. Cartry, D. Kogut, K. Achkasov, J.-M Layet, **T. Farley** *et al.*, "Alternative solutions to caesium in negative-ion sources: a study of negative-ion surface production on diamond in H 2 /D 2 plasmas", *New Journal of Physics*, **19**(2), 25010 (2017)

[3] J. Zacks, U. Fantz, **T. Farley**, I. Turner, R. McAdams and D. WÃijnderlich, "Characterisation of the SNIF ion source", *AIP Conference Proceedings*, **1869**(1), 30047. (2017)

Acknowledgements

There are many people I would like to thank for their generous help, guidance and support during the course of my PhD. I would like to start by thanking my two supervisors, Fulvio Militllo and James Bradley for their guidance, enthusiasm and expertise that have made this PhD possible.

During the course of the PhD I have had the great fortune to work with many colleagues and collaborators, whose help has been greatly appreciated. In particular, I have been very lucky to work closely with Nick Walkden, who I cannot thank enough for his help, advice, support, encouragement and generosity of time. I would also like to thank Andrea Piccione and Sara Mendez for their input developing the filament tracking algorithm, Ryan Chaban for his good humour and enthusiasm working on the stereoscopic camera data, the group at Cagliari University for their work on applying neural networks to the filament identification problem and James Young for helping to lay the ground work for the camera technique. I would further like to thank James Harrison, Scott Silburn and Chris Bowman for many valuable discussions, Matt Carr and Alex Meakins for their programming guidance and the many other colleagues at CCFE whom I have worked with during my time here.

In addition I would like to thank the colleagues I worked with on my original PhD project, before its change of scope, in particular Jamie Zacks, Kostiantyn Achkasov, Gilles Cartry and Ken McClements. I would also like to thank the Fusion CDT for providing me with many great experiences and opportunities and the EPSRC for funding my work under grants EP/L01663X/1 and EP/P012450/1. Huel^{TM} also deserve my thanks for keeping me 'nutritionally complete' during the intense latter stages of my PhD.

I have been fortunate to be a part of a great student community at Culham and I am very grateful for all the new friendships, laughs and lunch time conversations. In particular I would like to thank Joe Allcock, Daljeet Ghali and Alistair McGann for being great office mates, Jack Lovel and Sam Murphy-Sugrue for their camaraderie during the Fusion CDT courses and TJ Petty for all the joy he brought during my time in Liverpool. All have been amazing friends to whom I am very grateful.

I have always received endless love and support from my parents, Jenny and Paul, who have been instrumental in getting me to where I am today. Their generosity,

encouragement and understanding have been limitless and I owe them so much for all they have given both during and before the PhD.

Finally, I would like to thank my fiancée, Eleanor-Rose Corney, for all her love and support and for being there for me throughout all the trials and tribulations of the PhD. I promise to stress less from now on and look forward to spending the rest of my life with you!

Chapter 1

Introduction

1.1 World energy demand

Access to energy is strongly linked to a population's educational and economic opportunities, having wide ranging effects on health and general quality of life [1]. As wealth increases so does the demand for energy. Currently an eighth of the world's population does not have access to electricity [2]. Increases in global energy consumption are therefore both desirable and inevitable given present trends.

The Worldwide Energy Council predictions indicate global annual energy consumption is likely to increase between 27% and 61% by the middle of the century from 2010 levels depending on a number of socio-political factors [3]. This projected increase is due to both a rising global population, surpassing 9.7 billion by 2050 [4] and rising standards of living, particularly in Asia and Africa. These projections are illustrated in Figure 1.1, with developing countries outside the Organisation for Economic Co-operation and Development (OECD), particularly China, India and African countries, accounting for the majority of projected increases. Unfortunately our current methods of energy generation are not compatible with scaling to future energy requirements. In 2017 65% of global electricity production was from fossil fuels (coal, oil and natural gas), with the remainder provided by renewables (25%) and nuclear fission (10%) [5]. Irrespective of their finite reserves, the burning of fossil fuels has a range of serious ramifications that make divestment from fossil fuels of utmost and immediate importance.

The most pressing of these problems is the generation of man-made green house gases, primarily CO_2, resulting in average global temperature rises. An average global temperature increase of 1.5°C is predicted to lead to increases in extreme weather, water shortages and crop failures, resulting in unprecedented humanitarian crises. Droughts, storms and floods together displaced over a million people in 2017 and adversely affected many millions more, primarily in Africa [7], with these effects predicted to become far more severe and wide reaching with the continued effects of climate change. Furthermore, melting of the ice sheets in Greenland and Antarctica is projected to cause sea level rises of around 0.5 m, rising to as much as 2.7 m in worst case 'business as usual'

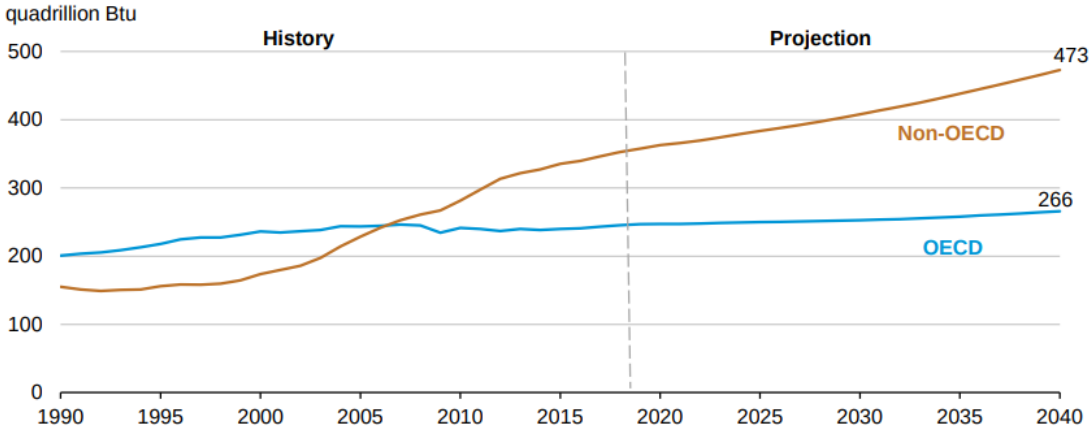

Figure 1.1: Historical and projected annual energy consumption for developed (OCED) and developing (non-OECD) countries. OECD refers to the Organisation for Economic Co-operation and Development. Figure Reproduced from Ref. 6.

scenarios [8]. With many of the world's most densely populated cities located in low lying coastal areas, many 10s of millions of people are at risk of displacement from rising sea levels.

In addition to global warming, the World Health Organisation estimates that outdoor [9] and indoor [10] air pollution each cause around 4 million premature deaths worldwide each year. Habitat change, ocean acidification and acid rain are also playing a central roll in the current mass extinction event which is estimated to threaten over a quarter of all land species with extinction by 2050 if current behaviour continues [11].

At the United Nations Climate Change Conference, COP21, in 2015, 174 nations signed the Paris Climate Accord committing to take action to prevent the average global temperature rise exceeding 1.5°C.

Figure 1.2 shows the Intergovernmental Panel on Climate Change's (IPCC) projection for the rate of greenhouse emissions and resulting temperature increase anticipated under the IPCC's most optimistic climate scenario. These projections require greenhouse gas emissions to peak in the next few years, with net zero emission achieved by the end of the century, at the same time as global energy consumption increases at the fastest rate in history.

While the initial reduction in greenhouse gas emissions can likely be met through a combination of renewable energy sources, energy storage, nuclear fission and improvements in energy efficiency and consumption patterns, these are likely to be insufficient to achieve the required level of zero emissions in the second half of the century.

Renewable energy will play a vital roll in future energy generation, but is fundamentally limited by its intermittency, large and geographically limited footprint and the strain of distributed energy generation on electricity grid infrastructure. While potential technological advances in energy storage, demand side response and transmis-

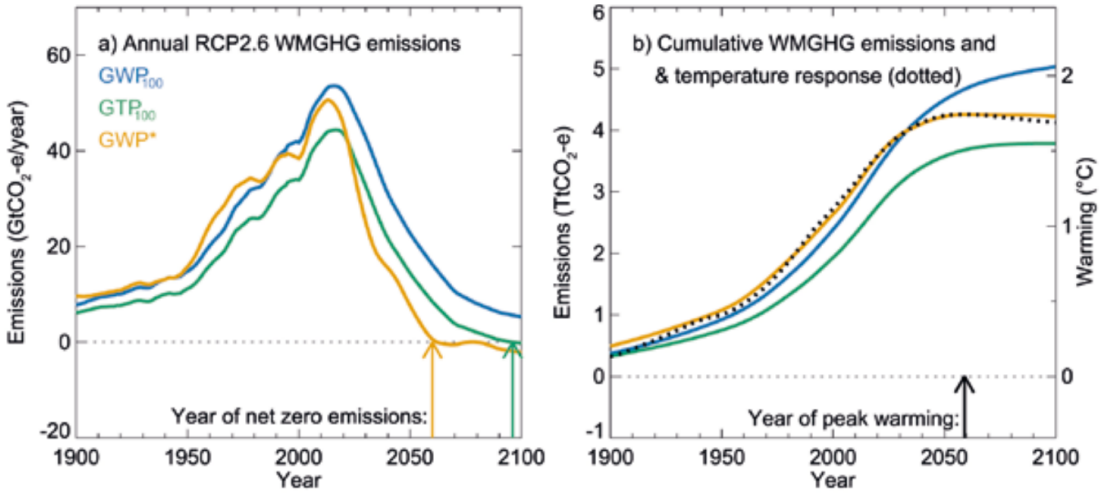

Figure 1.2: a) Aggregate emission of well-mixed greenhouse gases (WMGHGs) under the IPCC's most optimistic scenario. b) Cumulative emissions of WMGHGs and the corresponding global temperature response. The Global Warming Potential over a 100 year time horizon (GWP_{100}), Global Temperature-change Potential over a 100 year time horizon (GTP_{100}) and revised Global Warming Potential (GWP^*) are different parametrisations of combined climate force functions that are used as inputs to climate models. Representative Concentration Pathway (RCP) 2.6 is an ambitious IPCC climate change mitigation scenario in which human-caused climate change is severely limited and provides a better than two thirds likelihood of staying below 2°C temperature increase. Figure Reproduced from Ref. 12.

sion infrastructure are likely to mitigate some of these problems, high capacity, reliable baseload power generation is required to confidently avoid black outs and major social and economic disruption.

Nuclear fission, a strong candidate for carbon free base load generation, is unpopular both politically and publicly due to safety concerns after accidents such as Fukushima and Chernobyl, as well as concerns about storage of radioactive waste and proliferation issues. The energy policies of many countries include minimal use of nuclear fission, such as Germany which has committed to shutting down all nuclear fission power plants by 2022 [13], indicating that nuclear fission adoption will not be sufficient to provide base load power requirements.

Therefore, an alternative form of energy generation is required that meets the critical requirements of providing large quantities of dependable, carbon neutral energy at a price competitive with fossil fuel sources. Furthermore, the source must not depend on limited reserves of fuel, occupy a large geographical footprint, or present the safety, proliferation or radioactive waste concerns accompanying nuclear fission. Perhaps the only energy source that has a clear prospect of meeting each of these criteria together is nuclear fusion.

1.2 Nuclear fusion

Nuclear fusion is the process that powers stars and thus, via the sun, is already the direct or indirect source of most forms of renewable energy and is the historical source of energy stored in fossil fuels. Nuclear fusion is the process by which atomic nuclei fuse together to form heaver nuclei and smaller by-products, typically neutrons and alpha particles, in addition to releasing large quantities of energy.

Figure 1.3 shows the average nuclear binding energy per nucleon for the common and significant isotopes of the naturally occurring elements. For fusion reactions forming nuclei up to the mass of ^{56}Fe, there is an increase in total nuclear binding energy, resulting in a reduction in the mass of the fusion products relative to the reactants. The change in mass-energy, δm, is converted to kinetic energy, E, of the fusion products according to Einstein's famous equation, $E = \delta m c^2$, where c is the speed of light. Fusion reactions producing elements heavier than ^{56}Fe lead to a loss in binding energy and are thus endothermic and only occur in the extreme conditions of supernovae [14]. The change in binding energy per nucleon with atomic number is greatest for the light elements, particularly fusing isotopes of H to form He, making them ideal for energy production. Fusion reactions with light elements also require less extreme conditions, making them technically feasible.

The main fusion reaction of interest for energy production is that between the two isotopes of hydrogen, deuterium (^{2}D) and tritium (^{3}T):

$$^2\mathrm{D} + {}^3\mathrm{T} \rightarrow {}^4\mathrm{He}(3.5\mathrm{MeV}) + \mathrm{n}(14.1\mathrm{MeV}). \tag{1.1}$$

This reaction is favourable due to its high reaction cross section at attainable temperatures. As can be seen from Figure 1.4, the reaction cross section of the D-T fusion reaction is almost two orders of magnitude higher than the next best fusion reaction for the same reactant temperature. In addition, the energy released from the D-T fusion reaction is especially high compared to alternative reactions, releasing a total of 17.6 MeV per reaction (1 eV = 1.60×10^{-19} J). This is of the order of a million times greater than the energy yield from a chemical combustion reaction using fossil fuels ($\sim$10 eV). While future technological advancements may make other fusion reactions feasible, many consider the D-T fusion reaction to currently be the only clearly viable candidate for commercial energy production.

Deuterium is readily available, composing 0.0115% of hydrogen atoms in sea water (33 g per ton), corresponding to enough accessible D fuel for 150 billion years of current world energy consumption [17]. Tritium however is scarce, due to its short radioactive half life of 12.3 years, with very little natural production via cosmic ray interactions with nitrogen in the Earth's atmosphere. Most tritium has historically been produced in dedicated heavy water fission reactors, with most current production coming from CANDU fission reactors [18]. The world tritium inventory can be measured in 10's of

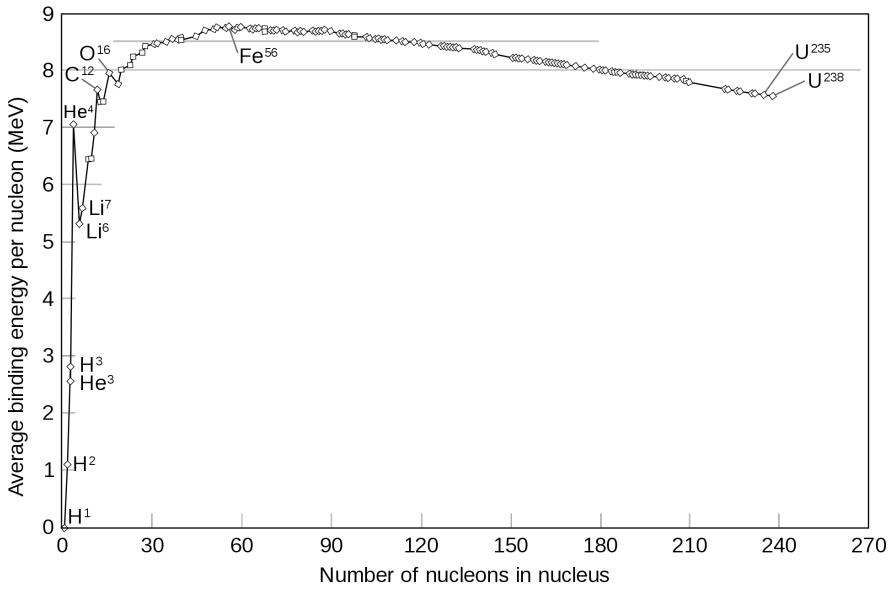

Figure 1.3: Average binding energy per nucleon for the common elements. Reproduced from Ref. 15.

kilograms [19] with only a fraction of this available for civil applications. Fortunately, the neutron produced in the D-T fusion reaction can be used to fission lithium nuclei present in breading blankets in the walls of the fusion reactor via the reactions:

$$\begin{aligned} {}^{6}\mathrm{Li} + \mathrm{n} &\rightarrow {}^{3}\mathrm{T} + {}^{4}\mathrm{He} + 4.8\mathrm{MeV} \\ {}^{7}\mathrm{Li} + \mathrm{n} &\rightarrow {}^{3}\mathrm{T} + {}^{4}\mathrm{He} + \mathrm{n} - 2.5\mathrm{MeV}. \end{aligned} \tag{1.2}$$

Thus it is possible in principle to 'breed' tritium from readily available lithium producing a self-sustaining fuel cycle. In fact it is preferable to breed slightly more tritium than is consumed, termed a tritium breeding ratio (TBR) > 1, in order for each generation of fusion reactors to generate the starting inventory of tritium for the next generation. In order to achieve this every fusion neutron must, on average, lead to the production of at least one new T atom. Given many neutrons will be absorbed by structural materials or lost through ports, materials such as lead and beryllium are needed as neutron multipliers. Given this process cannot be properly tested without a real fusion neutron source, tritium breeding blanket design is a major untested component of a fusion reactor. With tritium produced in this way, known reserves of Li are sufficient for around 3000 years at current energy consumption levels, with this number rising to 60 million years if the more expensive extraction of lithium from sea water is included [20]. Therefore, nuclear fusion has the potential to provide an inexhaustible energy source for our current energy needs.

In order for nuclear fusion reactions to occur, the fusing nuclei must come into sufficiently close proximity for the short ranged strong nuclear force to bind them together. Achieving this requires first overcoming the long ranged electrostatic repulsion of the positively charged nuclei. The classical Coulomb barrier is prohibitively large, of the

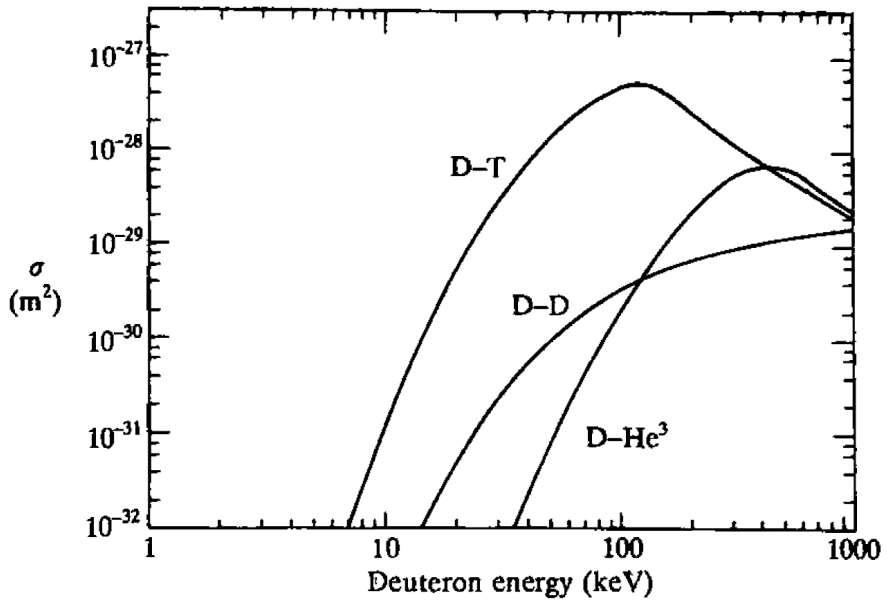

Figure 1.4: Fusion reaction cross sections for three reactions of interest for energy production. The D-D cross section shown is the sum of the cross sections for the two reaction possible D-D reaction channels. Reproduced from Ref. 16.

order of 1 MeV (1 eV = 11600 K), but fortunately quantum tunnelling results in a finite probability for fusion reactions to occur at much lower energies. However, the required temperatures of 10-15 keV (115-175 MK) are still extreme, meaning that the reactants are plasmas in which the electrons have high enough energies to escape the atomic nuclei. Provided the necessary reaction rate and temperature confinement are met, the internal heating of the plasma from fusion alpha particles (^{4}He nuclei) becomes sufficient to offset thermal losses from the plasma, enabling the external heat sources to be reduced or turned off. This point at which the heating of the plasma becomes self-sustaining is termed 'ignition'. In order for this to occur the plasma must be maintained at sufficient temperatures and densities for long enough for a high enough rate of fusion reactions to occur. This requirement is expressed by the fusion triple product, $n_i T_i \tau_E$, where n_i and T_i are the peak density and temperature of the ions in the fusing plasma and τ_E is the energy confinement time, expressing the time scale over which energy is lost from the plasma. It can be shown [16] that to achieve ignition this product must satisfy the Lawson criteria,

$$n_i T_i \tau_E > 5 \times 10^{21} \text{ m}^{-3} \text{ keV s}. \tag{1.3}$$

Stars rely on gravitational confinement and compression to achieve high densities ($\gtrsim 10^{30}$ m^{-3}) and temperatures ($\gtrsim 10^7$ K ≈ 1 keV). However, due to their reliance on direct proton-proton fusion even they are incredibly inefficient fusion reactors; the half life for a proton in the core of the sun to be involved in a successful fusion reaction is of the order of a billion years. Only due to their size do they produce sufficient fusion to

sustain themselves, with the average energy per unit volume produced the in sun being only around $\sim$1 MWm^{-3} (276 Wm^{-3} in the sun's core), over 3 orders of magnitude lower than that required for commercially viable energy production on Earth.

Due to the prohibitive size required for gravitational confinement of a fusion plasma, alternative confinement approaches are required for controlled fusion to be achieved on Earth. Most approaches fall into the broad categories of magnetic confinement fusion (MCF) and inertial confinement fusion (ICF) or a combination of the two. ICF approaches attempt to rapidly and uniformly compress fusion fuel with lasers or heavy ion beams in order to satisfy the Lawson criteria through very high densities and relatively short confinement times. Here the inertia of the fuel, which limits the expansion of the ignited fuel to the material's sound speed, is exploited to provide the brief energy confinement required to produce a fusion burn. MCF by contrast aims to maximise τ_E at more modest densities by making use of the fact that plasmas are composed entirely of charged particles and thus can be confined by magnetic fields, in place of solid materials that would be incompatible with the required operational temperatures. It is with the MCF approach to fusion that this work is concerned.

1.2.1 Magnetic confinement fusion

Charged particles, in this case atomic nuclei and electrons, in the presence of a magnetic field $\mathbf{B}$ and an electric field $\mathbf{E}$ are subject to the Lorentz force,

$$\mathbf{F} = Ze\left[\mathbf{E} + (\mathbf{v} \times \mathbf{B})\right], \tag{1.4}$$

where Ze and $\mathbf{v}$ are the charge and velocity of the particle respectively. Thus the motion of charged particles are unrestricted parallel to $\mathbf{B}$, but are constrained to Larmor orbits around magnetic field lines in the plane perpendicular to $\mathbf{B}$. The radii of the orbits, ρ, are given by [16]

$$\rho = \frac{v_\perp}{\Omega} = \frac{mv_\perp}{|Ze|B}, \tag{1.5}$$

where m, Ω and $v_\perp$ are the particle's mass, gyration frequency and velocity perpendicular to the magnetic field respectively. Thus for strong magnetic fields of several Tesla, the trajectories of the particles are closely constrained to helical paths along magnetic field lines, with ion and electron Larmor radii in a fusion device like MAST (Mega Ampère Spherical Tokamak) typically of the order $\sim 10^{-2}$ m and $\sim 10^{-4}$ m respectively [21]. By shaping the magnetic field such that it does not impinge on the reactor's main walls, confinement of the plasma can be optimised. An early approach to MCF was the magnetic mirror which relied on conservation of a gyrating particle's magnetic moment [22],

$$\mu = \frac{mv_\perp^2}{2B}. \tag{1.6}$$

Conservation of μ implies that $v_\perp$ depends on the local value of B. Given a time-invariant magnetic field performs no work on a particle, the total energy of the particle

does not change, so that $v^2 = v_\parallel^2 + v_\perp^2$ is constant. Thus $v_\parallel$ also varies with B. Therefore, given a magnetic field with regions of sufficiently high magnetic field, B_{max}, at each end, some particles travelling to the end of the device will have their velocities transferred entirely to their perpendicular components, $v_\perp = v$, with $v_\parallel = 0$, leading to the reflection of the particles back towards the centre of the device. The condition for a particle to be confined in this manner is given by [22]

$$\frac{B_{max}}{B_{min}} < \frac{v^2}{v_{\perp,min}^2}, \tag{1.7}$$

where $v_{\perp,min}$ is the perpendicular component of the particles velocity in the region of minimum magnetic field strength, B_{min}. Particles not satisfying this condition escape from the ends of the device. Even for large variations in field strength, giving large initial trapped particle fractions, collisions within the plasma lead to redistribution of particle velocities and excessive particle losses leading to very poor confinement.

Most modern approaches to MCF avoid the problem of end losses by wrapping the magnetic field into a torus shape so that the magnetic field lines close on themselves. In such devices the distance from the centre of the device to the magnetic axis is refereed to as the *major radius*, R, while the radius in the cross section of the plasma is called the *minor radius*, r. Positions can be specified in cylindrical polar coordinates, where the *toroidal angle*, ϕ, specifies the angle around the torus and the position in the resulting *poloidal plane* is given by (R, Z), where Z is the vertical Cartesian coordinate relative to the magnetic axis. A position in the poloidal plane can also be specified in polar coordinates by the poloidal angle, θ, and minor radius, r.

Unfortunately the introduction of the curved magnetic geometry introduces drifts that transport particles perpendicular to the magnetic field [22]. The first, termed the *grad-B drift*, arises from the radial variation in magnetic field strength inherent to the toroidal magnetic field. Assuming no electric field ($\mathbf{E} = 0$), taking a linear expansion in magnetic field, $\mathbf{B} = \mathbf{B_0} + \delta\mathbf{B}$ and time averaging over a gyro orbit, the resulting drift velocity can be shown to be given by [16]:

$$\begin{aligned} \mathbf{v}_{\nabla\mathbf{B}} &= \frac{1}{2}\rho\frac{\mathbf{B}\times\nabla\mathbf{B}}{B^2}v_\perp \\ &= \frac{mv_\perp^2}{2ZeB^3}\mathbf{B}\times\nabla\mathbf{B}. \end{aligned} \tag{1.8}$$

Additionally, the *curvature drift* due to the centrifugal force experienced by the particles as they traverse the curved magnetic field can similarly be shown to result in a drift given by [16]

$$\mathbf{v}_R = \frac{mv_\parallel^2}{ZeB^2}\frac{\mathbf{R}_c\times\mathbf{B}}{R_c^2}, \tag{1.9}$$

where $\mathbf{R}_c$ is the curvature vector of the magnetic field. In the absence of currents in the plasma $\nabla\mathbf{B} = -\frac{B}{R_c^2}\mathbf{R}_c$ and the two drifts are aligned. The combined drift perpendicular

to the magnetic field and its gradient can thus be written as [22]

$$\begin{aligned}\mathbf{v}_d &= \mathbf{v}_{\nabla\mathbf{B}} + \mathbf{v}_R \\ &= \frac{v_\parallel^2 + \frac{1}{2}v_\perp^2}{\Omega}\frac{\mathbf{B}\times\nabla\mathbf{B}}{B^2} \\ &= \frac{mv_\perp}{ZeB}(v_\parallel^2 + \frac{1}{2}v_\perp^2)\frac{\mathbf{B}\times\nabla\mathbf{B}}{B^2}.\end{aligned} \tag{1.10}$$

Given the drift direction depends on the sign of the particle's charge, ions and electrons drift in opposite directions producing an electric field transverse to the magnetic field. It can be shown from Equation 1.4 that the presence of an electric field produces a guiding centre drift,

$$\mathbf{v}_{E\times B} = \frac{\mathbf{E}\times\mathbf{B}}{B^2}, \tag{1.11}$$

termed the $\mathbf{E}\times\mathbf{B}$ *drift*. This drift is identical for election and ion species and directed radially outward. Therefore, while the curvature and grad-B drifts do not produce a bulk plasma drift in and of themselves, the charge polarisation and resulting electric field they produce, drives an $\mathbf{E}\times\mathbf{B}$ drift that does produce the bulk transport of particles radially across the magnetic field lines and out of the plasma, loosing confinement.

The loss of confinement resulting from curvature and grad-B drifts can be avoided by employing a poloidal magnetic field introducing what is termed a *rotational transform*. The rotational transform, ι, is a measure of the field line pitch angle, with $\iota/2\pi$ equal to the number of poloidal (short circumference) transits, m, of a field line per single toroidal (long circumference) transit, $n = 1$, around the machine. An alternative parametrisation is the *safety factor*, $q = 2\pi/\iota = n/m$, and can be approximated by the circular plasma case for which,

$$q \approx \frac{rB_\phi}{RB_\theta}, \tag{1.12}$$

where B_ϕ and B_θ are the toroidal and poloidal components of the magnetic field respectively. The combination of the toroidal and poloidal fields produces a helical magnetic field which carries particles back and forth between regions of high and low magnetic field. While on the low field side, the region of 'bad' curvature, the E × B drift carries particles radially outwards, while on the high field side the drift is directed into the plasma with the net result that the drifts cancel out as the particles traverse the field. Thus the radial loss of particles from these drifts is avoided and much higher confinement times can be achieved.

A poloidal field can be achieved in many different ways through complex shaping and placement of external magnetic field coils. These approaches fall under the family of devices called *stellarators* [24], which have many highly desirable qualities. As the magnetic field of a stellarator is fully prescribed by its external magnetic coils, the confining magnetic field can be maintained as long as the currents are maintained in the coils, facilitating a steady state confinement regime. The absence of large gradients

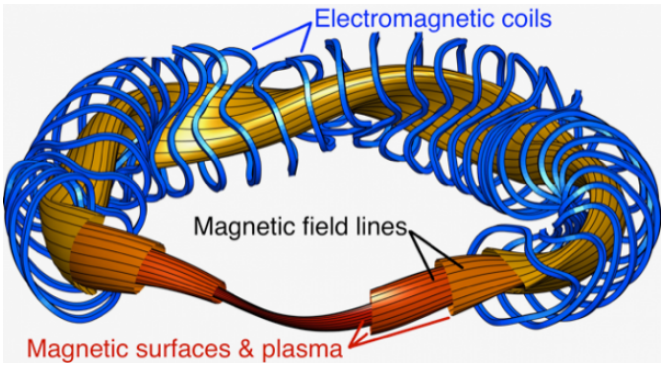

Figure 1.5: Arrangement of magnetic coils and resulting magnetic flux surfaces in the Wendelstein 7-X stellarator. Reproduced from Ref. 23.

in the toroidal currents in the plasma leads to a lack of magnetohydrodynamic (MHD) modes and vertical disruptive instabilities which are also very favourable. However, stellarators face many significant engineering challenges. Designing a highly optimised magnetic coil configuration is extremely challenging, typically requiring sophisticated computational modelling. The geometrical complexity of the resulting machine designs (see Figure 1.5) encounter severe limitations set by current engineering precision and capabilities, in particular in the manufacture of highly shaped superconducting magnets [25]. The level of optimisation and precision required in stellarator design and operation also limits their flexibility, requiring increased design iterations to explore operational parameter spaces. The complex geometry and lack of toroidal symmetry also presents challenges in the modelling and diagnosis of stellarator plasmas. Furthermore, current stellarators present even greater challenges in tritium breading blanket designs [26]; a field already subject to significant uncertainties.

Due to these and other challenges, much of the world's fusion effort is directed at an alternative approach to MCF called the tokamak. It is with the tokamak route to fusion energy production that this work is concerned.

1.3 The tokamak

The tokamak is currently the most widely pursued approach to MCF. Its name derives from a Russian acronym which translates as "toroidal chamber with magnetic coils". Figure 1.6 shows a schematic of the core components of a tokamak. In contrast to stellarators in which toroidal currents in the plasma are minimised in order to avoid perturbations to the highly optimised magnetic field produced by the external magnetic coils, tokamaks instead rely on a large toroidal current in the plasma, I_p, of the order of megamps, to generate the poloidal field. This current is driven by a step down transformer circuit provided by a solenoid running down the centre of the machine. By ramping the current through the central solenoid, a flux change through the torus is

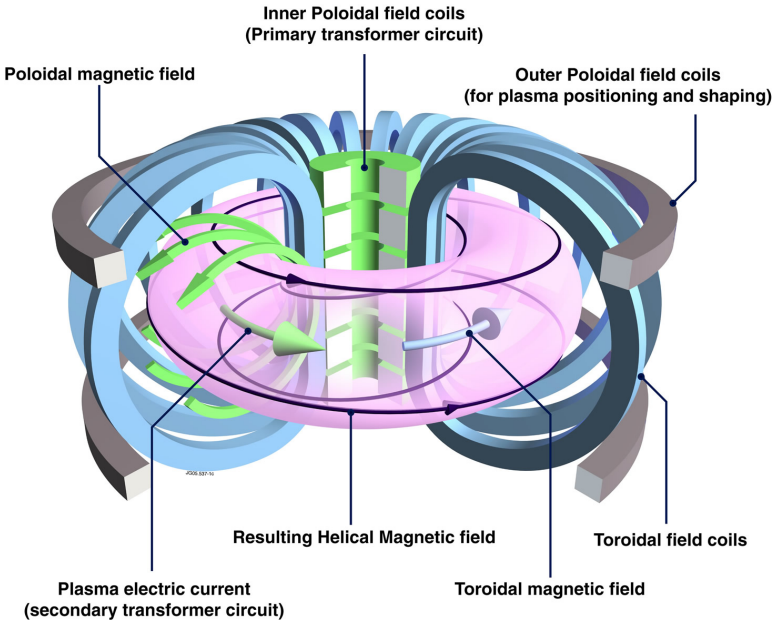

Figure 1.6: The tokamak magnetic field configuration resulting from the combination of external coil currents and the internal plasma current. Reproduced from Ref. 27.

generated, with the plasma acting as a single turn secondary winding. In addition to this current generating the aforementioned poloidal magnetic field, it also provides a source of ohmic heating for the plasma. As the extent to which the central solenoid can be driven is limited by the range of its power supplies, the plasma current cannot be driven indefinitely by induction alone, make the tokamak a pulsed device. Difficulty controlling the plasma current profile with external actuators, combined with a range of MHD and tearing mode phenomena can drive large scale instabilities that can lead to complete loss of control of the plasma in events known as "disruptions". However, it is important to note that such a loss of control does not entail the safety risks associated with loss of control of a fission power plant. Whereas fission reactor cores contain fuel sufficient for several years of operation that can be released over very short times scales due to the nuclear chain reaction going super critical, a fusion reactor would only contain fuel sufficient for several tens of seconds of operation at any moment [17]. This leads to inherent passive safety, whereby loss of control of the fusion reaction is to be avoided to prevent damage to the machine, rather than prevent a catastrophic Chernobyl-type nuclear accident.

Fusion reactors are being designed to avoid long lived radioactive waste, such that any radioactive materials generated inside the reactor will have sufficiently short half lives that they are safe for reuse on timescales of around ∼100 years. This is achieved by restricting materials used in the fusion reactor to a limited set of elements for which bombardment by 14.6 MeV fusion neutrons does not result in transmutations to long lived radioactive isotopes [28]. However, building a commercially viable fusion reactor

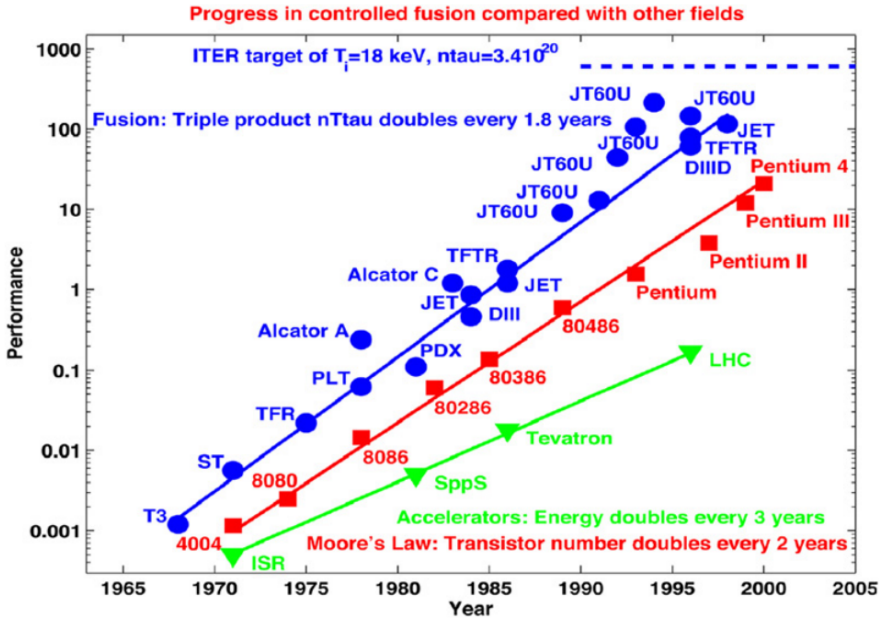

Figure 1.7: The fusion triple product has on averaged doubled every 1.8 years since the adoption of the tokamak in the 1970s. This rate of progress has exceeded that of particle accelerator design and the famous biennial doubling of microprocessor transistor density under 'Moore's law'. Reproduced from Ref. 29.

that can operate with a high duty cycle, necessarily avoiding frequent or extended shut downs for maintenance and repairs is going to be incredibly challenging given the thermal and neutron loads the machine must operate under.

1.3.1 Future fusion devices

Historical progress in fusion performance as measured by the fusion triple product has shown a doubling time of 1.8 years, exceeding the exponential progress seen in microprocessor design under 'Moore's law' [29] (see Figure 1.7).

Another important figure of merit for reactor performance is the fusion energy gain factor, $Q = P_{fusion}/P_{heating}$, where $P_{fusion} = P_\alpha + P_{neutron}$ and $P_{neutron}$ is the energy carried by the fusion neutrons which can, in part, be converted into electrical power via the blanket. A value of $Q = 1$ is referred to as *break even*, while $Q = \infty$ is called *ignition* in analogy to combustion. While this metric does not account for energy consumption in the magnets, cooling and other auxiliary systems, or inefficiencies in the heating or electricity generation systems, it gives a useful measure of the efficiency of the plasma physics aspects of a fusion reactor. As such, a fusion reactor will likely require $Q \gtrsim 40$ to be economically viable.

The current fusion energy record was set by the JET (Joint European Torus) tokamak in 1997, with a fusion output of around 16 MW corresponding to $Q \approx 0.66$ [30][†].

[†]However, it should be acknowledged that the *engineering* fusion gain, accounting for non-heating

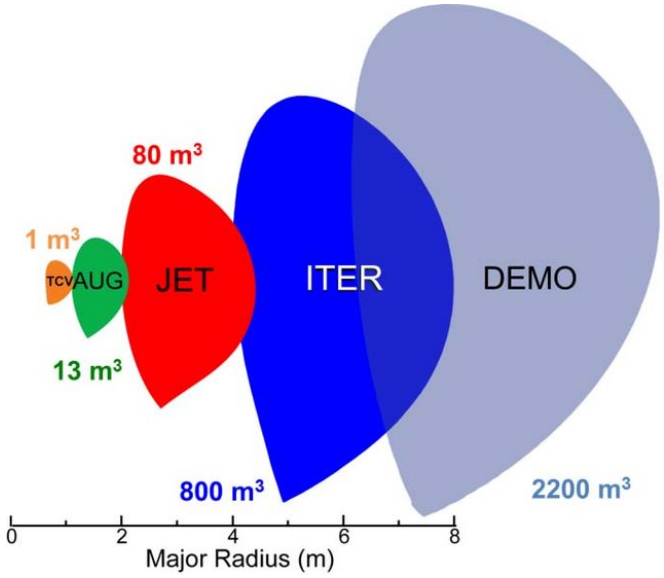

Figure 1.8: Schematic showing the increase in major radius, poloidal cross-sectional area and plasma volume with increasingly higher performance tokamaks. TCV (Tokamak à configuration variable) and AUG (Axially Symmetric Divertor Experiment Upgrade) are currently operating medium sized tokamaks located in Switzerland and Germany respectively. DEMO is the primary European demonstration power plant design. Reproduced from Ref. 31.

Projecting the performance of pure D plasmas to effective performance in D-T plasmas, JT-60U achieved $Q_{effective} \approx 1.25$ in 1998. Following this rate of progress it is projected that commercial reactor relevant fusion performance should be achievable within two iterations of large future devices.

The next step on the roadmap to fusion energy is ITER (*'the way'* in latin), an international project under construction in the south of France, with the goals of achieving $Q \geq 10$ for discharges of $300 - 500$ s and in-principal steady state (nominally $\sim$3000 s) operation with $Q \geq 5$ [32]. ITER will demonstrate solutions to a range of engineering and physics challenges. In particular, ITER will push the boundaries of current plasma exhaust solutions. Tritium retention precludes the use of carbon plasma facing components (PFC) in devices routinely operating with tritium. Therefore, ITER will operate with metallic PFCs composed of tungsten and beryllium, chosen for their thermal and low Z properties respectively. These materials introduce problems with melting, cracking and high Z impurity radiation that were not faced in carbon walled machines. Materials science imposes practical operational limits on acceptable heat fluxes to PFCs made of these materials of 10 MWm^{-2} in heavily armoured exhaust regions and $\sim$1 MWm^{-2} on the main walls. Furthermore, the high neutron fluences in fusion energy plans will likely lower these limits by a further factor of two. The plasma stored energy

auxiliary systems and electricity generation would give a value closer to $Q_{eng} \sim 5 \times 10^{-3}$ for this discharge. Furthermore this record discharge operated in a transient "hot ion" regime that is not relevant to steady state reactor operation, making the record 4s, 4 MW discharge a more relevant milestone.

in ITER will be 100s MW, with the consequence that damage from transient heat fluxes from disruptions and instabilities such as edge localised modes (ELMs) will be a limiting factor on the machine's operability without effective mitigation schemes.

These challenges will be even greater in DEMO, the demonstration power plant planned for the operations around the middle of the century. DEMO is intended to be the last state funded machine before the growth of commercial fusion sector, demonstrating electricity production to the grid, tritium self sufficiency and high duty cycle, reliable operation.

In order for ITER and DEMO to be a success, further advances are required in our understanding of heat and particle exhaust.

1.3.2 The exhaust problem

In order for the temperatures required for fusion reactions to be sustained, energy, be it from fusion alphas, P_α, or external auxiliary heating, $P_{heating}$, must be supplied to the plasma to offset energy losses from the plasma P_{loss}. The conservation of energy in the plasma can be expressed as

$$\frac{dW}{dt} = P_{heating} + P_\alpha - P_{loss} \tag{1.13}$$

where W is the stored thermal energy in the plasma.

$P_{heating}$ includes the Ohmic heating previously discussed, but given plasma resistivity drops with temperature, is insufficient alone to reach fusion conditions. Therefore additional heating sources contributing to $P_{heating}$ are used. Neutral beam injectors (NBI) produce beams of high energy atoms which can penetrate the tokamak's magnetic field before being ionised in the core and distributing their energy to the rest of the plasma via collisions. Other methods use microwaves at the resonant frequencies of the electron or ion gyro-frequencies to excite specific species and heat the plasma. These heat sources can also be used to drive current in the plasma, with the intention of prolonging tokamak discharges or achieving steady state operation. Given a fusion power plant will require a $P_{heating}$ of many tens of megawatts and P_α goes as $1/4 P_{neutron}$ ($\gtrsim$ 200 MW), the reactor must be able to exhaust very large power losses from the plasma boundaries in steady state conditions.

Power and particle losses from the plasma occur via a number of mechanisms. Radiative losses deposit energy uniformly around the reactor vessel. Free electrons emit Bremsstrahlung radiation as a result of collisions with ions and synchrotron radiation due to their gyration around magnetic field lines. Ions' interactions with electrons and neutral atoms produce line and continuum radiation. Heavy impurity species, with high atomic numbers, Z, whether seeded intentionally or eroded from the walls of the reactor, produce very high levels of recombination emission and large radiative losses. Heavy elements such as tungsten produce several orders of magnitude more emission

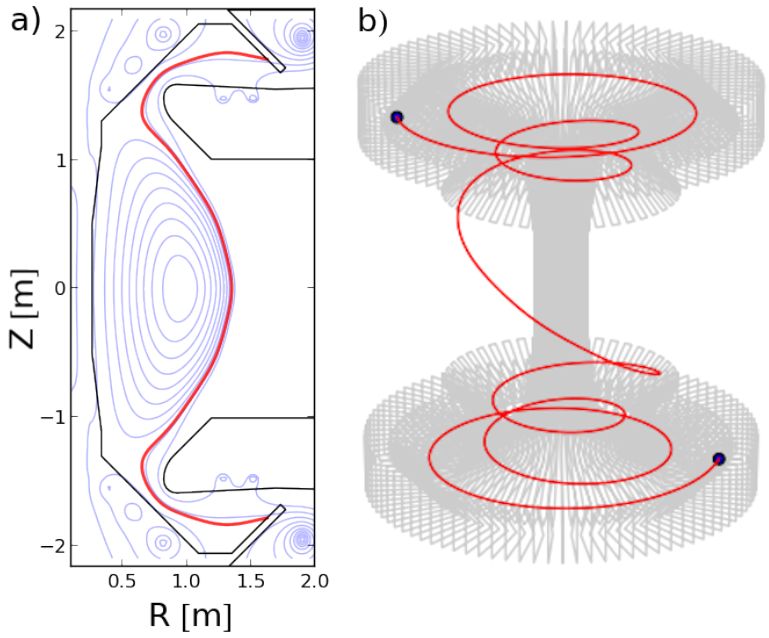

Figure 1.9: Illustration of a path taken by ions and electrons escaping the confined plasma. The particles follow magnetic field lines connecting the edge of the plasma to solid material surfaces. Despite the radial distance to the walls often being of the order of centimetres, as seen in panel a), the total distance traversed by the particles along the connecting field lines in red (the connection length) is often many meters, as illustrated in panel b).

than light elements [33], leading to radiative collapses of the plasma from very low concentrations ($\sim 10^{-3}$).

As electrons and ions in the magnetically confined plasma gyrate along magnetic field lines and follow the path of the local magnetic field, their paths out of the plasma are typically much longer than their radial distance from the walls. Particles typically travel many meters along open magnetic field lines before they encounter solid surfaces, despite the walls being situated only centimetres from the edge of the plasma. The length of field line connecting a point in the plasma to a material surface is termed the *connection length* and is illustrated in Figure 1.9.

Collisions between particles on different magnetic flux surfaces facilitate the transport of energy and particles across the magnetic field over length scales of the Larmor orbit. This classical diffusion is sufficiently small that early attempts at fusion power considered small devices with major radii, $R \approx 40$ cm, (distance from the machine centre to the toroidal magnetic axis) sufficient to achieve a burning plasma. However, radial variation in the magnetic field strength produces trapped particle orbits, reminiscent of those in magnetic mirrors discussed in §1.2.1, with particles in the loss cone given by Equation 1.7 executing so called *banana orbits* that never reach the inner high field side of the device. These trapped particle orbits are not confined to the concentric flux surfaces to which passing particles are confined, providing a mechanism of trans-

porting heat across flux surfaces on length scales of the order of several Larmor radii, depending on the strength of the poloidal magnetic field. This additional transport of particles across the magnetic field is termed *neoclassical* transport, in contrast to *classical* transport in the case of non-toroidal configurations. Besides increased transport, friction between trapped and passing particles generates the *bootstrap current*, which can be beneficial for contributing to non-ohmic current drive.

However, even neoclassical transport heavily underestimates the transport that is experimentally observed. Steep gradients in the plasma produce a broad range of microscopic plasma instabilities with a variety of different drives and scale lengths. These instabilities contribute to a low level broadband turbulence that pervades the plasma. Turbulent eddies efficiently transport particles and energy across the magnetic field. The resulting enhanced cross field transport, termed *anomalous* transport, can be orders of magnitude higher than neoclasical transport [34].

Given the engineering limits on the strength and thus stiffness of the magnetic field that can be produced with current superconducting magnet technology, fusion reactors must be built to be very large, in order that the large length scales over which turbulence must transport particles increase τ_E. Thus future fusion reactors will be very large, high power devices (see Figure 1.8) for which handling plasma exhaust becomes many times more challenging than in current devices.

1.4 Plasma surface interactions

A balance must be struck in the proportion of in vessel components that interact with the plasma. On the one hand the proportion of surfaces interacting with the plasma should be minimised to reduce the amount of impurities that enter the plasma. In addition to their thermal energy, ions are also accelerated through the Debye sheath, increasing the rate of sputtering of material from the reactor walls. Light ions, in particular helium, also can penetrate into the atomic lattice of structural materials causing swelling and embrittlement [35], while heavy impurity ions from the walls are especially effective at causing further erosion due to more efficient momentum transfer to like atoms in the walls. At the same time the high heat fluxes present in a fusion reactor necessitate that the area of interacting materials should not be so small as to overly focus the heat exhaust and exceed wall material thermal limits, causing permanent structural damage. It can be argued that the high plasma confinement that is so beneficial in the plasma core is a hindrance in the plasma edge where it makes it very challenging to spread the plasma exhaust over sufficiently large areas in appropriate areas of the device.

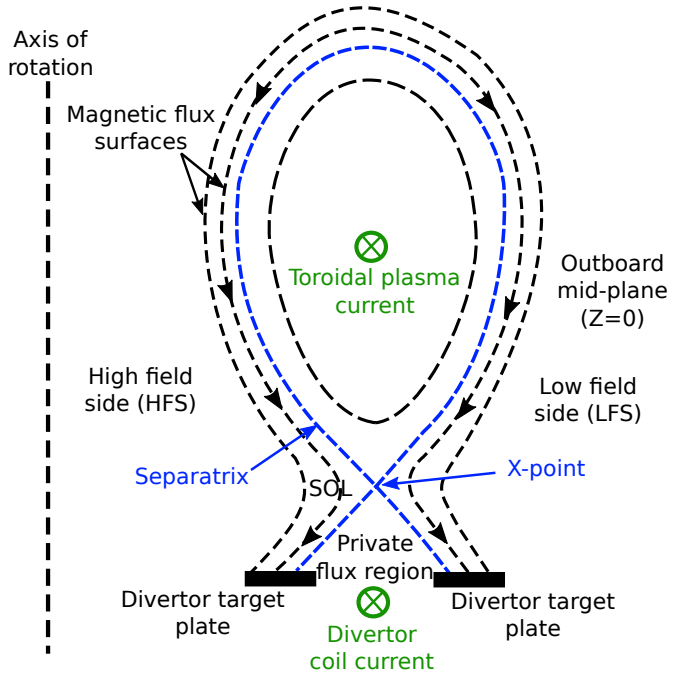

Figure 1.10: Schematic of the poloidal cross section of a diverted tokamak. Magnetic flux surfaces are shown as dashed lines. The separatrix, separating the open and closed magnetic flux surfaces is shown in blue. The toroidal plasma and divertor coil currents are shown in green. Allows on flux surfaces indicated the rapid flow of plasma parallel to the magnetic field, carrying plasma from the mid-plane to the divertor targets.

1.4.1 Limiters

A simple method of prescribing where plasma surface interactions occur is through the use of *limiter* structures. Limiters, so called because they *limit* the extent of the plasma, are armoured structures that intrude into the edge of the plasma and thus determine the location of the last closed flux surface (LCFS). Given the narrow width of the unconfined edge plasma, this allows the rest of the vessel to be lightly armoured despite its proximity to the plasma edge. Unfortunately, this configuration has the disadvantage that the plasma surface interactions (PSI) occur at the very edge of the confined plasma, enabling sputtered impurities to easily penetrate into the core and degrade confinement. Therefore, while limiters are often used in the start up and ramp down phases of a discharge, the most common magnetic configurations used in a tokamak utilise a *divertor*.

1.4.2 Divertors

A divertor configuration is achieved through the creation of a magnetic null point in the poloidal plane, within what would otherwise be the LCFS. The magnetic null is created by a divertor coil carrying current parallel to the plasma current and manifests as an 'X' in a poloidal plot of flux surfaces, giving rise to the name '*X-point*'. In a

diverted plasma the concept of the LCFS is replaced with that of the *separatrix*, which is the flux surface passing through the magnetic X-point null. This configuration is illustrated in Figure 1.10. By moving the primary plasma facing components, referred to as the divertor *targets*, to the other side of the X-point, plasma surface interactions are moved far from the core plasma. This increased separation, combined with the counter-flow of plasma into the divertor region, reduces the influx of impurities into the core, improving performance. Indeed the benefits of a divertor are such that, ASDEX (Axially Symmetric Divertor Experiment), the first large tokamak to install a divertor, quickly discovered a new operational regime, the *high-confinement mode* or *H-mode*, in contrast to the *low confinement mode* (*L-mode*) familiar to limiter machines [36]. H-mode typically results in a doubling of τ_E due to the formation of an edge transport barrier that leads to steep temperature and density gradients in the edge of the confined plasma, termed the *pedestal* [37]. A consequence of these steep gradients is that they are accompanied by the aforementioned edge instability, the ELM, which imposes large transient power loads on the divertor targets. Provided the ELMs can be mitigated, H-mode is a favourable regime of operation for future fusion reactors.

1.5 The scrape-off layer

The scrape-off layer (SOL) is the name given to the outer, unconfined region of the tokamak plasma outside the separatrix, whose magnetic field lines intersect with material surfaces. These field lines are considered *open* due to the fact they meet a solid surface before they can close back on themselves. Material surfaces act as recombination sinks for the plasma, enabling the plasma to maintain parallel velocities of the order of the plasma sound speed, $v_\parallel \approx c_s \approx [k_B(T_e + T_i)/m_i)]^{1/2}$, where k_B is the Boltzmann constant, m_i is the ion mass and T_e is the electron temperature. For typical edge temperatures of $T_e \approx T_i \sim 25$ eV and D^+ ions, $c_s \sim 50$ kms^{-1} [38]. By contrast filamentary structures in the edge transport particles across the magnetic field at $\sim$1-5% of the sound speed. Thus the width of the SOL is governed by the balance between parallel and perpendicular transport processes, with the rapid streaming of plasma parallel to the field resulting in a very narrow SOL. The plasma density, n, has historically been described by a diffusion-advection equation for the steady state SOL [39],

$$\begin{aligned} \frac{\partial n}{\partial t} &= \nabla \cdot (D_\perp \nabla n) - \nabla \cdot (\mathbf{v} n) + \xi n \\ &\implies \frac{\partial}{\partial r}(D_\perp \frac{\partial n}{\partial r}) - \frac{\partial}{\partial s}(v_\parallel n) + \xi n = 0 \\ &\implies D_\perp \frac{\partial^2 n}{\partial r^2} + \xi n \approx \frac{n c_s}{L_\parallel}, \end{aligned} \tag{1.14}$$

where r and s are the radial and parallel coordinates respectively and the cross field diffusion coefficient, $D_\perp$, is assumed to be constant across the SOL. The term $\xi n = n_0 n \langle \sigma v \rangle_i$

is the source of particles from ionisation, where n_0 is the density of neutral particles and $\langle \sigma v \rangle_i$ is the integrated ionisation cross section. Here we have made the assumption that gradients parallel to the magnetic field are small, so that $\frac{\partial n}{\partial s} \approx \frac{n}{L_\parallel}$, where the connection length, $L_\parallel$, is the distance along the field line to the solid surface. The origin for the radial coordinate, r, is at the separatrix. Equation 1.14 can be solved by an exponential density profile of the form

$$n(r) = n_s \exp\left(-\frac{r}{\lambda_n}\right), \tag{1.15}$$

where $n_s = n(r = 0)$ is the density at the separatrix and the density width of the SOL is given by

$$\lambda_n = \sqrt{\frac{D_\perp L_\parallel}{c_s}}. \tag{1.16}$$

Through similar reasoning, the other thermodynamic properties of the SOL such as temperature and parallel heat flux, $q_\parallel$, are also typically taken to have exponential profiles described by equivalent e-folding lengths, λ_T and λ_q. For representative values of $D_\perp \sim 1$ m^{2}s^{-1}, $c_s \sim 50$ kms^{-1} and $L_\parallel \sim 50$ m, we get a SOL density width of $\lambda_n \sim 3$ cm. A consequence of this rapid exponential radial variation in thermodynamic properties is that the majority of the heat and particles escaping the plasma are deposited over narrow annular regions of the machine, leading to concentrated heat fluxes and erosion. Given the width of the SOL determines which areas of the machine interact with the plasma and are subject to erosion and damage processes, a strong understanding of SOL physics is required to accurately predict and ultimately tailor its properties for the development of future machines.

1.5.1 Scrape-off layer regimes

The properties of the SOL depend strongly on the importance of collisional processes in producing parallel gradients in the SOL, with SOL conditions falling into three broad regimes [40]. The characteristic parameter describing the behaviour of the SOL is the *SOL collisionality* [38],

$$\nu^* = \frac{\ell_\parallel}{\lambda_e}, \tag{1.17}$$

where $\ell_\parallel$ is the parallel length scale of the SOL (typically the connection length between mid-plane and the target, $L_\parallel$) and λ_e is the electron mean free path. Thus ν^* gives an estimate of the number of collisions an electron will incur as it travels from the outboard mid-plane to the divertor target or limiter.

The *sheath limited* regime occurs when $\nu^* \lesssim 10$ corresponding to low density discharges with moderate input power. Under these conditions flows to the target are relatively unimpeded by collisions and are limited by the heat transmission properties of the sheath. Therefore temperature and density profiles are relatively uniform along field lines.

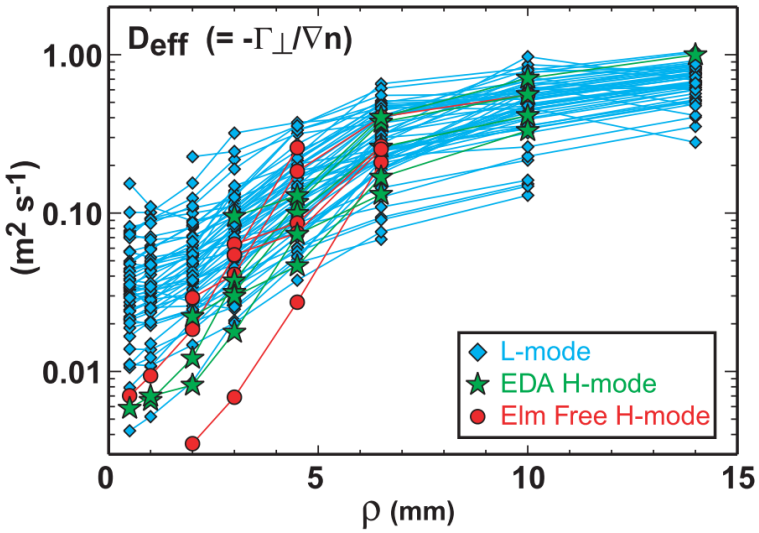

Figure 1.11: Radial profiles of the effective cross-field particle diffusivity ($D_\perp$) for ohmic L mode discharges in Alcator C-Mod over a range of line averaged plasma densities (1.0×10^{20} m^3 $< n_e < 2.3 \times 10^{20}$ m^3) and several ohmic H-mode discharges. The radial coordinate ρ is zero at the separatrix and positive in the SOL. The unphysical variation in $D_\perp$ over two orders of magnitude indicates the invalidity of the diffusive SOL model. Reproduced from [41].

For intermediate density and power discharges for which $10 \lesssim \nu^* \lesssim 85$, the SOL is in the *conduction limited* regime. Under these conditions collisions are sufficient to limit parallel heat conduction and produce significant thermal gradients in the SOL close to the target. The reduced temperature at the target increases cooling from charge exchange processes, further lowering the temperature. Pressure conservation along flux tubes requires that the density increases at the target, with the high target density levels maintained by strong recycling. Thus the conduction limited regime is often synonymous with the *high recycling regime* [38]. Although the total power to the target is relatively unaffected, the lower temperature at the target has the beneficial effect of lowering the sputtering yield.

For sufficiently high densities with $\nu^* \gtrsim 85$ the SOL enters the *detached* regime. Under these conditions the temperature at the target drops bellow a few eV, so that volume recombination and neutral frictional drag effects become important. Detachment can be aided by impurity seeding in the divertor to increase radiative cooling. The target is considered 'detached' due to the cloud of neutrals that forms, shielding it from the upstream plasma. In this situation the pressure drops close to the target, leading to a reduction in both the plasma temperature and density. Volumetric radiative heat losses spread the exhausted power over larger areas, making operation in detached regimes very favourable for future fusion devices.

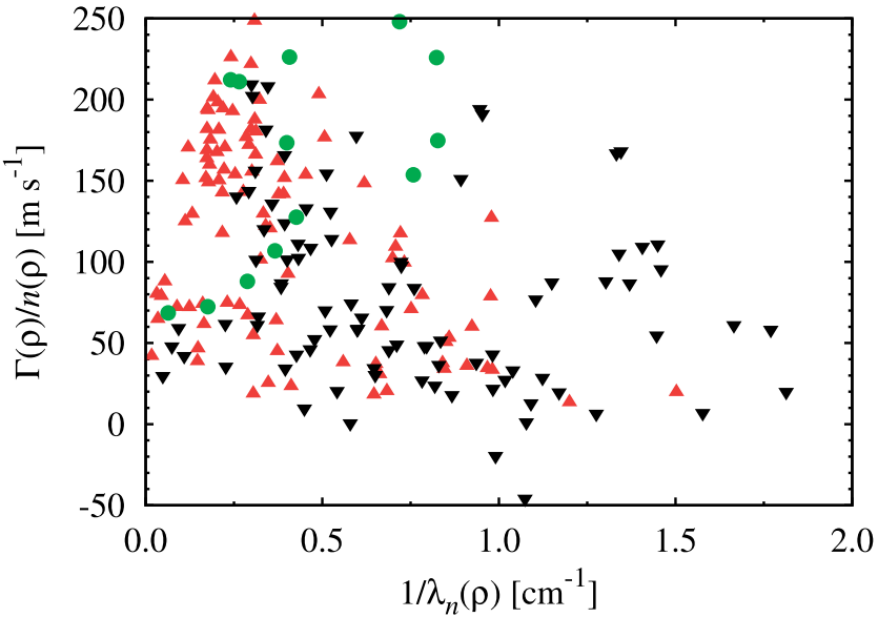

Figure 1.12: Scatter plot showing the relationship between the edge particle flux, Γ, normalised to the edge density, n, and the reciprocal density decay length, λ_n, as measured with a reciprocating Langmuir probe in TCV. Black ▼ and red △ are measurements taken at low and high densities respectively, while green • correspondent to ESEL simulations. The absence of linear structure indicates the data cannot be explained via a simple diffusive-advective model. Reproduced from Ref. 44.

1.5.2 Anomalous transport

The simple diffusive model of the SOL discussed above fails to account for experimentally observed levels of cross field transport, requiring that empirically derived *anomalous* diffusion coefficients be used [16]. In order to fit experimental data with the diffusive-convective model it is typically necessary for $D_\perp$ to be very large and vary strongly both across the radial extent of the SOL and between different plasma regimes [41], as demonstrated in Figure 1.11. These $D_\perp$ are higher than those given by Bohm diffusion [42, 43] and the strong variations over two orders of magnitude cannot be physically justified, indicating such a model is incorrect. Furthermore, Langmuir probe measurements have shown that SOL transport is incompatible with a model of transport described purely in terms of an effective particle diffusivity and a convective velocity [44]. This is demonstrated in Figure 1.12.

Experimental measurements often demonstrate *flattening* of density profiles, whereby the steepness of the SOL profiles decreases with increasing depth in the SOL [45]. *Broadening* is also often observed, corresponding to a large increase in λ_n at high fuelling. There is often a notable shoulder or discontinuity in the profiles, leading to the SOL being divided into *near* and *far* SOL regions depending on the distance from the separatrix [46]. Figure 1.13 illustrates these phenomena and indicates the density shoulder marking the transition from the near to the far SOL. The anomalous transport leading

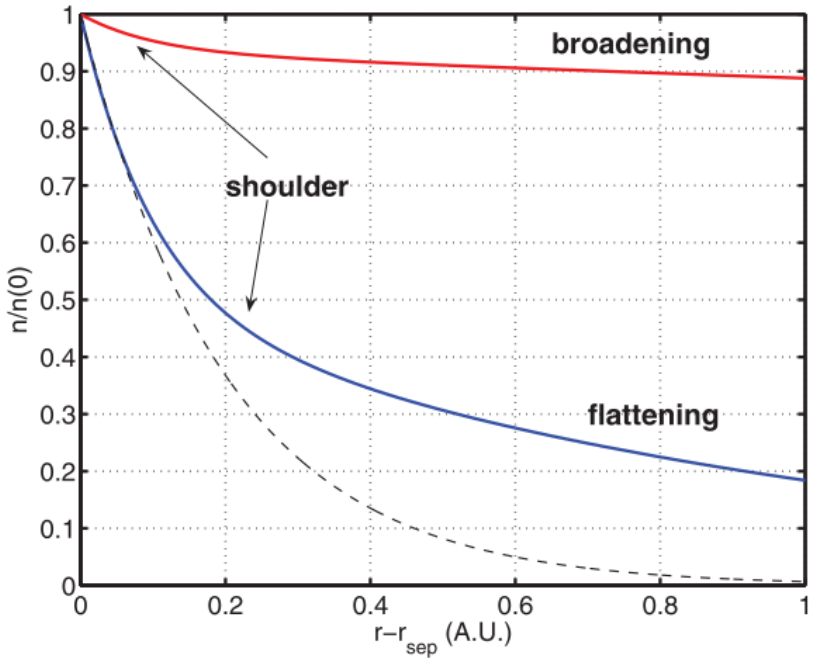

Figure 1.13: Illustration of flattening and broadening phenomena in SOL density profiles. Flattening describes the change in decay length between the 'near' (inner) and 'far' (outer) SOL regions, separated by a density 'shoulder', which becomes more pronounced with increased fuelling and decreasing plasma current. Broadening corresponds to the broadening of the far SOL region to encompass most or all of the SOL. Reproduced from Ref. 45.

to these phenomena is believed to arise from the coherent radial propagation of turbulence induced structures in the edge of the plasma. On crossing the separatrix, these turbulence induced density fluctuations rapidly spread out along field lines, manifesting as filaments, which advect particles and, to a lesser extent, power and momentum across the SOL.

1.5.3 SOL filaments

In the last decade or so it has become increasingly clear that a large portion of the radial flux of particles and heat in the SOL is carried by these filamentary structures referred to in literature as filaments, blobs, avaloids and intermittent plasma objects (IPOs) depending on the manner of observation [48, 49]. Filaments are ubiquitous in magnetically confined plasmas and are found in tokamaks, stelarators, linear plasma devices and reversed field pinches, having been observed in over 40 devices [49,50]. These coherent filaments of plasma propagate out radially from the plasma edge, draining through parallel flows to the target as they travel. Thus the plasma travels in large discrete units with densities of the order of the background plasma through which they move, in contrast to the classical picture of a smoothly varying diffusive SOL. It is the time average of the ensemble of many radially propagating and draining filaments that

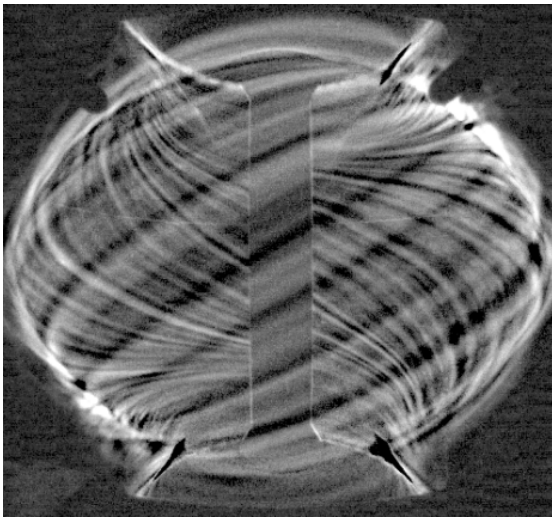

Figure 1.14: Background subtracted visible camera image of filaments in the MAST tokamak. The background subtraction is achieved by subtracting the pixel-wise minimum intensity across a set of frames preceding the frame of interest (see [47] for further details). Reproduced from Ref. 47.

produce the familiar two-part exponential profiles. Recent forward modelling [45] on MAST has shown that particle [51] and heat-flux [52] profiles at divertor targets can be reconciled with a transport based on filament motion. The presence of filaments leads to a non-local relationship between fluxes and gradients in the SOL [44] and a fuller understanding of the physics underlying filaments is required before predictions can be made for future devices. This requires extensive experimental measurements of filaments for comparison with modelling and the informing of analytic filament frameworks.

The SOL profiles resulting from filamentary transport impact many aspects of machine operation. The broadening of the SOL and the fact that high density structures travel much further out in the SOL than would be anticipated through classical transport, can lead to much higher particles fluxes into the far SOL [42]. This has implications for increased interactions and erosion of weakly armoured wall strictures such as radio frequency antennas, wall panels and limiters than would otherwise be anticipated [53]. The SOL is also an important interface between the external actuators and the core plasma, through which fuel particles must penetrate to reach the core [54] and He and impurity ions must be flushed to prevent accumulation in the core [55]. Filaments thus play an important role in these transport processes. Similarly, resonant heating waves must first traverse the SOL to reach the intended resonant surfaces, but can be scattered by filamentary perturbation of similar scale length, resulting in inefficient coupling of ion cyclotron and lower hybrid waves to the plasma [56]. Finally filaments play an important role in SOL processes influencing divertor detachment onset [57] and tritium retention [58, 59]. With this in mind, it is important that a sound understanding of perpendicular transport in the SOL is established.

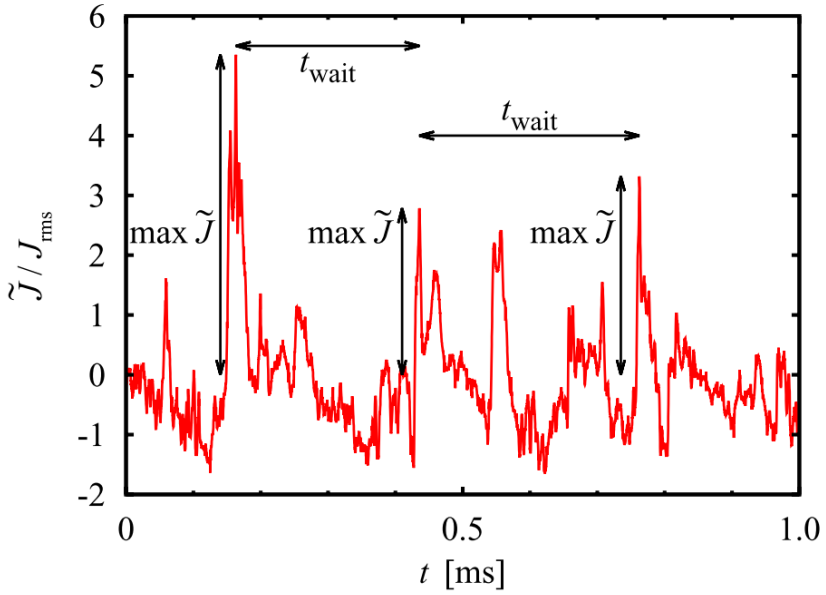

Figure 1.15: Time series of ion saturation current ($\tilde{J}$) fluctuations from a Langmuir probe in TCV. Filaments are visible as large intermittent events with peak amplitudes above 2.5 times the rms value of the signal. The amplitude of the filaments is indicated by 'max $\tilde{J}$' and the time between filament events by τ_{wait}. Reproduced from Ref. 61.

1.6 Advantages and disadvantages of filament measurement techniques

Experimental investigations of filaments are usually made through measurement of local plasma density in the SOL, with filaments identified as density and temperature perturbations matching a given set of criteria. The vast majority of filament measurements in the literature can be divided into the two categories of Langmuir probe (LP) diagnostics and optical imaging diagnostics, although observations of filaments have also been made with Thomson scattering and electron cyclotron emission diagnostics [60]. Of the optical diagnostics, the most common are beam emission spectroscopy (BES), gas puff imaging (GPI) and direct visible imaging (DVI).

This thesis focusses on the analysis of direct wide-angle fast visible imaging diagnostics, which have not been widely exploited in the past. To understand the role that DVI can play in filament studies, each of the main diagnostics used to study filaments is discussed in more detail below, along with their relative strengths and weaknesses. The experimental findings from measurements with these diagnostics are reserved for Chapter 2.

1.6.1 Langmuir probes

Langmuir probe measurements are most often made with individual probe heads either mounted on reciprocating pistons that can perform radial sweeps through the plasma,

or fixed probes that are flush mounted to plasma facing components. A smaller number of probe measurements have been made with 2D arrays of probes enabling spatially resolved instantaneous remeasurements of filaments [62,63], although these have typically been restricted to smaller devices. The main signal of interest for filament measurements with Langmuir probes is the ion saturation current [64],

$$\begin{aligned} I_{sat} &= \frac{1}{2} A_{pr} e n \sqrt{k \frac{T_i + T_e}{m_i}} \\ &\propto n \sqrt{T_i + T_e}, \end{aligned} \tag{1.18}$$

where A_{pr} is the area of the probe head and e is the electronic charge. This is the current drawn by the probe when it is biassed sufficiently negatively that all electrons are repelled. The current that is drawn is limited by the flux of ions entering the probe sheath at the ion sound speed, as a result of the Bohm criterion, which results from the acceleration of the ions in the pre-sheath surrounding the probe. Given I_{sat} depends linearly on n, but only weakly on temperature, it can be treated as a proxy density measurement. Thus filaments propagating past the probe tip register as spikes in I_{sat} (see Figure 1.15), whose statistical properties can be analysed.

Langmuir probe filament measurement techniques

Rather than attaining in-depth information about the physics of individual filaments, Langmuir probe analyses typically rely on statistical techniques to asses the properties of filament populations. The probability density function (PDF) of the ion saturation current measurements in the SOL show non-Gaussian behaviour characteristic of intermittent events.

Important measures of the deviation of a distributions properties from the Gaussian case are the 3$^{\text{rd}}$ and 4$^{\text{th}}$ statistical moments of a PDF, $\mathcal{P}(x)$:

$$S(x) = \int \frac{(x - \bar{x})^3}{\sigma^3} \mathcal{P}(x) dx = \frac{1}{N} \sum_{i=0}^{N-1} \frac{(x_i - \bar{x})^3}{\sigma^3}, \tag{1.19}$$

$$K(x) = \int \frac{(x - \bar{x})^4}{\sigma^4} \mathcal{P}(x) dx = \frac{1}{N} \sum_{i=0}^{N-1} \frac{(x_i - \bar{x})^4}{\sigma^4}, \tag{1.20}$$

where N is the total number of samples in a discrete PDF and the mean, $\bar{x}$, and variance, σ^2, are the familiar 1st and 2nd statistical moments. The skewness, S, provides a measure of how biased the distribution is towards positive or negative values with respect to the mean and is zero for a Gaussian or any other symmetrical PDF. The kurtosis, K, is a measure of how pronounced the tails of the distribution are, with larger values indicating broader tails and hence a predominance of extreme events with respect to Gaussian behaviour. Given the kurtosis of a Guassian is $K = 3$, the excess Kurtosis, $K_{ex} = K - 3$, is often used to more clearly highlight differences from the Gaussian case.

PDFs of I_{sat} are observed to become increasingly positively skewed with increasing depth in the SOL [65], indicating the rise of large positive density fluctuations out into the far SOL. Positively skewed fluctuation statistics are observed universally in toroidal plasma devices [49], with in-depth studies performed on tokamaks including Alcator C-Mod [46], DIII-D [48], TCV [66], MAST [67], JET [68] and TEXTOR [65] among others [69]. An example of this universality is demonstrated by the common skewed PDF shapes in Figure 1.16, which take the form of a Gamma distribution.

A common analysis technique used to study the form of intermittent fluctuations is *conditional averaging*. This involves setting a threshold value of I_{sat}, typically a number of standard deviations above the mean signal level. Peaks that exceed this threshold are labelled as filaments. A fixed time window is taken around each peak in the time series and these windows are averaged to produce an 'average' filament profile. These filament profiles can often be modelled as double exponentials. An example of conditionally averaged waveforms observed on different machines is given in Figure 1.17. Cross conditional averaging can also be applied where identical time windows are averaged in the time series for other parameters such as the floating potential, V_p, or radial component of the velocity, $\mathbf{v_r}$. In order to calculate v_r, the floating potentials from two nearby probe tips are used to approximate the local perpendicular electric field, $\mathbf{E}_\perp$, from which the radial drift velocity of the filament (discussed further in §2.3.2) is inferred, $\mathbf{v_r} = \mathbf{E}_\perp \times \mathbf{B}/B^2$ [22].

Another property of interest for understanding the level of filamentary transport is the distribution of waiting times, τ_{wait}, describing the time between successive filaments passing the probe tip. These measurements are made by considering the time differences between peaks satisfying the conditional averaging conditions (see Figure 1.15). A range of other time series analysis can also be applied such as measurements of autocorrelation times [61] and fluctuation power spectra [69] which provide information about the spatial structure and motion of the filaments. A number of these techniques are discussed further in Chapter 6.

While Langmuir probes are very versatile and have yielded a great deal of information about filaments from a range of machines, they have a number of limitations. The intrinsic 0D/1D nature of Langmuir probe measurements limits the spatial information they can gain. With the exception of a minority of multi-probe studies [62, 63] which have facilitated direct measurements of filament widths and velocities, Langmuir probes, whether fixed or reciprocating, are limited to studies of filament population statistics and cannot provide measurements of individual filaments over their lifetimes. This precludes detailed studies of filament dynamics, which are vital in order to understand the generation of filament profiles. Reciprocating Langmuir probes (RCPs), with which radial measurements of filament statistics are made, are also intrusive diagnostics, necessarily perturbing the plasma they are measuring. As a consequence the

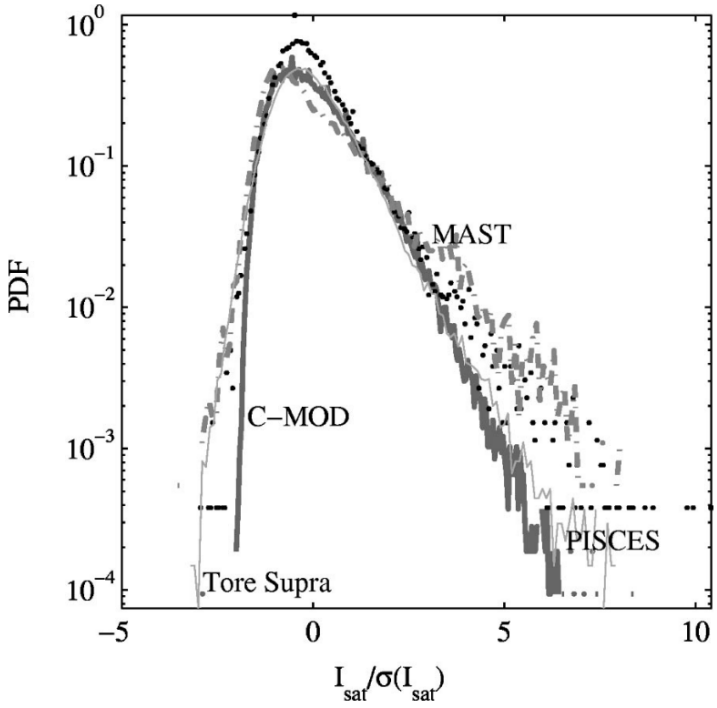

Figure 1.16: PDF of the ion saturation current normalized to its standard deviation for Langmuir probe measurements in the the Tore Supra, Alcator C-Mod, MAST and PISCES devices. The integral of the four PDFs have been normalised to 1. Reproduced from Ref. 69.

duration and timing of the measurements is limited and they are typically only used in dedicated experiments, limiting the quantity and variety of data they collect. Many of these limitations are overcome by using optical diagnostics which are discussed next.

1.6.2 Optical diagnostics

Optical measurements of filaments most commonly look at line emission from interactions between the SOL plasma and neutral atoms. The intensity of line emission from the collisional excitation of neutral atoms, I_{opt}, is given by [70]:

$$I_{opt} = n_0 f(n, T_e) \tag{1.21}$$

where n_0 is the density of neutral atoms and $f(n, T_e)$ is a function of plasma density and temperature that depends on the ratio of the density of the upper level of the radiative transition to the ground state density. Due to its brightness and use of the native ion species, D_α light (656.28 nm) resulting from the Balmer-alpha $n = 3$ to $n = 2$ transition is commonly used for these measurements. For the case of D_α, $f(n, T_e)$ is given by [70]

$$f(n, T_e) = \left[\frac{n_D(n=3)}{n_D(1s)}\right] A_{3\to 2}, \tag{1.22}$$

where $n_D(n = 3)$ is the density of the excited state, $n_D(1s) = n_0$ is the density of the ground state and $A_{3\to 2}$ is the Einstein coefficient for the transition. Evaluation of

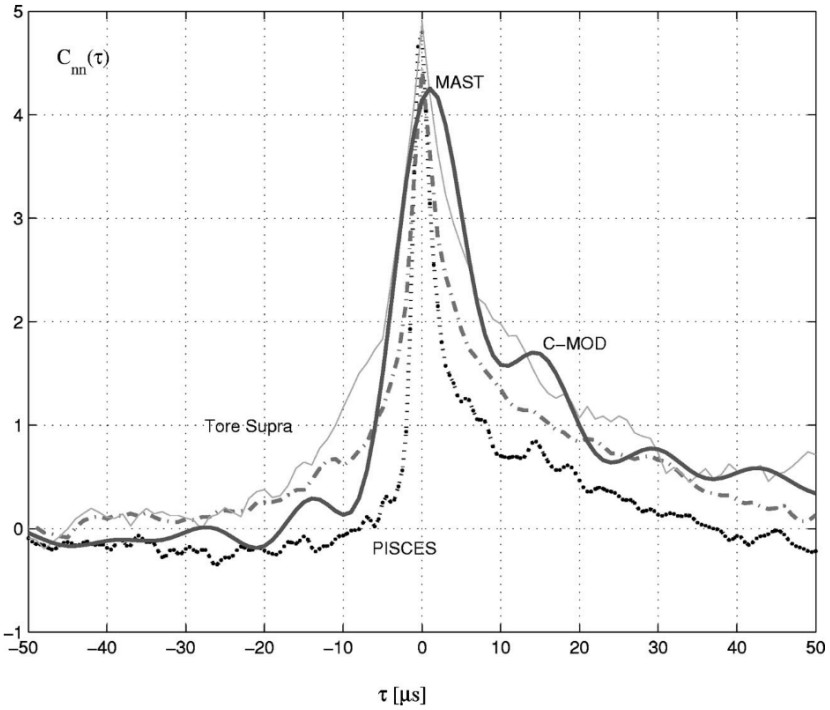

Figure 1.17: Conditionally averaged filament waveforms for Langmuir probe measurements in the the Tore Supra, Alcator C-Mod, MAST and PISCES devices. The integral of the four PDFs have been normalised to 1. Reproduced from Ref. 69.

$f(n, T_e)$ requires solving a collisional radiative model, but can be approximated at a given point in parameter space by $f(n, T_e) \propto n^{\alpha_n} T_e^{\alpha_T}$. At densities of $n \leq 10^{19}$ m^{-3} relevant to the SOL, $\alpha_n \approx 1$, while $\alpha_T \lesssim 1$ over the temperature range 5-200 eV. Thus, as for I_{sat}, I_{opt} depends linearly on n_e and more weakly on T_e. Therefore, in the absence of independent measurements from another diagnostic to disentangle the effects of n and T_e, I_{opt} can be used as a proxy measurement for the plasma density under the assumption that n_0 is slowly varying. Given that D is almost entirely ionised above ∼100 eV [71], n_0 rapidly drops to negligible values a few centimetres inside the separatrix. Therefore, D_α emission is concentrated in the SOL and edge of the plasma where filamentary density fluctuations can be observed.

Beam emission spectroscopy

Beam emission spectroscopy makes use of heating or dedicated diagnostic neutral beams to measure localised density fluctuations at the intersection points between the neutral beam path and optical lines of sight. Heating neutral beams are typically operated using D, although diagnostic Li beams [72–74] are also widely used. Neutral deuterium in the beam is collisionally excited by electrons in the plasma producing D_α line emission that is Doppler shifted with respect to thermal D_α emission from the bulk edge plasma, enabling it to be filtered and isolated [75]. The intensity of the light measured by

photodiodes is, as previously discussed, predominantly dependent on the plasma density, so it can be used to measure density fluctuations, correlation lengths, decorrelation rates, mean flow rates, flow fluctuations, and particle fluxes [75]. Optical sight lines are typically arranged in radial, poloidal or 2D arrays and are aligned tangential to the local magnetic field lines in the beam path to optimise the cross-field spatial resolution. BES systems are installed on many large tokamaks including MAST [76], JET [77], EAST [74], TCV [78], COMPASS [79] and TEXTOR [72], although only those on NSTX [80, 81], DIII-D [48, 64, 82, 83] are appropriate and have been used for in-depth edge turbulence studies.

An example of the poloidal layout of BES channels in the NSTX tokamak is given in Figure 1.18, showing coverage of both the SOL and plasma edge. By comparing the signals across the array of channels, a 2D map of density fluctuations can be produced as shown in Figure 1.19. However, the narrow spatial extent ($< 25 \times 20$ cm) and low spatial resolution ($\sim$3 cm) of the measurements mean that BES measurements are rarely used this way. Typically comparisons of the fluctuation statistics between channels are performed to calculate correlation lengths, decorrelation times and blob velocities [81]. Therefore, as for Langmuir probes measurements of filament dynamics are limited. BES is also subject to temporal and spacial limitations on its application. Measurements can only be made when the beams are on, perturbing the plasma and limiting the plasma scenarios that can be explored. The location of the measurements is limited by the position of the beam, the magnetic geometry and the available lines of sight, limiting the flexibility of the diagnostic.

1.6.3 Gas puff imaging

Gas puff imaging (GPI) produces images of edge turbulence in the radial vs. poloidal plane near the outer edge of the plasma. A gas nozzle or manifold, usually located near the outer mid-plane, is used to inject a cloud of D, He or N_2 gas, which enters the SOL and interacts with the plasma edge, dissociates and becomes collisionally excited by plasma electrons [80]. The excited atoms radiatively decay, producing visible light which is observed by a camera with a line of sight parallel to the magnetic field at the location of the gas puff. A filter is used to isolate the light from the injected gas. The radial resolution is limited by the toroidal size of the gas cloud and the curvature of the magnetic field lines. Therefore, for a fixed GPI set-up only a limited range of magnetic geometries give lines of sight sufficiently tangent to the magnetic field so as to achieve sufficient spatial resolution. The size of the window in the poloidal plane that can be imaged is restricted by the size of the gas plume and the local magnetic shear, limiting viewing areas to a few tens of centimetres squared [84]. Modelling is usually required to understand the propagation and dissipation of the gas plume, so that the transient neutral density profile in the measurement region is known.

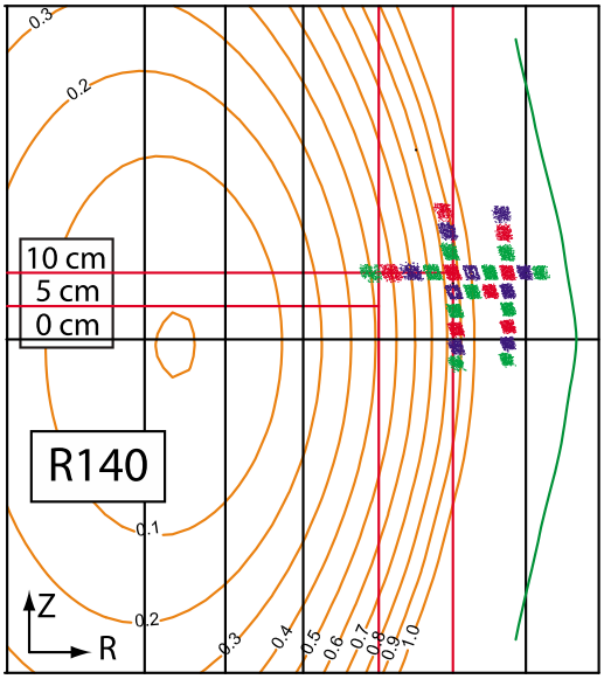

Figure 1.18: Poloidal channel layout for the BES diagnostic on NSTX overlaid on typical flux surfaces. Reproduced from Ref. 75.

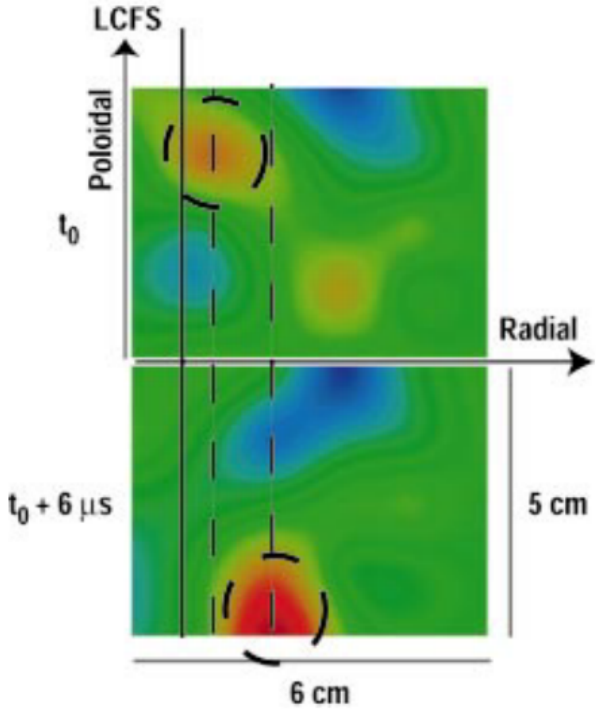

Figure 1.19: 2D map of density fluctuations from BES measurements on DIII-D. Red indicates high density and blue low density. A filament is labelled with a dashed circle, showing radial and poloidal motion over the 6 μs between measurements. Reproduced from Ref. 48.

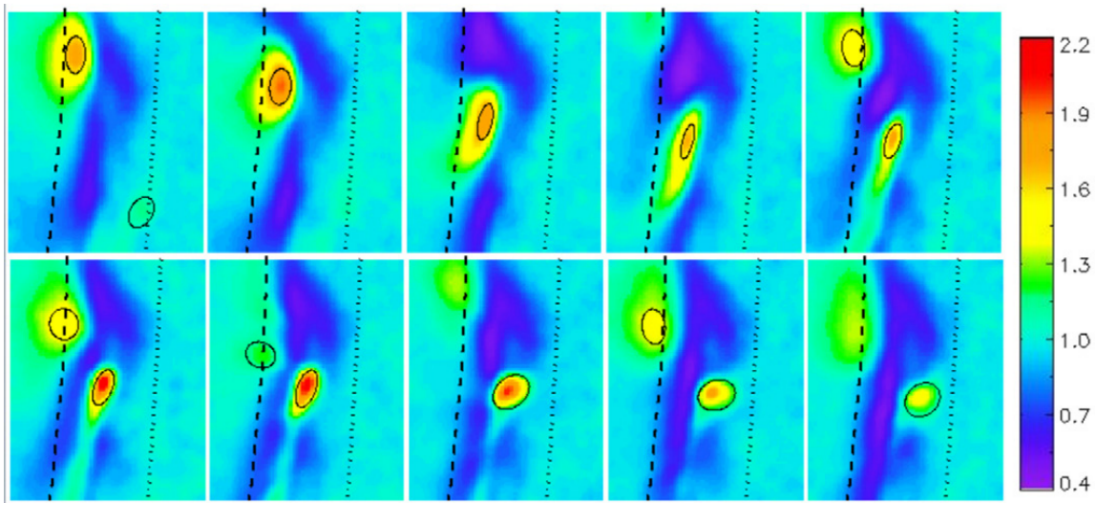

Figure 1.20: Examples of GPI D_α light emission for an NBI heated H-mode discharge in NSTX. The GPI view covers an area spanning ∼24 cm radially and ∼30 cm poloidally at a frame rate of 100 kHz. The images have been normalised relative to the time averaged image. Several filaments are shown with black ellipses, while the separatrix and limiter shadow are shown as dashed and dotted lines respectively. Reproduced from Ref. 87.

GPI systems have been installed on NSTX [80,81,84–87], Alcator C-Mod [85,88–90] and ASDEX-U [91], amongst others, and can achieve spatial resolutions of around ∼0.5 cm and temporal resolutions of up to 2.5 μs [81]. Figure 1.20 gives an example of GPI images taken in NSTX, with filaments labelled with black ellipses. GPI produces very high quality data from which large quantities of filaments can be identified and tracked. Measurements of the same flux-tube with Langmuir probes and GPI have shown highly correlated signals [92], confirming that both diagnostics show manifestations of the same filamentary structures.

GPI's main limitations are that measurements are confined spatially to a small section of the plasma that can be viewed tangentially to the magnetic field and temporally to the duration of a gas puff ∼80 ms. As these gas puffs are also perturbative, they are not used routinely.

1.6.4 Direct wide-angle visible imaging

Direct visible imaging operates on the same principals as GPI, with the difference being that DVI relies passively on the presence of native recycled neutrals at the plasma edge, rather than the active injection of gas at a fixed location. The condition of sufficiently high passive edge neutral density is not easily satisfied in many machines, with the consequence being that passive D_α emission at the edge can be too faint to resolve filaments with a sufficient signal-to-noise ratio. Therefore DVI benefits from large gaps between the plasma and the wall, providing space for the redistribution of neutrals from the main

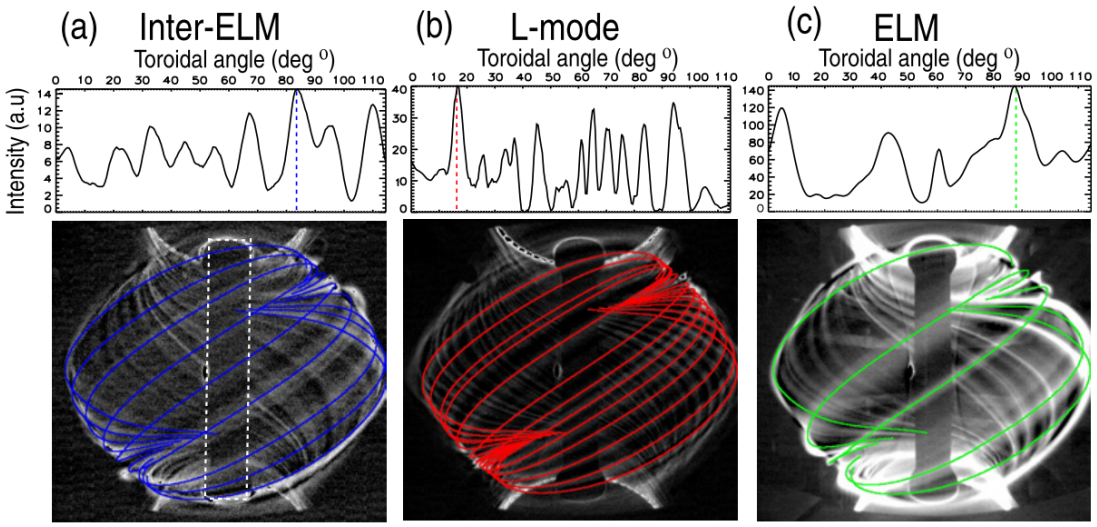

Figure 1.21: Examples wide-angle fast visible imaging of inter-ELM, L-mode and ELM filaments in MAST. While the whole plasma is imaged in this case, only the variation in emission in the white dashed box around the centre column is used to locate the filaments. The panels above the images show the toroidal variation in emission across the centre column, whose peaks are used to label the filaments. Radial information about the filaments cannot be recovered with this technique. Reproduced from Ref. 97.

regions of PSI and thus higher main chamber neutral pressures. Wide-angle imaging of the plasma is also improved in tight aspect ratio machines for which large regions of the plasma can be viewed at once. As a result, the most studies with DVI have been on MAST [47, 51, 68, 93–99], which is well suited to these measurements due to its compact spherical design and lack of a plasma conforming first wall. However, fast camera measurements of filamentary structures have been made on a broad range of tokamaks including TCV [100], COMPASS [101], KSTAR [102], ASDEX [103], TORPEX [104] and TFTR [71].

Fast framing camera technology has improved dramatically since the earliest visual observations of filaments in the 1980s [71, 105]. Modern fast framing cameras are now capable of recording the whole plasma at resolutions, frame rates and exposure times sufficient for resolving filament motion.

Past implementations of direct wide angle visible imaging

The flexibility of fast camera measurements have enabled pixel-wise statistical analysis of filaments in the divertor regions of MAST and TCV [68, 100]. These studies have explored the differing properties of filaments in the near and far SOL and private flux regions of the divertor legs. Upstream filaments are observed to be heavily sheared as they pass the X-point, while different types of filaments are born in the private flux region and a quiescent region is observed in the near-SOL. However, while these measurements shed light on the processes affecting filaments in different regions of the

machine, they do not explore the dynamics of individual filaments. Past fast camera analyses of individual filaments have typically either been performed manually, making them time consuming and subjective, or have achieved automation of the measurement of a limited subset of filament parameters. Past simultaneous measurements on MAST of both toroidal and radial filament properties (positions, widths and velocities) [47, 94, 97] have involved the manual alignment of projected field lines with filamentary structures in camera images by eye; a challenging exercise owing to their translucent, diffuse nature and their interactions with one another, greatly reducing the quantity of data that has been analysed. Automated or semi-automated measurements of radial (toroidal) positions, widths and velocities have previously been made by considering the time varying intensity of specific radial [51,94,97] (toroidal [94,103]) arcs of image pixels to infer the passage of filaments along the length of the arc. These techniques rely on finding areas of an image where the camera line of sight is such that pixel coordinates along a line can approximately be mapped directly to radial (toroidal) coordinates and typically assume that only one filament is passing the region at any one time. More sophisticated techniques set a fixed radial position *a priori*, typically close to or at the separatrix, and located maxima in the average intensity of superimposed field lines as a function of their toroidal position [47,51,96,97]. For instance, Figure 1.21 shows how toroidal variation in the emission across the centre column of MAST has been used to automatically toroidally locate filaments. However, since this technique relies on treating all filaments as being located on the separatrix, their radial positions, widths and velocities could not be recovered. A consequence of this is that uncertainty in the precise radial location of the filaments would incur greater uncertainties in the toroidal measurements [47]. Furthermore, with the exception of Ref. [51], all applications of these autonomous or semi-autonomous techniques have focused on narrow, windowed camera views of specific areas of the plasma (e.g. the outboard mid-plane [96,97], centre column [97] or top outside corner of the plasma [94]), in order to maximise the camera's frame rate, thus limiting the number of observable filaments and restricting the length of the filaments along which parallel structure can be explored.

Fast camera measurements on Alcator C-mod have been used to study filament motion near the X-point [106]. A grid of field lines intersecting different points in the R-Z plane was used to identify the locations of filaments in the images. The field lines were picked by hand in a "labour-intensive quantitative procedure", limiting the quantity of data that could be gathered and preventing an assessment of filament widths and amplitudes.

Very recently a group in Prague has published a technique for tomographic inversion of fast camera data in the COMPASS tokamak [101]. This work built upon an earlier paper [107] which presented a complex wavelet-vaguelette decomposition technique for the inversion of helically structured emission. Combined with the proprietary TRACK

[108] blob tracking software, the technique has been used to make measurements of radial and poloidal velocity distributions of L-mode filaments.

1.6.5 Benefits of inverting fast camera data

In machines such as MAST, where the edge neutral density is high enough to facilitate it, direct wide-angle fast visible imaging of the plasma holds great potential for scrape-off layer studies. Visible imaging benefits from being a passive technique that can operate reliably over a wide range of conditions, for the whole duration of a pulse, while measuring a large section of the plasma.

This leads to a number of advantages with respect to LP, BES and GPI diagnostics, an overview of which are given in Table 1.1. Firstly, the measurements are not constrained to a single toroidal location so the full 3D structure of the filaments can be explored. This opens up possibilities for studying the structure of the filaments parallel to the magnetic field, as well as measuring the toroidal distribution of the filaments. Furthermore, a much wider region of plasma can be observed, without the constraints of gas puff nozzle proximity, beam path or local magnetic geometry. Therefore, measurements need not be isolated to a small portion of the SOL, meaning variations in filament properties with poloidal position can be explored, both at the mid-plane and in the divertor. The fact that the system can operate in all magnetic configurations and is passive, not requiring perturbative probe, gas puff or beam injection, means it can be run as standard on all shots, collecting far more data, over a greater range of conditions. The passive nature of the measurement also reduces aspects of its complexity, avoiding the complications of understanding complex non-equilibrium phenomena such as gas plume propagation or neutral beam-plasma interaction, instead relying on relatively uniform and slowly varying excitation and ionisation rates. The trade off with this increased information and flexibility is greater complexity in processing and interpreting the line integrated emission in the images.

This thesis describes a new tomographic inversion technique that has been developed to automate the identification and analysis of filaments in fast camera data, with the goal of retrieving as much individual filament data as possible. By inverting the emission in the camera images onto a set of field aligned basis functions, both the radial and toroidal properties of the filaments can be found together. The technique is designed to simultaneously measure the positions, widths, amplitudes and velocities of individual filaments, across a large extent of the plasma. Further quantities can be derived from these measurements such as a quasi-toroidal mode number and the distributions of filament separations and waiting times. As all quantities can be measured for individual filaments, inter-relationships between different filament properties can be investigated. In this way a large database of filament properties can ultimately be generated, enabling big data approaches to the understanding of the statistical properties

	Langmuir probe	Beam emission spectroscopy	Gas puff imaging	Direct visible imaging
Measurement parameter	Ion saturation current	Emission from NBI	Emission from local neutral gas puff	Unfiltered D_α light
Spatial dimensions	0D/1D	1D/2D	2D	3D
Temporal durations [ms]	Reciprocation ∼400	Duration of beam	Duration of gas puff ∼80	Whole pulse
Spatial extent [cm] (radial×toroidal)	20	25×20	24×30	25×100
Temporal resolution [MHz]	∼2+	∼0.4	≤ 0.4	≤ 0.2
Spatial resolution [cm]	N/A	∼3	∼0.4	∼0.5
Perturbative/ passive	Requires probe reciprocation	Requires NBI	Requires gas puff	Passive

Table 1.1: Comparison of key attributes for the four main diagnostics used for filament measurements. Representative values for the LP, BES and GPI diagnostics have been taken from [109], [81] and [87] respectively. The maximum DVI camera frame rate of 200 kHz has been given for the Photron SAX-2 camera that will be used on MAST-U.

of filament parameters and their dependence on engineering and physics parameters. A good understanding of these statistics will enable the validation of analytic frameworks such as that developed in Ref. [45] and provide the inputs for these frameworks, greatly improving our understanding of the filamentary generation of SOL density profiles.

1.7 Thesis outline

In Chapter 2 a review of scrape off layer filamentary physics is presented, covering the birth, motion and typical parameters of SOL filaments. The existing theoretical body of work on filaments from the literature is discussed, both on analytic descriptions of filaments and modelling using complex turbulence codes.

The experimental setup concerning the work in this thesis is introduced in Chapter 3. Details are given of the camera diagnostic used to observe the filaments, together with the field of view analysed and the procedure for spatially calibrating the 3D location of structures in the cameras field of view. This is followed by a an overview of the experimental discharges that have been analysed.

The tomographic inversion technique that has been developed to identity and analyse filaments in fast camera images is presented in Chapter 4. This is broken down into the preprocessing, matrix inversion and filament detection steps of the analysis procedure.

Chapter 5 discusses the work that has been carried out to benchmark the new in-

version technique and quantify its associated errors and artefacts. A synthetic cameras diagnostic is introduced that is used to test the inversion technique against data containing filaments with known properties. This data is utilised to assess the accuracy with which the probability density functions of filament parameters can be recovered.

The technique is then applied to experimental data in Chapter 6. Langmuir probe analysis techniques are applied to data from both the camera and reciprocating probe diagnostics in order to perform a direct comparison. This is done to establish the effect of the measurement technique on the statistical conclusions that can be drawn and provide a basis for the comparison of new fast camera analysis results with existing Langmuir probe measurements in literature. Fast camera measurements are described that facilitated a benchmarking exercise comparing four turbulence codes, which, in addition to assessing the capabilities of the codes, further elucidated the experimental measurements. Unique measurements are presented of the toroidal distribution of the filaments which have important implications for the verification of analytic filament frameworks. Experimental PDFs of filament positions, widths, amplitudes, velocities and lifetimes are presented and these results are interpreted with the insight gained from the synthetic benchmarking of the technique in Chapter 5.

Finally, Chapter 7 draws conclusions, summarising the main results from the preceding chapters. The potential for application of the inversion technique to future data from the MAST Upgrade tokamak is discussed, along with other regimes of interest for analysis and potential further improvements to the analysis technique.

Chapter 2

Review of Scrape Off Layer Filament Physics

2.1 Introduction

Filaments feature in the literature under a variety of names including: blobs, intermittent plasma objects (IPOs), radial streamers and avaloids, generally depending on the manner in which they were observed, whether as point density fluctuations or regions of raised density in the poloidal plane. While high amplitude intermittent fluctuations had been observed in Langmuir probe measurements for some time, the first observation of coherent plasma structures were with analogue burst fast cameras in the mid 1980's [110] and 2D probe arrays in the mid 1990's [111, 112].

Despite these early observations of filaments, it was not until after the discovery of the "main chamber recycling regime" in Alcator C-Mod in 1998 [113] that their importance in edge transport was considered in detail. Examination of the particle balance at high line averaged density revealed that, contrary to expectations at the time, plasma flow in the SOL was not dominated by parallel flows to the divertor, but rather by perpendicular transport to the main chamber walls [41]. The historical description of perpendicular transport due solely to diffusion required huge increases in $D_\perp$ across the SOL to unphysical levels, similar to modelling results on ASDEX-Upgrade [114]. The realisation of the lack of direct external control over density profiles and mid-plane neutral pressures, led to concerns about increased sputtering of the first wall in a reactor. The failure of the existing theoretical models led to the conjecture that it was "likely that the transport mechanism is convective in nature, perhaps involving a rapid transport of unstable plasma flux tubes toward the wall surfaces" [113]. These results triggered intensive studies of transport and turbulence in the plasma edge, so that these ideas were quickly expounded by others [57]. The first qualitative theory of individual filament dynamics was proposed shortly afterwards by Krasheninnikov, describing how polarisation of filaments drives their radial advection [115]. This was quickly built upon to include the transport of heat and vorticity and to discuss the

consequences for SOL profiles and neutral recycling.

Filaments have since been observed in over 40 devices, including both toroidal magnetic fusion devices and basic plasma devices, using a range of diagnostics [49]. In particular, camera data from MAST, QUEST [116], Alcator C-mod [106] and COMPASS [101] have unambiguously shown the presence and prevalence and 3D structure of filamentary structures in the SOL.

Filament definition

There are a range of definitions in the literature as to what constitutes a filament, although many of them take a similar approach to D'Ippolito *et al.* [49] in defining them as plasma structures with the following defining properties:

1. A positive single-peaked density distribution with peak density values significantly higher than the surrounding root-mean-square (rms) fluctuations of the background plasma.

2. Alignment to the local magnetic field, $\mathbf{B}$, such that their dimensions and scale lengths parallel to the magnetic field are far larger than those perpendicular to it.

3. Radial $\mathbf{E} \times \mathbf{B}$ motion due to charge polarisation arising from an outward species summed net force ('bad curvature' effects in the case of tokamaks).

Each of these criteria warrant some further discussion. In many cases, particularly in Langmuir probe studies, condition (1) is made more stringent such that filaments are taken to be only those density fluctuations that are 2-3 times higher than the rms background level. Such a definition can be useful in these cases, due to the cleaner statistics produced by limiting the analysis to only the largest perturbations, and the fact that it circumvents the difficulty in fully asserting conditions (2) and (3). However, such a rigid definition is not necessary. In this thesis we shall avoid setting an arbitrary defining threshold on the required amplitude of the fluctuations, in order to include all structures that exhibit the same common properties and behaviours, acknowledging that diagnostic limitations will always make measurement of low amplitude filaments most challenging. Condition (1) is also responsible for filaments producing positively skewed fluctuation statistics, such that the presence of these statistics is often taken to be synonymous with that of filaments [117].

Condition (2) is a consequence of the disparity between the timescales of the parallel and perpendicular transport transport mechanism discussed in §1.3.2. As the plasma tends to stream rapidly along field lines, parallel non-uniformities tend to be dispersed at the ion sound speed, while perpendicular structures vary over tens of microseconds. The condition is often alternatively phrased in terms of filaments being 'mesoscale' structures, in reference to the fact that filaments' perpendicular dimensions are intermediate between the ion gyroradius and the macroscopic machine dimensions.

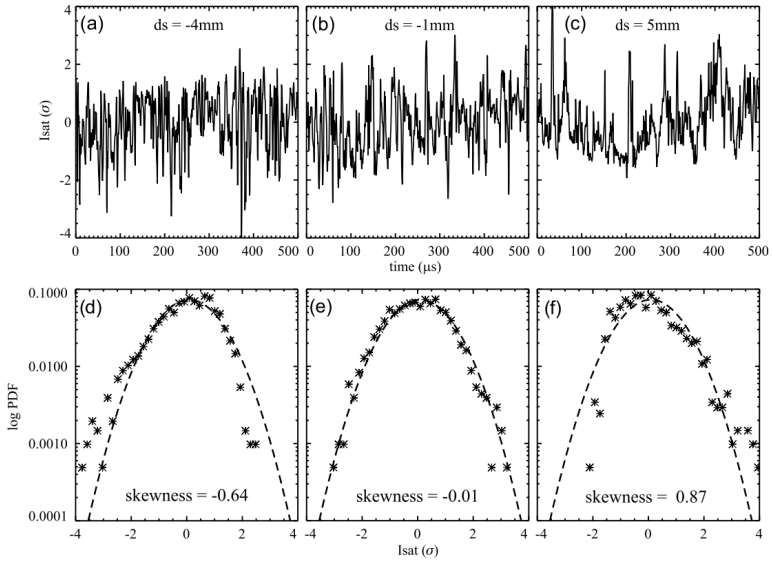

Figure 2.1: Time series of ion saturation current (a), (b), (c) and their respective fluctuation PDFs (d), (e), (f) for three radial displacements from the separatrix, $ds = -4, -1, 5$ mm, in ASDEX-U. Measurements at $ds = 5$ mm exhibit positive skewness indicating the presence of filaments, while measurements at $ds = -4$ mm show negative skewness characteristic of holes. The I_{sat} signals have been normalised with respect to their standard deviations. Gaussian curves with zero skewness are shown as dashed lines. Reproduced from Ref. 119.

The motion of the filaments captured in condition (3) is responsible for much of the complicated dynamics that the filaments exhibit. It is how this radial $\mathbf{E} \times \mathbf{B}$ drift depends on and competes with external SOL conditions that in large part determines the shape of the SOL profiles and the maximum radial distance the filaments can travel before they drain along the magnetic field lines. Close to the separatrix the filaments are in a region dominated by small scale turbulence where the magnitudes of the filaments are not substantially higher than the local near-diffusive background, producing Gaussian-like statistics and short density scale lengths. However, in the far SOL the filaments' relative contributions to the density are far higher, resulting in a second scale length determined by the filament dynamics and strongly positively skewed statistics.

In analogy to positive filament density perturbations, holes (alternatively 'dips') can be defined by the same three criteria, but instead with, (1) a *negative* single-peaked density distribution, and (3) an *inward* species summed net force, due to the reversal in direction of the $\nabla \mathbf{B}$ force on a negative density perturbation [118]. Due to the inward propagation of holes from their formation region near the separatrix, they are responsible for producing *negatively* skewed fluctuation statistics inside the separatrix. An example of the variation in skewness of fluctuation PDFs is given in Figure 2.1 for Langmuir probe measurements in ASDEX Upgrade.

Holes potentially provide an important mechanism for transporting impurities from

the edge into the core [48]. However, unlike filaments whose maximum amplitudes can be as high as the local edge density, holes are necessarily smaller due to the fact that negative density fluctuations cannot exceed the magnitude of the background level, which is small. This, combined with the fact that modelling has shown that vorticity effects tend to slow down the inward propagation of holes and break them up, means that the importance of holes in impurity transport is still unclear [120].

2.2 Typical filament parameters

Observations of filaments are mainly confined to the low field 'bad curvature' side and divertor regions of toroidal fusion devices, with filaments not observed in the high field side inner SOL [106]. Plasma filaments observed at the mid-plane are typically divided into the three categories of ELM, inter-ELM and L-mode filaments. Of these, ELM filaments have the largest amplitudes, velocities and lifetimes and display strong radial accelerations [97]. While L-mode and inter-ELM filaments are produced by turbulence, ELMs are a result of MHD instabilities. ELM filaments are highly energetic and carry large parallel current currents, while L-mode filaments are not observed to carry significant current, with Ref. 94 placing an upper limit on L-mode filament currents of 5 A. Therefore, ELMs are expected to experience significant ballooning, distorting the filament's local magnetic field from that of the magnetic equilibrium. As shall be discussed in detail in Chapter 4, the Elzar fast camera analysis technique developed in this thesis makes use of the property that current free filaments are well aligned with the equilibrium magnetic field. Therefore ELM filaments are not prime candidate for analysis with this technique and shall not be discussed further here. L-mode and inter-ELM filaments have largely similar properties to one another, although L-mode filaments typically have slightly higher amplitudes, radial widths and effective toroidal mode numbers, n_{eff}, where n_{eff} corresponds to the average number of filaments observed per toroidal circuit [97]. The greater brightness, frequency and size of L-mode filaments, combined with the simpler L-mode edge physics, which lacks a complicating pedestal and edge transport barrier, makes L-mode filaments ideal objects to study with the Elzar technique. Therefore, we shall focus here on L-mode filament properties.

Table 2.1 provides a breakdown of typical important L-mode filament parameter observed in a range of tokamaks and a toroidal plasma device (TORPEX), using a mix of probe and optical diagnostics. As shall be discussed in §2.3, these parameters are of particular importance for a theoretical understanding of the role of filaments in SOL particle transport. The data in Table 2.1 was primarily compiled in Refs. [121] and [122]. Despite a wide variation in geometries and operational parameters in these machines there is a great deal of commonality between the filaments found in each device, indicating common physical origins and dynamics in each.

In general, filaments are observed to survive for tens of microseconds, equivalent

Device	Perpendicular width [cm]	Density fluctuation, $\delta n/n_{bg}$	Radial velocity [km s^{-1}]	Lifetime [μs]	Sources
MAST	0.5-10	0.1-4	0.5-2.0	40-60	[51, 96]
Alcator C-Mod	0.5-1.5	-	0.15-1.5	1-50	[88, 89, 106, 125]
NSTX	3-5	-	0.2-1.4	60-70	[80, 86]
DIII-D	0.5-3.9	0.05-1	0.33-2.6	15-20	[64, 126]
JET	0.5-2	0.3-1.5	0.2-1.5	20-130	[118, 127]
ASDEX-U	0.2-7	-	0.1-3.5	-	[128]
TORPEX	1.3-2.5	0.6-0.85	0.2-1.85	-	[129, 130]

Table 2.1: Comparison of L-mode filament parameters across a range of devices.

to ∼tens of ion collision times, ν_i, [106] and have density fluctuation amplitudes δn of the order of the background plasma density, n_{bg}, with $\delta n/n_{bg}$ largest in the far SOL. The filaments are seen to typically have perpendicular widths in the range 0.5-5 cm corresponding to ∼ 1-100 ion Larmor radii [122] and radial velocities of up to several kilometres per second, or around 1-10% of the sound speed, c_s.

As filaments propagate through the SOL, they drain through parallel losses, leading to a reduction in amplitude which has also been observed to be accompanied by a reduction in average filament size toward the walls [123]. In some machines, the filaments' parallel correlation lengths have been observed to be more than a full toroidal turn [106]. Langmuir probe measurement have shown that in L-mode T_e and n fluctuations are largely in phase [117]. However, due to the speed of parallel electron heat conduction compared to particle drainage times, filaments' temperatures normally decay faster than their densities. Correspondingly, at low collisionalities [124], the temperature decay length in the SOL is generally shorter than the density decay length [48]. This difference is further increased through the ionization of neutrals which can further slow the parallel loss of particles, while enhancing energy losses [123]. A further consequence of the high electron heat conduction is that much of the energy flux coming into SOL from the core is transported rapidly to the divertor by electrons. Therefore, while filaments are the dominant driver of particle transport, their effect on energy transport is only through the ion component of the temperature, so that their role in heat transport is more complicated.

2.3 Filament theory

The first analytic models of 'blobs' were produced in the early 2000's by Krasheninnikov [115] and D'Ippolito *et al.* [123]. More recently stochastic frameworks have been developed for interpreting SOL profiles in terms of filament dynamics [45, 131].

2.3.1 Filament generation

L-mode filaments are believed to be formed as part of the non-linear saturation process of edge instabilities, particularly drift-resistive interchange instabilities. These processes are not well understood and remain an active area of investigation.

To date there have not been any quantitatively accurate analytical models of filament generation [49]. Therefore insight must come from experimental data and computer simulations.

In plasmas devoid of significant poloidal sheared flows the turbulence manifests as 'radial streamers' which extend out from the separatrix, but remain connected to the birth zone. For individual filamentary behaviour to be observed, the time for shearing of the radial streamers into disconnected filaments must be shorter than the radial convection time. For circular cross section filaments, this has been shown to occur when the velocity shear gradient, v'_y, is greater than the growth rate, y, of the underlying linear instability producing the turbulence [132]. The shear flows found in L-mode plasmas typically satisfy this criteria. However, if the shear flows are too strong, as in H-mode, simulations have shown they act to suppress the transport of the filaments by either tearing them apart or trapping them in the sheared-flow layer [133]. Momentum transport is a poorly understood field owing to the complexity of the full viscosity terms in the Braginski equation that are rarely modelled in full and the difficulty developing detailed models of turbulent moment transport. However, it is believed that sheared flows can be initiated by momentum transport by ReynoldâĂŹs stress. Filaments can then potentially lead to further transport of momentum across the separatrix to where it is lost, reinforcing the sheared flows resulting in a coupled self-regulating system [49].

Equal numbers of filaments and holes are generated in the birth zone, reflected in the near zero skewness of fluctuations in the birth region. The filaments and holes then separate and propagate in opposite directions along the magnetic field and density gradients resulting in the positive and negative regions of skewness either side of the birth region.

Filaments are likely born inside the separatrix, although large experimental uncertainties in the position of the last closed flux surface, which are typically of the order of centimetres, mean that this has been difficult to confirm. Therefore, observations in Alcator C-mod of the location of filament birth moving inwards, crossing from open to closed field lines as the density is increased for a fixed magnetic geometry [134] may be due to uncertainties in the magnetic equilibrium. However, a transition from closed to open field lines is not the only factor playing a role in determining the location of the filament birth region. Similar turbulent filament generation dynamics have been observed in the TOPRPEX toroidal plasma device, despite all field lines being open and terminating on solid surfaces at the ends of the device [135]. Rather, the pressure gradient profile is important in determining the location of the maximum in the tur-

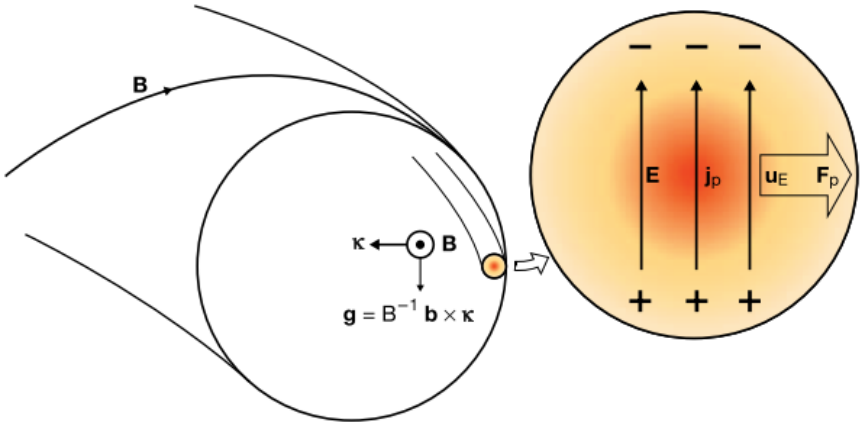

Figure 2.2: Illustration of charge polarisation of filaments leading to radial $\mathbf{E} \times \mathbf{B}$ propagation. Bad curvature drifts result in a species dependent force, $\mathbf{g}$, perpendicular to the curvature vector, $\boldsymbol{\kappa}$. The resulting charge polarisation current, $\mathbf{j}_p$, produces an electric field $\mathbf{E}$, which exerts a net species summed force, $\mathbf{F}_p$, on the filament plasma, which moves the filament radially at the $\mathbf{E} \times \mathbf{B}$ velocity, $\mathbf{u}_E$. Red illustrates high densities and yellow lower densities. Reproduced from Ref. 141.

bulent instability's linear growth rate, which in turn determines where the filaments are born. Several experiments [86, 136] have confirmed through skewness and conditional averaging measurements that filament generation occurs at the maximum in the logarithmic pressure gradient $-d(\ln p)/dr$, where p is the pressure and r is the radial coordinate. In tokamaks this typically occurs just inside the separatrix.

2.3.2 Filament dynamics

The radial motion of filaments is of great importance as it enables them to transport particles [115], heat [123,137], momentum [138,139] and parallel current [140] across the SOL. Most attention has been devoted to particle transport as it is most easily measured and has the greatest effect on SOL widths. By contrast, the other plasma parameters decay much more quickly via parallel transport and result in smaller 'anomalous' differences relative to classical diffusive transport [49]. Therefore we shall focus here on particle transport.

Filaments are approximated as ballistically propagating flux tubes that travel radially across field lines. As previously touched upon, filament radial motion is driven by $\mathbf{E}\times\mathbf{B}$ advection, with a corresponding radial drift velocity $\mathbf{v}_{E\times B}$ given by Equation 1.11 of the order of ~ 1 km s^{-1} [87]. The electric field across the filament is maintained by charge dependant curvature and $\nabla\mathbf{B}$ drifts illustrated in Figure 2.2, which cause electrons and ions to drift in opposite directions. A filament's radial velocity therefore depends on the strength of the electric field across the filament. The magnitude of the electric field depends on the closure of parallel currents that dissipate the diamagnetic current which generates the filament's electrostatic dipole. Common models of filament

motion are built on the principle that, for quasi-neutrality of the plasma to be maintained and no large charges to accumulate, there can be no divergence of the current density $\mathbf{J}$ in the plasma. This is expressed by the current continuity equation

$$\nabla \cdot \mathbf{J} = \nabla \cdot \mathbf{J}_\perp + \nabla \cdot \mathbf{J}_\parallel = 0, \tag{2.1}$$

where $\mathbf{J}_\perp$ and $\mathbf{J}_\parallel$ respectively are the perpendicular and parallel components of the current density relative to the magnetic field. Therefore, the currents in filaments must form closed loop circuits, with the size of $\mathbf{J}_\perp$ determined by the non-linear properties of the circuit.

An expression for the perpendicular current density can be found by considering the plasma force balance equation:

$$m_i n \frac{d\mathbf{v}}{dt} = -\nabla p + \mathbf{J} \times \mathbf{B}. \tag{2.2}$$

where $p = nT$ is the plasma pressure. Multiplying both sides of Equation 2.2 by $\times \mathbf{B}/B^2$ yields

$$\begin{aligned} \frac{m_i n \frac{d\mathbf{v}}{dt} \times \mathbf{B}}{B^2} &= -\frac{\nabla p \times \mathbf{B}}{B^2} - \left(\frac{(\mathbf{B} \cdot \mathbf{B})\mathbf{J}}{B^2} - \frac{(\mathbf{J} \cdot \mathbf{B})\mathbf{B}}{B^2} \right) \\ &= -\frac{\nabla p \times \mathbf{B}}{B^2} - \mathbf{J}_\perp. \end{aligned} \tag{2.3}$$

which on rearranging for $\mathbf{J}_\perp$ and substituting into Equation 2.1 results in:

$$\nabla \cdot \underbrace{\left(\frac{mn \frac{d\mathbf{v}}{dt} \times \mathbf{B}}{B^2} \right)}_{\mathbf{J}_p} = \nabla \cdot \underbrace{\left(J_\parallel \frac{\mathbf{B}}{B} \right)}_{\mathbf{J}_\parallel} - \nabla \cdot \underbrace{\left(\nabla p \times \frac{\mathbf{B}}{B^2} \right)}_{\mathbf{J}_d}. \tag{2.4}$$

Thus the perpendicular current has been broken down into diamagnetic, $\mathbf{J}_d$, and polarisation, $\mathbf{J}_p$, components: $\mathbf{J}_\perp = \mathbf{J}_p + \mathbf{J}_d$. The divergence of the diamagnetic current can be re-expressed as follows:

$$\begin{aligned} \nabla \cdot \mathbf{J}_d &= \nabla \cdot \left(\cancel{en} \left[\frac{\nabla p \times \mathbf{B}}{\cancel{en} B^2} \right] \right) \\ &= \cancel{\frac{\mathbf{b}}{B} \cdot \nabla \times \nabla p} - \nabla p \cdot \nabla \times \frac{\mathbf{b}}{B} \\ &\approx -\nabla p \cdot 2 \frac{\mathbf{b} \times (\mathbf{b} \cdot \nabla)\mathbf{b}}{B} \\ &= -\nabla p \cdot 2 \frac{\mathbf{b} \times \boldsymbol{\kappa}}{B}, \end{aligned} \tag{2.5}$$

where $\mathbf{b} = \frac{\mathbf{B}}{B}$ is the unit vector along the magnetic field, $\boldsymbol{\kappa}$ is the curvature vector of the magnetic field and in the second to last step we have made the assumption of small plasma beta, $\beta = \frac{p}{B^2/2\mu_0} \ll 1$, where μ_0 is the permeability of free space. Therefore, we see that in a tokamak the diamagnetic current arises due to magnetic

curvature. Expressing the diamagnetic current in the form, $\mathbf{J}_d = \frac{1}{B^2}\mathbf{F} \times \mathbf{B}$ it can be seen that the effect of the magnetic curvature is equivalent to an 'effective gravity' force, $\mathbf{F} = nm_i\mathbf{g}$, producing a drift $\mathbf{v}_g = -\frac{m\mathbf{g}\times\mathbf{B}}{eB^2}$, in a system without magnetic curvature. Taking $\mathbf{g}$ and $\mathbf{B}$ to be uniform and the plasma to be isothermal we find that, $\nabla \cdot \mathbf{J}_g = \nabla \cdot (mn\mathbf{v}_g) = -\nabla n \cdot m\frac{\mathbf{b}\times\mathbf{g}}{B}$, where $\mathbf{g}$ has taken on the roll of $\boldsymbol{\kappa}$ in Equation 2.5. In devices other than tokamaks, the diamagnetic current can be driven by alternative non-curvature based mechanisms such as thermal gradients [142, 142] or the frictional 'neutral particle wind' interactions [143], such as is the case in the PISCES linear plasma device. The diamagnetic drift is very important in tokamak plasmas, as edge gradients in density and temperature are always present. The diamagnetic current acts to cancel out part of the externally applied magnetic field, reducing the strength of the magnetic field in the plasma.

Returning to the polarisation current and substituting, $v = \frac{\mathbf{E}\times\mathbf{B}}{B^2} = \frac{-\nabla\Phi\times\mathbf{B}}{B^2}$, for the drift velocity, using the same identity as in Equation 2.3, we arrive at the following expression for the polarisation current:

$$\mathbf{J}_p = \frac{d}{dt}\left(\frac{nm_i}{B^2}\nabla_\perp\Phi\right), \tag{2.6}$$

thus expressing the current in terms of the electrostatic potential across the filament, Φ. Taking the polarisation current in this form and applying some vector identities to rearranging terms, we arrive at the vorticity equation [49],

$$\nabla_\perp \cdot \frac{d}{dt}\left(\frac{nm_i}{B^2}\nabla_\perp\Phi\right) = B\nabla_\parallel\left(\frac{J_\parallel}{B}\right) + \nabla_\perp p \cdot \nabla \times \left(\frac{\mathbf{B}}{B^2}\right) \tag{2.7}$$

When considering the plasma equations in relation to filament physics, it is common to discard small terms by taking the drift ordering approximation [144] that the flow velocity is $v \sim \delta v_{t,i} \sim \frac{\rho_i}{L} v_{t,i}$, where $v_{t,i}$ is the ion thermal velocity, ρ_i is the ion Larmor radius, L is the scale length of the system and $\delta = \frac{\rho}{L}$ is the ordering parameter. The drift approximation is typically used to describe the slow growing "micro-instabilities" that drive turbulent transport in tokamaks, in contrast to rapid MHD phenomena and therefore are appropriate for the study of filaments [145]. Under the drift ordering assumption, the polarisation drift is typically smaller than the diamagnetic drift by a factor of δ^2 and so the diamagnetic drift dominates in most situations. However, when considering the divergence of the two drifts, a component of the diamagnetic current cancels, so that the polarisation and diamagnetic terms in Equation 2.7 become comparable and the polarisation current becomes important.

Equation 2.7 describes how the combination of perpendicular polarisation currents and parallel currents closing further down the filament, balance the build up of charge from the diamagnetic current source, to determine the strength of the electrostatic potential. The possible current paths through the filament and its bounding solid surface can be visualised in an 'equivalent filament circuit' such as that shown in Figure 2.3. The

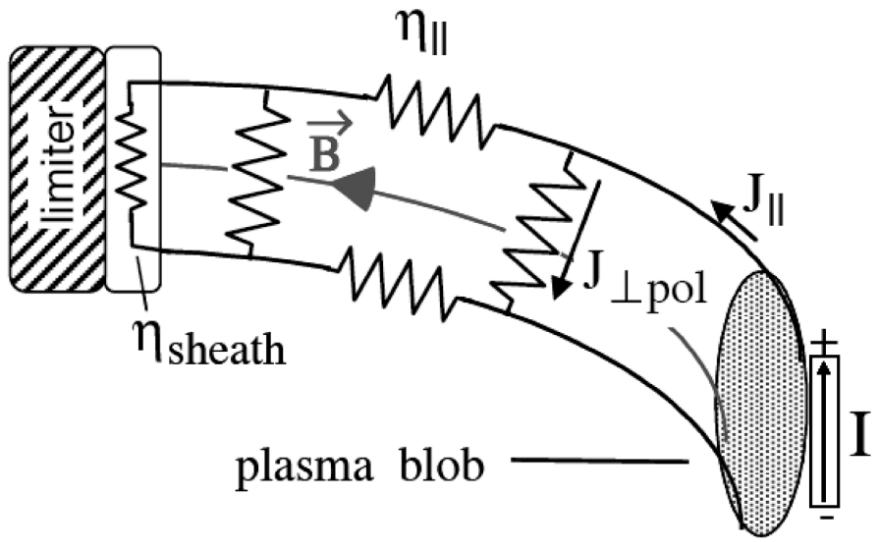

Figure 2.3: Equivalent filament circuit showing the possible perpendicular paths that close the parallel filament currents. The current curvature driven current source at the mid-plane is indicated by I. Reproduced from Ref. 42.

constant current source I at the outer mid-plane is provided by the filament diamagnetic current, while the dominant path of current closure depends on the plasma regime and has consequences for the scaling of the filaments' velocities.

The first qualitative analytic model of individual filament motion was provided in Krasheninnikov's seminal 2001 paper [115] which was able to successfully predict the magnitudes of radial velocities in the *sheath limited* regime. In the sheath limited regime the plasma resistivity $\eta_{\|}$ is low, so that the parallel current closing the diamagnetic current generated by the filament at the mid-plane, flows unimpeded along field lines to the bounding solid surface, where the resistance of the resulting filament circuit is determined by the surface sheath resistivity, η_{sheath}. For ions entering the sheath at the ion sound speed and electrons following the Boltzman relation, the boundary condition for the current accepted by the sheath is given by [146]

$$J_{\|} = -nec_s \left[1 - \exp\left(-\frac{e\phi}{T_e}\right)\right]. \tag{2.8}$$

Krasheninnikov worked with a simplified system, describing a filament as homogeneous in the parallel direction and propagating through vacuum, with negligible parallel flows. By assuming a separable solution of the filament (blob) density distribution of the form $n_b(t,x,y) = n_b^{(x)}(t,x)n_b^{(y)}(y)$, where $n_b^{(x)}(t,x)$ gives the time dependent radial profile of the filament and $n_b^{(y)}(y) \propto \exp(-(y/\delta)^2)$ describes a Gaussian filament density profile in the poloidal direction, with poloidal width δ, he arrived at the following expression for the radial velocity of the filament [115]:

$$V_b = c_s \left(\frac{\rho_i}{\delta}\right)^2 \frac{\ell_b}{R}\frac{n_b}{n_t} \propto \delta^{-2}. \tag{2.9}$$

Here ℓ_b is the parallel length of the filament and n_b and n_t are the filament and target densities respectively, which were assumed to only differ by a constant factor. Importantly, this established a qualitative understanding of the dependence of the filament's radial motion on its cross field size.

Krasheninnikov's model was extended in Ref [123] to include the transport of heat and vorticity, as well as further expounding the consequences of the model and their relevance to experimental measurements. Several phenomena such as high SOL collisionality, large X-point magnetic shear and electromagnetic effects can generate parallel gradients in the electric potential and effectively disconnect the filament from the wall so that the sheath is no longer the limiting resistance in the circuit [128]. When the filament is not longer electrically connected to the sheath it enters the *inertial* regime where the sheath closure is replaced by polarisation currents. In this case the filament velocity dependence of filament width changes to $v_b \propto \delta^{1/2}$ [49].

Probably the most general model of filament motion across a range of plasma regimes is the *two-region model* set out in [86]. In the two-region-model the plasma is divided into distinct regions between the mid-plane and the X-point and the X-point and divertor sheath and the coordinate mapping between the two is used to account for the effects of magnetic field fanning and shear near the X-point. The curvature driven current source is located in the first mid-plane region, while the current loop can close at different positions either in the first or second region depending on the properties of the system. The model employs four scale invariant parameters to characterise the behaviour of the filaments.

The first of these is the *collisionality parameter* [128],

$$\Lambda = \frac{L_\parallel/c_s}{1/\nu_{ei}}\frac{\Omega_i}{\Omega_e}, \tag{2.10}$$

where ν_{ei} is the electron-ion collision frequency and $\Omega_{i/e}$ is the ion/electron gyro-frequency. Λ is related to the SOL colissionality parameter discussed in §1.5.1 according to $\Lambda = \nu_*(m_e/m_i)^{1/2}$ and takes on a similar role in describing the size of the plasma bulk resistivity with respect to the effective sheath resistivity.

The second scale invariant parameter is *filament size parameter* $\Theta = \hat{\delta}^{\frac{5}{2}}$, where $\hat{\delta} = \delta_b/\delta_*$ is a dimensionless filament width, normalised to the fundamental characteristic filament size. The *fundamental filament size* is given by [147]

$$\delta_* = 2\rho_s\left(\frac{L_\parallel^2}{\rho_s R}\right)^{\frac{1}{5}}, \tag{2.11}$$

where $\rho_s = c_s/\Omega_i$ is the ion sound gyro-radius. This corresponds to the filament size at which all three terms in the vorticity equation become comparable and filaments travel with the greatest coherency, propagating large distances across the SOL.

For small filaments with sizes, $\rho_s < \delta_b < \delta_*$, the polarisation current term dominates and the path of least resistance uses predominantly cross-field currents.

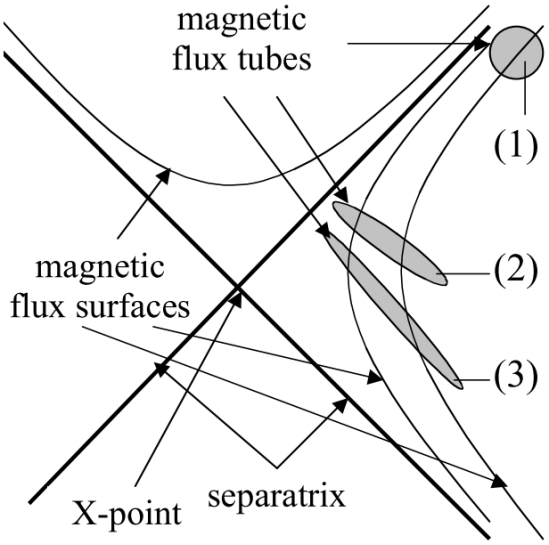

Figure 2.4: Illustration of the shearing of filaments as they pass the X-point. Filaments with circular cross sections at the mid-plane (1) become highly elliptical as they pass the X-point (3). Reproduced from Ref. 148.

For large filaments of width $\delta_b > \delta_*$, the parallel current to the sheaths dominates and there is an increase in filament velocity with collisionality [49]. Large filaments are susceptible to being broken up into fingers or lobes by curvature-driven Rayleigh-Taylor (RT) instabilities [49].

The third parameter, ϵ_x is the ratio of the perpendicular scale lengths in the two regions providing a measure of how much the magnetic field lines in a flux tube fan out in the divertor region. This accounts for how the changing magnetic geometry along the filament's length can enable polarisation currents to flow more easily in certain regions. Near the magnetic X-point the poloidal magnetic field is low and the local magnetic shear is high, so that the filament's cross section becomes highly distorted. At the mid-plane the poloidal and radial filament widths are typically comparable $\delta_y \sim \delta_x \sim \delta_b$, but near the X-point the filament cross section is mapped to a thin elliptical fan, as illustrated in Figure 2.4. A value of $\epsilon_x \approx 1$ corresponds to little shear, while $\epsilon_x \to 0$ for very high shear. The reduced dimensions of the potential dipole allow polarisation currents to readily flow and close the parallel currents in the X-point region. Filaments very close to the X-point can be sufficiently distorted that the electrostatic potential varies on a spatial scale smaller than the ion Larmor radius, in which case the ions' and electrons' average $\mathbf{E} \times \mathbf{B}$ velocities over a Larmor orbit become unequal, providing an additional source of current closure [149].

Finally, the *scaled filament Mach number*, $\hat{v} = v_x/v_*$, is the dimensionless ratio of the filament velocity to the scale invariant characteristic velocity,

$$v_* = c_s \left(\frac{\delta_*}{R}\right)^{\frac{1}{2}}. \tag{2.12}$$

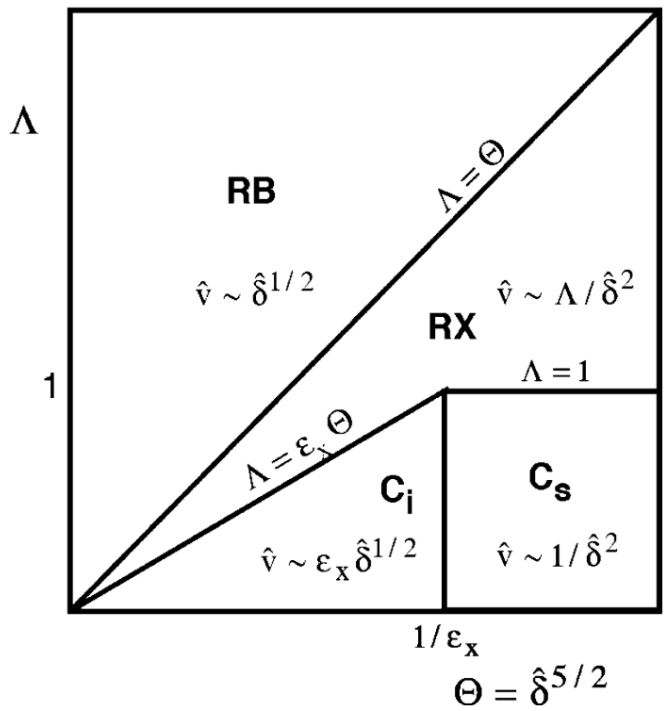

Figure 2.5: Regime diagram for two region model of filament velocity scaling. The scalings of the radial filament velocity parameter with normalised poloidal filament width are shown for each region, along with the boundary conditions for each region. The four regimes and their abbreviations are given in the text. Reproduced from Ref. 86.

The phase space described by these parameters can be divided into four transport regimes, each with a different relationship between filament Mach number, $\hat{v}$ and size $\hat{\delta}$, as shown in Figure 2.5. The four regimes are as follows [86]:

1. *Sheath Connected Interchange Regime* (C_s)
   This regime corresponds to the sheath limited regime from the simpler sheath dissipation model discussed previously. This regime is applicable for large filaments ($\hat{\delta} > \epsilon_x^{-2/5}$) at low collisionalities ($\Lambda < 1$) for which polarisation currents are small and the path of least resistance is for the filament circuit to close at the target, so that the current is limited by the sheath resistivity, leading to filament velocities that scale with filament size as $\hat{v} \propto \hat{\delta}^{-2}$.

2. *Connected Ideal Interchange Regime* (C_i)
   For low collisionality ($\Lambda < \epsilon_x \hat{\delta}^{5/2} < 1$) and small filament sizes, significant coupling to the divertor region persists. However, polarisation currents in the divertor region become the dominant form of current closure. The level of magnetic shear, determines how much the filament is compressed, shortening the cross-field path of the return current and slowing the filament, resulting in velocities scaling as $\hat{v} \propto \epsilon_x \delta^{1/2}$.

3. *Resistive X-point Regime* (RX)
   At intermediate collisionality and sufficiently large filament size, the high parallel resistivity fully disconnects the mid-plane region from the target sheath, such that strong polarisation currents close the current loop up stream from the target. In this regime the parallel current is limited by the parallel resistivity, so that the filament velocity has a linear dependence on the collisionality, scaling as $\hat{v} \propto \Lambda\hat{\delta}^{-2}$.

4. *Resistive Ballooning Regime* (RB)
   When the collisionality is high or the filament widths are small ($\Lambda > \hat{\delta}^{\frac{5}{2}}$), the filaments are sufficiently small for the polarisation currents to decouple the mid-plane from the divertor region, such that the dynamics of the filaments are fully determined in the mid-plane region. Therefore distortion of the filament cross section close to the X-point is no longer important. Perpendicular currents close the current loop in the vicinity of the curvature current drive, causing the filaments to balloon locally at the mid-plane. Filaments in this regime have radial Mach numbers that scale with size as $\hat{v} \propto \hat{\delta}^{1/2}$.

Given that these velocity scalings are derived by approximating gradients as inverse scale lengths and use asymptotic approximations is should be noted that their normalized quantities are likely to have uncertainties greater than order unity [150]. Across these regimes there is a general increase in filament Mach number with increasing collisionality and a reduction with increasing shear at the X-point [151]. Therefore, increases in line-averaged density and reductions in plasma current which increase the SOL collisionality are thought to lead to increases in radial convective transport from filaments [146]. Similarly, experimental observations have shown higher filament velocities closer to separatrix, where the magnetic shear is greatest in proximity to the X-point [148]. The stronger electric fields for a given dipole potential of smaller filaments drives small filaments with the greatest velocities at all but the very highest collisionalities, at which the effects of rapid closure by polarisation currents become dominant [42].

From the two-region velocity scalings we see that the sheath connected (C_s) and resistive ballooning (RB) regimes provide simple analytical lower and upper bounds respectively on the radial filament velocity as a function of filament width,

$$\frac{1}{\hat{\delta}^2} < \hat{v} < \hat{\delta}^{\frac{1}{2}}, \tag{2.13}$$

which can be verified in experiments and numerical models. Figure 2.6 shows an experimental verification of these theoretical bounds using GPI data from NSTX. The points lie broadly within the theoretical limits, with most points closer to the C_s limit and some points below it. The theoretical velocity scaling are often found to be upper limits, since secondary mechanisms can slow the filaments, in particular the internal rotation of the filaments [128]. The rotation is driven by the radial electric field formed in

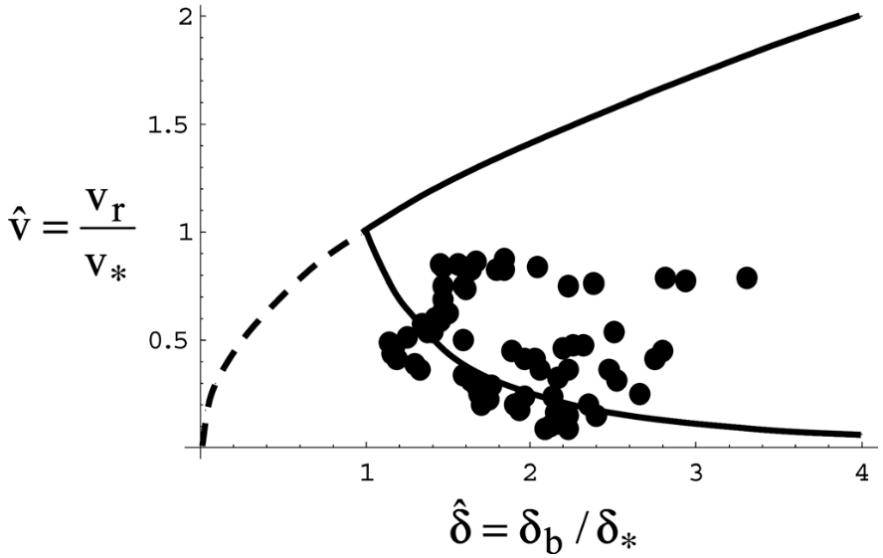

Figure 2.6: Normalised radial filament velocity as a function of normalised filament width for a set of bright, long-lived, isolated filaments observed in GPI data on NSTX. The solid lines show the lower ($\hat{v} \propto \hat{\delta}^{-2}$) and upper ($\hat{v} \propto \hat{\delta}^{1/2}$) theoretical limits on the filament velocity from the two region model, corresponding to the sheath connected (C_s) and resistive ballooning (RB) regimes respectively. The experimental data lines within or slightly below the theoretical bounds. Reproduced from Ref. 139.

the filament due to the greater temperature and thus floating potential with respect to the background plasma. This monopole potential out competes the dipole potential and produces significant total rotation when the temperature perturbation in the filament is much greater than its density perturbation $\frac{\Delta T}{T_{bg}} \gg \frac{\Delta n}{n_{bg}}$. The internal vortices spin with a frequency $\Omega_{spin} \sim (\Delta T/T_{bg})(c_s \rho_s/\delta_b^2)$, where T_{bg} is the temperature of the background plasma and ΔT is the temperature elevation of the filament [42]. These effects are strongest in the near SOL where the filaments have not yet equilibrated with the background plasma. Charge mixing from the internal vorticity introduced by these instabilities reduces the polarisation of the filaments and slows down their radial motion and reduces their distance of propagation [152]. Additionally, the finite background plasma density increases the filament density scale length compared to the models, reducing the charge polarisation drive, and making the filaments slower and more stable [153]. Large variations in internal filament vorticity are observed, which introduce significant scatter in experimental velocity measurements, making detailed confirmation of the theoretical velocity scalings challenging. However, studies in TCV [150] and TORPEX [130] have shown rough agreement with these scalings, despite significant scatter.

Less attention has been paid to the scaling of filament's velocities with their amplitudes, although several approaches have been explored [154–157]. The amplitude dependence for the inertial regime has been identified through multiple approaches [154, 157] to be $v \propto \sqrt{A}$, where A is the maximum fluctuation amplitude normalised to the

background, $A = \max[\delta n/n_{bg}]$. The scaling for the sheath limited case is less straightforward, with Omotani *et. al* [157] showing that the appropriate density scale lengths depend on a combination of the fluctuation and background amplitudes, such that an additional parameter $0 < \beta < 1$ is required, related to the relative importance of the filament amplitude ($\beta = 0$) and the total filament density ($\beta = 1$), which is not known *a priori*. Then the velocity scaling with amplitude goes as:

$$v \propto \frac{A}{1 + \beta A}. \tag{2.14}$$

Two dimensional simulations for large filaments well into the sheath limited regime indicate a value for the amplitude parameter of $\beta \approx 0.31$ [157].

2.3.3 Relationship between filaments and scrape off layer profiles

As was touched on in §1.5.2, SOL profiles show a number of important behaviours that cannot be understood with diffusive models, instead requiring a thorough understanding of filament dynamics. Mid-plane density profiles are often observed to *flatten* a certain distance from the separatrix, with a *shoulder* separating the near SOL and the far SOL where the flattening occurs. Importantly, this flattening leads to stronger plasma surface interaction with first wall. Experiments show the decay length becomes longer as fuelling is increased [43, 66, 158]. This is observed in both the near and far SOL, but is most pronounced in the far SOL, as can be seen in the left panel of Figure 2.7 showing measurements from MAST [158]. For high levels of fuelling the shoulder moves closer to the separatrix; an effect that is sometimes called *broadening*. At particularly high fuelling the far SOL can be almost flat and encompass most of the SOL [43, 44, 46, 158]. The broadening is also accompanied by the extending of the non-Gaussian fluctuation statistics typical of the far SOL to the near SOL region, indicating the role of filaments in this process.

Higher plasma currents are seen to produce steeper gradients in the SOL and can prevent broadening [158], as demonstrated in the central panel of Figure 2.7. This is likely due to a reduction in the connection length, as profiles are similarly seen to steepen in regions of the SOL with reduced connection length, where the field lines are in the 'shadow' of the first wall and impinge on limiters before they reach the divertor [44, 45].

Electron temperature profiles also show some level of flattening, but do not appear to be affected by fuelling levels [43, 46, 159]. The fact that the two parameters respond differently indicates that particle and heat exhaust may be regulated by different mechanisms. Therefore the electron temperature is unlikely to be an important parameter in determining the density profiles. The decoupling of the two profile behaviours also means it is unlikely that density broadening is a result of increased ionisation in the SOL as this would be reflected in a drop in electron temperature due to the ionisation energy imparted by the electrons.

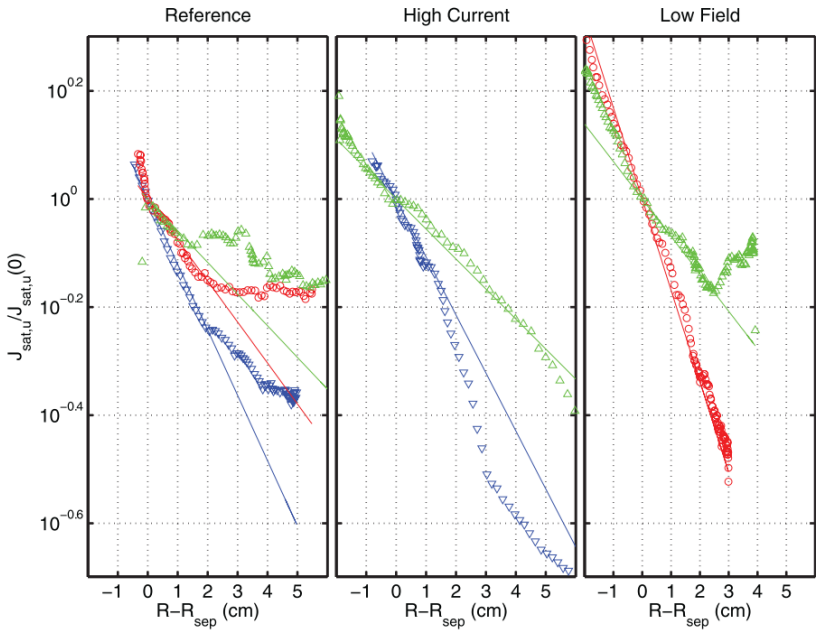

Figure 2.7: Radial profiles of ion saturation current from the MAST tokamak, normalised to their separatrix values, demonstrating *flattening* and *broadening* effects as a result of changes in fuelling and plasma current (c.f. simplified illustration in Figure 1.13). Blue ▽, red ◯ and green △ symbols represent low, intermediate and high densities respectively, while straight lines show exponential fits to the near SOL portions of the data. Reproduced from Ref. 158.

Measurements of Langmuir probe statistics across machines [66, 69, 160] have shown that the PDFs of ion saturation current under a range of conditions collapse onto a common curve when renormalised by subtraction of the mean and rescaling by the standard deviation of the fluctuations. The PDFs are fitted well by a Gamma distribution under all conditions as observed in other machines [66] (see Figure 1.16). The right-tailed, positive skewness of the Gamma distribution is indicative of the greater prevalence of positive density fluctuations (filaments) compared to negative density fluctuations (holes) in the SOL. The consistency of the shaping parameter of the fitted Gamma distribution across a range of conditions indicates that the intermittency of the fluctuations are approximately constant across a range of line averaged densities and plasma currents [160].

Measurements on ASDEX-U, accompanied by EMC3-EIRENE simulations [161], indicated that the formation of the flattened L-mode density shoulder at the outer midplane correlates with a transition between the sheath-limited and the inertial filamentary regimes. SOL divertor collisionality, Λ_{div}, was suggested as the parameter triggering the transition.

However, measurements on TCV have found changes in divertor resistivity and connection length to have little effect on filament sizes and velocities as the majority (80-90%) of filaments in TCV are found to populate the resistive ballooning (RB) regime, for

which the filament currents are completely closed before they reach the divertor [150]. However, this result may not hold in other machines with more closed divertors, where the collisionality in the divertor is likely to dominate over that in the main chamber.

A more complete picture has been presented by studies in Alcator C-mod [162] and JET [163] which investigated changes in Λ_{div} introduced through different channels. Increases in flattening and intermittency with upstream fuelling were found to be dependent on divertor strike point orientation and strike point sweeping. The absence of flattening in JET vertical target configurations [163] shows that upstream fuelling is insufficient to cause flattening and localises the mechanism driving the process to the divertor region. Increases in the Λ_{div} from nitrogen seeding were also found to not be sufficient to bring about flattening. Furthermore, it is possible the changes in Λ_{div} may be a consequence of flattening, rather than a controller of it. These conclusions are consistent with findings in ASDEX-U [161] and TCV [164]. This indicates a more complicated parametrisation of shoulder formation, where the role of neutral interactions in the divertor region needs further investigation, as the level of flattening is found to scale most consistently with the level of recycling in the divertor [163].

There is a great need to elucidate how changes in machine parameters affect filament properties which in turn modify profiles. Through the coupling of filaments and profiles in analytic frameworks progress can be made in pinning down what mechanisms are responsible.

2.3.4 Analytic frameworks

In order to fully understand the behaviour of SOL profiles it is necessary to have a model for how the filaments, with their statistically distributed dynamical properties, combine to generate time averaged SOL profile shapes. Taking filaments as the 'quanta' of turbulent transport, we can produce an alternative to the model diffusive-advective model which has been shown to be incorrect.

Garcia's 2012 paper [131] introduced a single point stochastic model of plasma fluctuations based upon the independent nature of the filaments. Independent events, which for stationary conditions are uniformly distributed in time, are governed by Poisson statistics on which the model is built. A common example of Poisson statistics is in describing the number of decay events per second from a radioactive source. Equivalently, the number of filaments, K, occurring in the time interval ΔT is given by the Poisson distribution,

$$P_K = \lambda^K e^{-\lambda}/K!, \tag{2.15}$$

where the expected number of filaments is given by $\lambda = \Delta T/\tau_w$, where τ_w is the average *waiting time* between filaments. It follows that the individual waiting times between

the filaments, τ, are exponentially distributed, as given by:

$$P_\tau(\tau) = \frac{1}{\tau_w} \exp\left(-\frac{\tau}{\tau_w}\right). \tag{2.16}$$

The amplitude of the density at a point in the SOL, $\theta(t)$, (e.g. as measured by a Langmuir probe) is modelled as the sum of the instantaneous bursts waveforms for each filament passing the point. The properties of the burst waveforms are described by their characteristic temporal duration, τ_d, and the average separation of the bursts τ_w. The level of variation in the signal can then be described by the intermittency parameter, $\gamma = \tau_d/\tau_w$, with bursts being well separated in time for $\gamma \to 0$ and strongly overlapping for $\gamma \gtrsim 1$. For large values of γ the PDF of the amplitude signal, $P_\theta(\theta)$ approaches a normal distribution, while for small values of γ, P_θ deviates strongly from a Gaussian, having positive skewness and kurtosis as is often observed for measurements in the far SOL. Garcia showed that as consequence of the Poisson properties of the system the flatness of the amplitude distribution has a parabolic relation to the skewness: $F \propto S^2$.

Further progress can be made by assuming the filament amplitudes η_0, with mean amplitude $\langle\eta_0\rangle$, follow an exponential amplitude distribution $P_{\eta_0}(\eta_0) = \langle\eta_0\rangle^{-1} \exp(-\eta_0/\langle\eta_0\rangle)$ as observed experimentally with multiple diagnostics [66, 90, 165]. In this case the variance, skewness and flatness are given by $\sigma^2 = \gamma^{-1/2}$, $S = 2\gamma^{-1/2}$ and $F = 3 + 6/\gamma$ respectively. The direct relationship between flatness and skewness is then given by,

$$F = 3 + \frac{3}{2}S^2; \tag{2.17}$$

a relationship found in experiments [66, 90, 162]. Furthermore, it was shown that the fluctuation amplitude distribution is described by a Gamma distribution,

$$P_\theta(\theta) = \frac{1}{\langle\eta_0\rangle\Gamma(\gamma)} \left(\frac{\theta}{\langle\eta_0\rangle}\right)^{\gamma-1} \exp\left(-\frac{\theta}{\langle\eta_0\rangle}\right), \tag{2.18}$$

with scale parameter given by the average filament amplitude $\langle\eta_0\rangle$ and shape parameter equal to the intermittency parameter, γ. The Gamma function is defined as: $\Gamma(\gamma) = \int_0^\infty \phi^{\gamma-1} \exp(-\phi) d\phi$.

Militello and Omotani [45, 167, 168] extended the single point model to include the radial dimension, x, and the dynamics of filaments and thus relate the dynamic statistical properties of filaments directly to radial profiles of thermodynamic variables. Their framework similarly describes the superposition of a train of individual filaments, each identified now by the subscript i and each with their own properties. The coordinate for radial position is defined such that x is zero at the separatrix and positive in the SOL. The filaments are assumed to be emitted from the separatrix at a constant rate, with average waiting time τ_w as before. The filaments are also assumed to have a well defined radial shapes, $\Lambda(x, w_i)$, depending on a single width parameter w_i for each filament. The initial amplitude of each filament is given by $\eta_{0,i}$, with the function

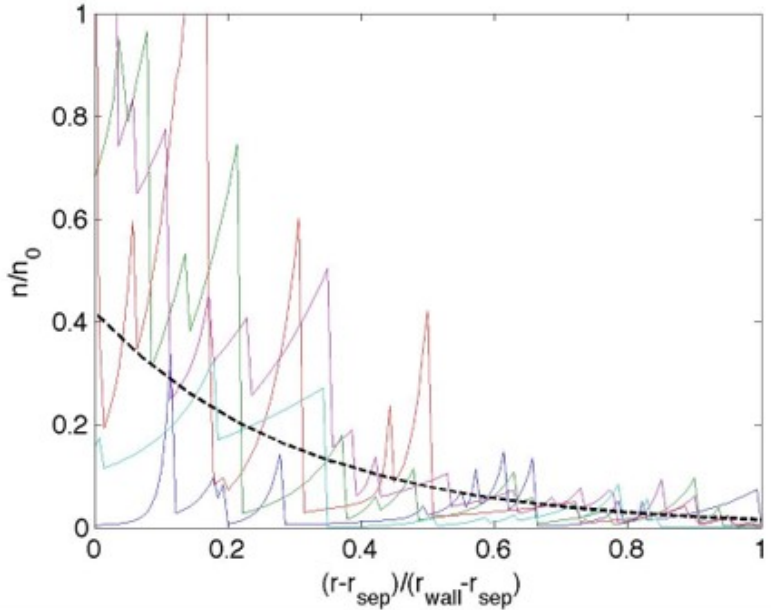

Figure 2.8: Illustration of the formation of a SOL density profile from the time averaging of many intermittent filamentary fluctuations. Different coloured solid lines show instantaneous profiles containing filamentary bursts, while the dashed black line shows the time averaged profile for long timescales. Reproduced from Ref. 166.

$F_i(t)$ describing how the amplitude falls with time due to generic parallel losses, over a characteristic time $\tau_{\|}$. The temporal variation of each filament's velocity is given by a function $V_i(t)$ that is assumed to be fully prescribed by the initial amplitude and size of the filament such that, $V_i(t) = V(t, \eta_{0,i}, w_i)$. The filament width and amplitude are further assumed to be independent.

Combining these elements the radial motion and parallel draining of each filament crossing the separatrix, $x = 0$, at time $t = 0$, can be represented as

$$\eta_i(x,t) = \eta_{0,i} F_i(t) \Lambda \left(x - \int_0^t V_i(t')dt', w_i \right), \tag{2.19}$$

where the filament is viewed in the Eulerian frame. The steady stream of filaments emitted from the separatrix at a constant rate, each at a time $t = t_i$, are assumed to not interact, with their amplitudes combining additively to produce a profile:

$$\theta(x,t) = \sum_{i=1}^{\infty} \eta_i(x, t - t_i). \tag{2.20}$$

Since radial plasma profiles are typically measured on times scales, ΔT, longer than the filament lifetimes, we seek to calculate the time-averaged profile $\Theta(x) = \overline{\theta(x,t)} = \lim_{\Delta T \to \infty} \Delta T^{-1} \int_0^{\Delta T} \theta(x,t)dt$. Figure 2.8 illustrates how the time average of many instantaneous bursty $\Theta(x)$ profiles produces the familiar exponential-like profiles from experiments. We wish to be able to calculate $\Theta(x)$ from knowledge of the probability

distribution functions of the filament parameters rather than specific knowledge of the exact parameters for each specific filament. The PDFs of the initial amplitudes and widths are denoted by P_{η_0} and P_w respectively. Given the ergodic nature of the system described above, the desired temporal average is equal to the ensemble average across these PDFs, $\Theta_{\Delta T}(x) = \langle\theta(x,t)\rangle$. Utilising the Poisson properties of the system it can be shown that in the limit $\Delta T \to \infty$, the radial profile is equal to

$$\begin{aligned}\Theta(x) &= \{\eta(x,t)\} \\ &= \frac{1}{\tau_w}\int_{-\infty}^{\infty} dt \int_0^{\infty} d\eta_0 \int_0^{\infty} dw[\eta(x,t)P(\eta_0)P_w(w)],\end{aligned} \tag{2.21}$$

where $\{...\}$ denotes the average operator and the subscripts have been dropped. Equation 2.21 provides us with the desired mapping from the statistical properties of the filaments to SOL profiles, with the filament dynamics describing the behaviour of both the near and far SOL in a consistent way. Therefore if the statistical properties of the filaments can be constrained through a combination of experimental data or theoretical models, equation 2.21 provides a means of making first principles predictions of the response of the SOL profiles to changes in plasma conditions.

The role of the different filament parameters can be elucidated further by considering the output of the framework in the simple case of filaments with fixed widths, w_*, and amplitudes, η_*, travelling with constant radial velocities and draining with a fixed parallel drainage time, $\tau_\parallel$, in a model where $F(t) = e^{-t/\tau_\parallel}$. In this case the resulting radial profile is given by

$$\Theta(x) = \frac{\eta_*}{\frac{\tau_w}{\tau_\parallel}\left(1 + \frac{V\tau_\parallel}{w_*}\right)} \exp\left(-\frac{x}{V\tau_\parallel}\right) \tag{2.22}$$

and relative amplitude of the fluctuations as measured by the ratio of the standard deviation to the mean is given by

$$\frac{\sigma(x)}{\Theta(x)} = \frac{\sqrt{2}}{2}\sqrt{\frac{\tau_w}{\tau_\parallel}\left(1 + \frac{V\tau_\parallel}{w_*}\right)}. \tag{2.23}$$

Equation 2.22 shows that the profile decay length depends on the combination $V\tau_\parallel$, so that in order for the flattening phenomena to occur, there must be radial variation in the velocity and/or parallel drainage timescales of the filaments. A radially increasing radial velocity could be the result of the reduction in the density of the background through which the filaments moves [157], or due to a radial increase in resistivity [169]. Radial variation of the filament speed also has the ability to explain the increase in relative fluctuation amplitude in the far SOL as can be seen from Equation 2.23. However, experimental evidence is often inconsistent as to whether filaments accelerate [84, 87], decelerate [86], or travel at constant velocity [51, 167].

A decrease with radius of $\tau_\parallel$ can also produce profile flattening, although it needs to be accompanied by an equal or greater increase in the waiting time τ_w in order to

reproduce the radial increase in relative fluctuation amplitude, $\frac{\sigma}{\Theta}$ according to Equation 2.23. This would require some filaments to fully dissipate before reaching the far SOL. In reality, flattening is most likely a consequence of a combination of increasing filament velocity and reduced parallel draining.

For broadening to occur, strong acceleration of the filaments is required. However an accompanying increase in filament amplitude is needed in order to not lower the separatrix density. Alternatively, $\tau_{\|}$ could increase with density as a result of increased charge exchange interactions between neutral particles and ions, slowing the filaments' drainage to the target. These interactions do not cool the electrons and so are consistent with the observed decoupled temperature profiles.

The framework was applied to the interpretation of profile flattening experiments in JET with great success and was able to simultaneously match density profiles, fluctuation PDF shapes and fluctuation auto-correlation functions using physically reasonable filament properties [170], as indicated in Figure 2.9. By eliminating possible variations in different free parameters in the model it was discovered that the onset of density shoulder formation could only be consistently produced by a strong local radial variation in the $\tau_{\|}$, corresponding to an increase by a factor of around 25 in the vicinity of the density shoulder. However, radial variation in $\tau_{\|}$ alone was insufficient to match subsequent evolution and widening of the density shoulder at the highest density. In this case a combination of a radially localised increase $\tau_{\|}$ and a radial variation in the filament velocity was required to match the profiles consistently. This suggests that while the onset and evolution of the profiles' density shoulders are related, they should perhaps be viewed as separate events, with the evolution at high densities evoking additional physical processes.

In order to advance our understanding of the origin of these profiles' behaviours, further experimental constraints on the statistical properties of the filaments are required, together with experimental verification of the assumptions underpinning the framework. These tasks play to the strengths of direct visible imaging measurements, such that inversion of visible camera data can play a crucial role in validating the framework and refining its inputs. To this end, Chapter 6.3 contains experimental results with these applications.

2.4 Numerical modelling of filaments

Due to the limitations of experimental measurements of turbulence and resulting filaments and the desire to gain first principle insight into the physics of these phenomena, there has been a large effort to develop a more sophisticated understanding of the important effects governing turbulent particle transport, that cannot be captured in simple analytic models. To this end a range of different 2D and 3D turbulence codes have been developed to explore the validity of different physical assumptions and the effects of

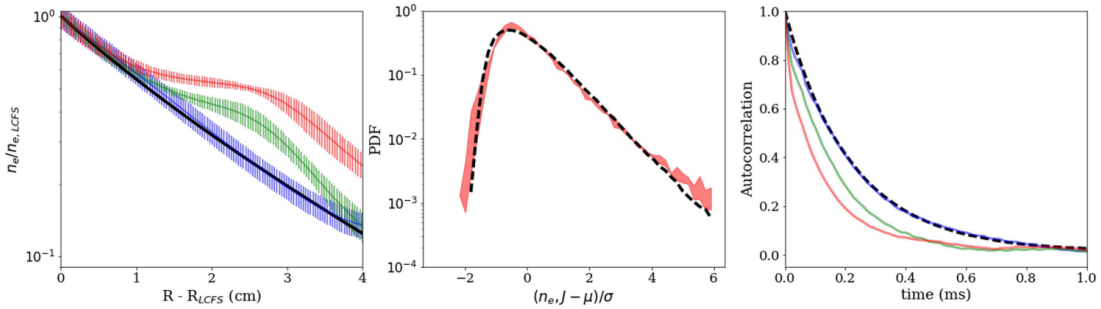

Figure 2.9: Application of the analytic filament framework to data from JET. Blue, green and red lines indicate experimental data taken at low, intermediate and high plasma densities respectively. The fits produced by running a tuned stochastic simulation utilising the analytic filament framework is shown in black. The left panel shows separatrix normalised density profiles measured with the lithium beam emission (LiBES) diagnostic. The central and rightmost panels show the rescaled rescaled fluctuation PDF and autocorrelation functions for Langmuir probe measurements made at the JET outer wall limiter. The model is able to simultaneously match all three measurements. Reproduced from Ref. 170.

realistic magnetic geometries. Computational filament studies typically consider either fully saturated turbulence from which filamentation occurs, or seeded isolated filaments propagating in a uniform background plasma, with each offering deeper insight on different aspects of filament behaviour.

It is common for models to employ a simplifying cold ion approximation. While such an approximation is known not to be valid, as experiments typically show $T_i \gtrsim T_e$ in the SOL [171–174], due to the higher parallel conductivity of electrons compared with that of the less mobile ions, the results of these models often still compare favourably with experiment [99]. A further common approximation is the *Boussinesq approximation* which is used to simplify the calculation of the ion polarisation currents, by taking the perpendicular density gradients to be small with respect to the gradients in the potential. Such an assumption is only valid for small density perturbations, which is not true for filaments, but nonetheless yields surprisingly consistent results when compared with experiment [122].

2.4.1 2D turbulence codes

A wide range of 2D turbulence codes including ESEL (Edge-SOL ELectrostatic) [175], HESEL (Hot Edge SOL Electrostatic) [176], SOLT (Scrape-Off-Layer Turbulence) [133], TOKAM2D [177] and STORM2D (Scrape-off layer TuRbulence Models) [122] have been applied to the study of filaments. Two dimensional numerical modelling work has often been applied in simplified slab geometries, with curvature effects supplied by effective gravity terms. Saturated turbulence is generated in a driving region, from which the perturbations spread into an open field line region representing the SOL, where their

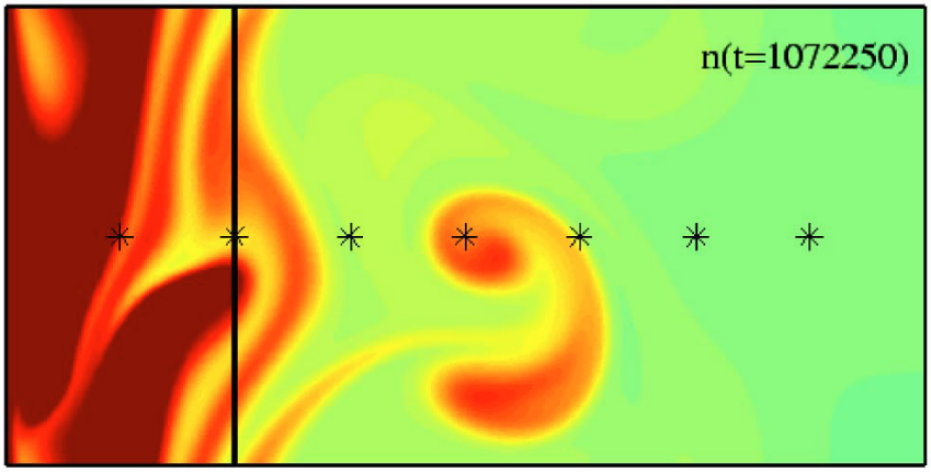

Figure 2.10: 2D simulation of turbulent density fluctuations leading to the formation of a filament with mushrooming characteristic of Kelvin-Helmholtz instabilities. Stars show points at which fluctuation statistics were compared. Reproduced from Ref. 178.

evolution is studied. An example of such a turbulent eddy detaching from the separatrix and propagating radially through the SOL is given in Figure 2.10. Simulations have been able to reproduce the characteristic PDFs of fluctuations and radial variations of statistical properties observed in experiments [175, 178]. 2D simulations have provided a great deal of insight into the stability of filament structures to Rayleigh-Taylor and Kelvin-Helmholtz instabilities. The mushrooming frequently seen in modelling work has been verified experimentally in GPI data and in toroidal plasma devices using Langmuir probe arrays [179].

2.4.2 3D turbulence codes

While 3D modelling of turbulence is significantly more computationally intensive than 2D modelling, advances in computing power over the last decade or so have made 3D turbulence simulations on supercomputers routine. 3D codes are able to take into account a range of parallel effects such as variation in the density between the mid-plane and the target due to sheath losses. Three dimensional turbulence codes used in filament studies include GBS (Global Braginskii Solver) [180], TOKAM3X [181], STORM3D [169], HERMES [182] and GRILLEX [183]. An example of a 3D turbulence simulation in GBS exhibiting parallel structure is shown in Figure 2.11. Full 3D codes are able to investigate realistic magnetic geometries and explore the effects of X-points and divertor configurations on filament dynamics. For instance recent 3D TOKAM3X simulations of plasmas in WEST [184] explored these effects and found agreement with the sheath connected (C_s) and resistive X-point (RX) regime scalings in the main SOL, while filaments close to the separatrix were found to follow the connected ideal-interchange (C_i) regime scaling.

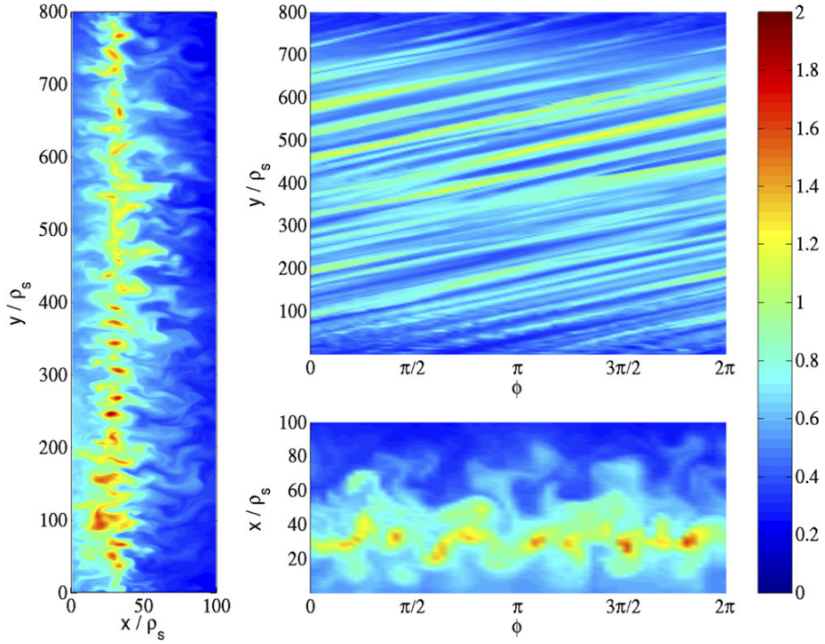

Figure 2.11: Normalised electron pressure in the SOL from 3D turbulence simulation in GBS. The left, upper right and lower right panels show the poloidal plane, a radial cut and a toroidal cut through the simulation domain respectively. Reproduced from Ref. 180.

2.4.3 Isolated filament modelling

Full simulations of saturated turbulence are particularly computationally expensive and do not offer any direct control over the size or amplitude of filaments arising from the turbulence. Therefore seeded isolated filament studies are valuable for their ability to control the size and amplitude of filaments in order to study filament propagation mechanisms in detail. Such analyses are most applicable to the far SOL where interactions between multiple filaments are likely to be less significant. Filaments are generally initialised as 2D Gaussian density perturbations in the perpendicular plane, extending from the mid-plane much of the way to the target. These are set on top of a uniform background and allowed to evolve subject to the physics included in the system.

3D simulations of filaments in the STORM module of BOUT++ have explored the effects of enhanced parallel resistivity close to the target [169]. Large filaments were observed to show large increases in radial velocities with increasing parallel resistivity as they transitioned from being sheath limited to closing via perpendicular polarisation currents in the inertial regime, with a resulting increase in resistance in the current loop. By contrast, small filaments were already in the inertial regime and so experienced only modest increases in radial velocity at higher resistivities. By contrast, especially large filaments retained their dominant connection to the sheath and presented larger increases in speed. This was attributed to entering the resistive sheath current regime,

in which the resistance of the plasma was large enough to generate a potential difference between mid-plane and the sheath, allowing larger driving potentials to form at the mid-plane for the same sheath current. Uniformly enhancing the resistivity along field lines also demonstrated that the variation in cumulative resistance along the field lines can lead to decoupling of the filament dynamics in the parallel direction, with the region of the filament at the mid-plane moving faster than the region downstream.

A mixture of 2D and 3D turbulence simulations in STORM have been used to study the interactions between pairs of seeded filaments. The dipolar electrostatic fields of the filaments were found to be too weak to interact significantly unless the distance between the filaments was less than around 5 filament widths. Filaments at closer separations for which interactions occurred were found to have their radial centre of mass velocities modified by less than $\sim 30\%$, so that their contribution to radial particle transport was minimally affected. Filament separations in experiments are generally observed to be equivalent to 4-10 filament widths, so it was argued that filament interactions are not anticipated to be important in determining filament dynamics. This was used as evidence in support the assumption of filament independence on which the theoretic filament framework discuses in section 2.3.4 is built.

2.4.4 Validation of modelling

A wide range of comparisons with experiments have been performed in order to individually validate turbulence codes with measurements from different machines [67,96]. However, cross comparisons of multiple codes with experiments are more unusual. Riva *et al.* performed a multi-code validation exercise of four different turbulence codes, comparing them with Langmuir probe array measurements of filaments in the TORPEX toroidal plasma device [185]. The STORM, GBS, HESEL and TOKAM3X codes were compared, which are each based on different combinations and implementations of assumptions used to simplify the drift-reduced Braginskii equations, such as cold ions, isothermal electrons and negligible electron inertia. By comparing each of the codes consistently it was possible to establish which details are most important in capturing experimentally observed filament behaviour. Neglecting electron inertia, applying the Boussinesq approximation for small density fluctuations and working in the infinite aspect ratio limit were each found to have minor impact determining filament velocities. However, due to the importance of parallel currents in TORPEX operating regimes, the vorticity advection closure implemented in HESEL was unable to correctly reproduce the filament dynamics, while the sheath dissipation closure implemented in BOUT++ 2D, gave good agreement with experimental measurements and 3D simulations. Furthermore, the value of the background electron temperature was found to be important in correctly determining the radial velocity of filaments, while the internal electron temperature was only necessary to describe the filaments' rotation.

The multi-code analysis was extended to the more complex tokamak SOL of MAST by Militello *et al.* [99], using the same four codes. In this case, the experimental data used to benchmark the performance of the codes was produced using the fast camera inversion technique presented in this thesis (see §6.2). Two large, bright, isolated filaments were studied, which each had similar perpendicular sizes and amplitudes, but which exhibited different radial motions. The evolution of the filaments was found to be relatively insensitive to the reinitialisation condition of the seeded filaments and their geometries. The absence of electron inertia terms in TOKAM3X was found to result in the immediate formation of a electrostatic monopole which retarded the early motion of the filaments with respect to the other codes and experiment. The exercise identified the biggest element of uncertainty in the modelling was the role of the ion temperature, due to the lack of direct experimental ion temperature measurements. Inclusion of hot ions in HESEL and GBS led to much larger radial and toroidal velocities than seen in experiment, suggesting that ion temperatures are lower than anticipated ($T_i \lessapprox 2T_e$), or the codes are missing physics that should slow the filaments with finite ion temperatures.

2.5 Conclusions

Coherent filamentary perturbations in density and temperature, with large dimensions parallel to the magnetic field compared to their perpendicular widths, are ubiquitous in the boundary regions of magnetically confined plasmas. In tokamaks, curvature-drifts lead to polarisation of the filaments which drives radial motion, carrying particles and heat out across the SOL. The motion of the filaments thus determines the density width of the SOL and can produce enhanced iterations with the walls of the machine. Extensive analytical and numerical modelling of filaments has elucidated many of the mechanics governing the dynamics of filaments. Through the application of analytic frameworks, understanding has been gained of how the statistical properties of filaments translate into SOL profiles and thus the level of plasma surface interaction with the main walls of the machine. In order to gain the predictive capabilities needed to inform the design of future fusion reactors, greater knowledge is required of the statistical properties of the filaments and their dependence on plasma conditions. This thesis describes a new diagnostic technique that is capable of making direct measurements of these important filament properties, with the potential to produce large databases of filament statistics in order to yield significant progress in this direction.

Chapter 3

Experimental setup

3.1 The Mega Amp Spherical Tokamak

The data used in this thesis have been collected from the Mega Amp Spherical Tokamak (MAST). A spherical tokamak has a very low aspect ratio, $A = R/a$, such that the minor radius is comparable to the major radius and the size of the centre column is minimised. This results in a near spherical plasma shape similar to that of a cored apple, rather than the conventional doughnut like shape of conventional higher aspect ratio machines. The spherical tokamak design was pioneered at the Culham Centre for Fusion Energy (CCFE) with the operation of the Small Tight Aspect Ratio Tokamak (START) in the 1990s [186]. A small aspect ratio has the advantage of being able to achieve a higher plasma pressure for a given magnetic field strength, as characterised by the toroidal plasma beta:

$$\beta_T = \frac{2\mu_0}{VB_T^2}\int pdV, \tag{3.1}$$

where B_T is the toroidal magnetic field, V is the plasma volume and p is the plasma pressure. As the magnetic coils are the most expensive part of a fusion reactor, tight aspect ratio, high B_T machines offer a potential path to smaller, more economical fusion power plants. Following on from the successful proof of principal in START, the larger MAST tokamak operated between 1999 and 2013. Given MAST ceased operations prior to the commencement of this work and delays in the completion of MAST-U pushed its operation to after the completion of the PhD, all experimental data analysis in this thesis was performed on archived data available from the last MAST campaign.

Table 3.1 compares the operating parameters of MAST to the current and future large conventional aspect ration tokamaks JET and ITER. MAST was designed with a particularly open vacuum vessel design, with the poloidal magnetic field coils suspending from the walls of the vacuum vessel as illustrated in Figure 3.1, so that the plasma edge and the vessel wall were typically separated by around 60 cm [96]. This allowed for flexible operation of the magnetic configuration and very good diagnostic access. The machine was up-down symmetric, with identical open carbon divertors at the top

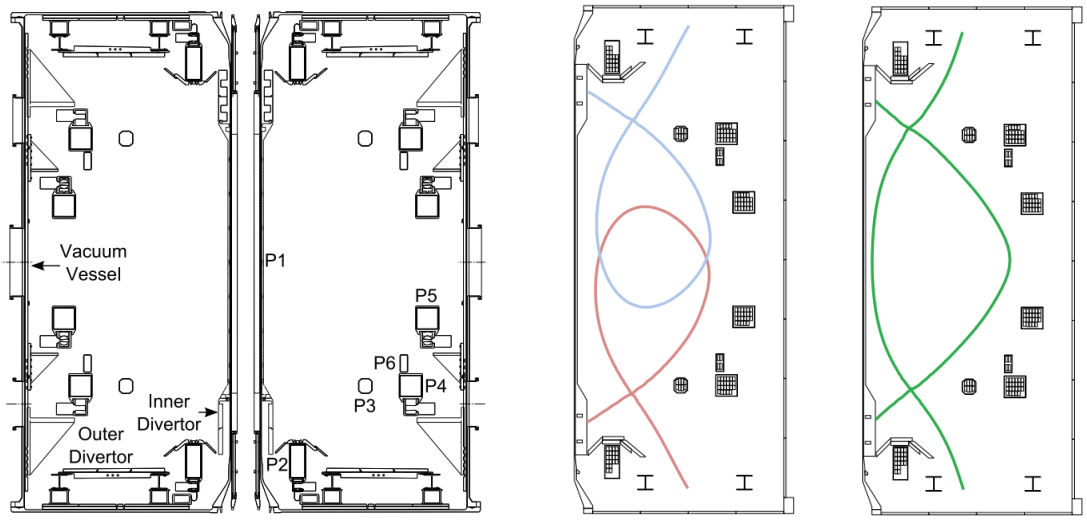

Figure 3.1: a) Engineering schematic of the MAST vessel, courtesy of the MAST design office. PX denotes the numbering of the poloidal field coils. b) Illustration of the separatrices for lower (red) and upper (blue) single null (SN) and connected double null (CDN, green) divertor configurations. Reproduced from Ref. 188.

and bottom. Combined with the flexibility of the magnetic coil design this facilitated operation with one or two X-points as illustrated in the single null and double null configurations in Figure 3.1b. The tight aspect ratio of MAST meant that the pitch of the field lines at the mid-plane typically ranged between 30° and 45° [51], compared to pitch angles of ∼10° in conventional aspect ratio devices [187]. On average, this resulted in the parallel distance from the mid-plane to the divertor targets, $L_{\parallel}$, being ∼14 m, with the X-point located approximately half way along this length.

MAST was particularly well suited to fast visible imaging of filaments due to several aspects of its open, compact design. Firstly, MAST's open 'tin can' vacuum vessel geometry meant that recycled neutrals could easily travel around freely resulting in a high, uniform main chamber neutral pressure. This provided increased plasma-neutral interaction at the plasma edge that facilitated constant imaging without active gas puffing. Under the assumption that the neutral density was roughly constant, the emission along a field line's length was not affected by variations in neutral density. This means that any variation in intensity along the parallel length of a field line should correspond to changing filament properties, so that the parallel structure of the filaments was captured in the camera images. The neutral density was typically around $\sim 10^{18}$ m^{-3} throughout much of the SOL, only varying significantly close to the separatrix where it dropped to $\sim 10^{16}$ m^{-3}, falling to negligible levels around 5 cm inside the separatrix [189]. Therefore, the drop in neutral density led to a rapid cut off in emission just inside the separatrix.

MAST's tight aspect ratio meant a very wide field of view over the plasma could

Table 3.1: Comparison of operating parameters for the MAST, JET and ITER tokamaks.

Quantity	MAST [190]	JET [191]	ITER [192, 193]
Major radius, R [m]	0.9	2.96	6.2
Minor radius, a [m]	0.5	1.0	2.0
Aspect ratio, A	1.3	3.1	3.1
Plasma current, I_p [MA]	≤ 0.8	≤ 3.7	≤ 15
Toroidal field, B_T ($R = 0.7$ m) [T]	≤ 0.63	≤ 3.4	≤ 5.3
NBI heating power, P_{NBI} [MW]	2.9	35	33
Pulse length, τ_{pulse} [s]	$\sim$0.5	≤ 25	≤ 3600
Plasma volume, V [m^3]	8	80	837

be achieved, so that the parallel dimension of filaments could be observed, giving the maximum information about 3D structure. The lack of a conventional first wall closely conforming to the edge of the plasma led to large gaps between the plasma boundary and the vessel surfaces ($\sim$50 cm). Therefore filaments could be observed propagating far out into the SOL without quickly interacting with or being concealed by first wall components. The fast camera system was capable of viewing the whole of the MAST plasma in one view, with only a small portion of the plasma obscured by the centre column. However, in order to achieve sufficient frame rates to resolve the motion of the filaments, analysis typically concerned views of half of the MAST plasma.

Finally, as will become apparent in the next chapter, the tight aspect ratio resulted in increased pitch of the magnetic field, which aids in resolving the radial widths of the filaments using information about the magnetic equilibrium geometry.

3.2 Camera specification

The analysis performed in this thesis relies on the use of fast camera imaging to measure the properties of the filaments. The Photron SA1.1[†] fast framing visible camera used in this work was installed during MAST's M9 campaign (May-September 2013), on loan via an EPSRC loan pool from the University of Swansea. This camera had significantly greater capabilities than CCFE's own Photron APX RS fast cameras which had previously been used.

The camera was mounted to port HM10, viewing the plasma tangentially from the mid-plane and was operated with simple collection optics, consisting of a 5.733 mm focal length lens. This choice of lens framed the plasma well, provided a full unobscured view of one half of the MAST plasma, while achieving the highest concomitant spatial resolution at the mid-plane. The camera was set back slightly from the vacuum vessel to avoid interference from the tokamaks magnetic field. The recording of the camera

[†]Full Photron SA1 specification: https://photron.com/fastcam-sa1-1 (accessed 23/05/2019)

was triggered via a remote control PC located in the MAST control room and relayed to the camera via optical fibre. The data from the camera was relayed to the control PC via return optical fibre, where the data was stored for future analysis.

The Photron SA1.1 used a monochrome CMOS sensor with a 20 μm pixel size and a full sensor resolution of 1024×1024 pixels. A global electronic shutter meant the sensor could scan the entire area of the image simultaneously, avoiding scanning artefacts that occur when using many CMOS sensors which have rolling shutters due to movement of the subject over the course of the image's acquisition. It recorded 12-bit intensities outputting digit levels in the range 0-4096. By operating only a subset of the sensor's pixels, the camera could image particular regions of interest (ROIs) at higher resolutions than the 5,400 fps at which the full sensor could be operated. The camera was unfiltered such that the light emission recorded was dominated by D_α emission at a wavelength of 656 nm. With the exception of the centre column which has a low surface area and poloidal field coils P4 and P5 (see Figure 3.1a), there are no first wall surfaces close the mid-plane plasma in MAST. Therefore significant releases of carbon from the walls in the main chamber are unlikely and the assumption of dominant D_α emission should be valid. The quantum efficiency of the sensor peaks around 640 nm, making the camera's sensitivity to D_α particularly high. A summary of the cameras operating parameters is given in Table 3.2. Overall the Photron SA1.1 camera provided the highest resolution and frame rate imaging capabilities of the full MAST plasma of any affordable camera at that time. As shall become apparent in Chapter 4, the ability to image the full height of the MAST plasma was vital properly constrain the inversion of the plasma emission, while the high spatial resolution and frame rate facilitated measurements of the filament's sizes and the tracking of filaments across multiple frames.

3.2.1 Field of view

The camera was installed in both 'main chamber' and 'divertor' views as shown in Fig. 3.2a, although this work focusses on the 'main chamber' viewing geometry shown in pink. In order to achieve the $\sim$10 μs time resolution needed to resolve filament motion

Parameter	Value
Focal length, f	5.733 mm
Dynamic range, d	$2^{12} = 4096$
Resolution, $(n_x \times n_y)$	256×160 pixels @ 100,000 fps
Pixel size, η	20 μm $\times$ 20 μm
Sensor size, $\eta(n_x \times n_y)$	3.20 mm $\times$ 5.12 mm
Frame rate, ω	100 kHz
Exposure time, ϵ	3 μs
Shutter type	Electronic global shutter

Table 3.2: Photron SA-1.1 camera specification .

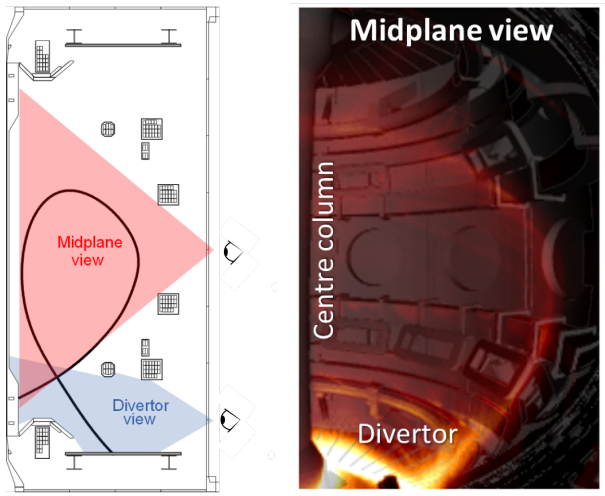

(a) (b)

Figure 3.2: a) Schematic of the field of view of the Photron SA-1 fast camera when installed at the mid-plane or divertor. b) False colour image of a lower single null MAST plasma, as viewed by the SA1 camera, with a CAD rendering of the MAST vessel components overlaid. Figures modified from [68].

across several frames, the camera was operated at a frame rate of 100 kHz, requiring the sensor to be sub-windowed to a 256×160 pixel ROI, viewing one half of the MAST plasma. Figure 3.2b shows a false colour image of a lower single null MAST plasma, superimposed on a CAD rendering of the MAST vessel components. The edge of the centre-column is visible on the left, while the P4 and P5 poloidal magnetic field coils (see Figure 3.1a) and oval shaped ELM control coils are visible on the chamber wall. A top down view of the camera's location in the machine is given in Figure 3.3. The camera was located 2.15 m from the centre of the machine and around 0.75 m from the edge of the plasma. The poloidal plane perpendicular to the camera line of sight (blue dotted line in Figure 3.3) was at a toroidal angle $\phi = 215°$, where the toroidal machine coordinate ϕ is 0° at the start of the first octant and continues round clockwise. The resulting spatial resolution at the tangency plane was around 5 mm. The camera was operated with an integration time of 3 μs, which a separate study has shown to be sufficiently short to avoid blurring of filaments at their typical velocities [51].

3.3 Camera spatial calibration

The 3D spatial coordinates of the scene observed by the camera is calibrated using the Calcam code [194]. Calcam makes use of functionality in the OpenCV computer vision library [195] to form a mathematical model that describes the mapping of real world 3D spatial coordinates to 2D image pixel coordinates.

The spatial calibration corresponds to a series of coordinate transformations, start-

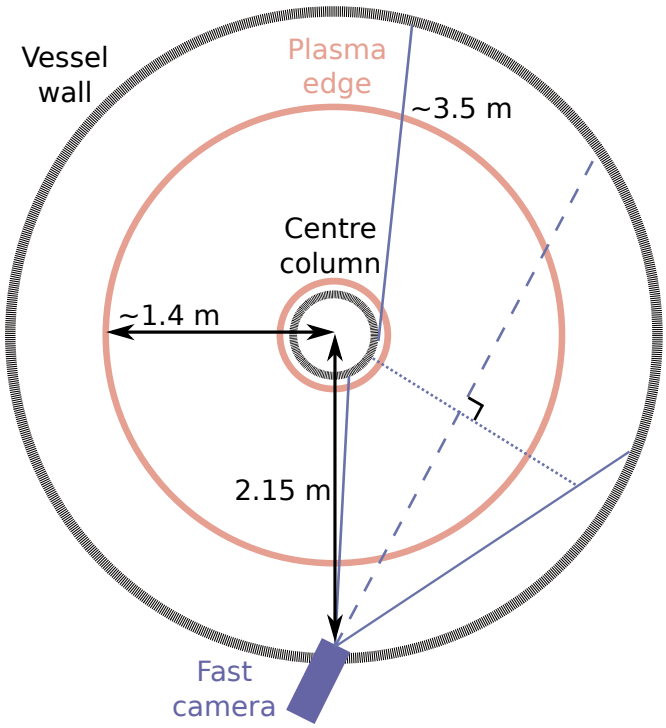

Figure 3.3: Top down schematic of MAST showing the field of view of the Photron SA-1 fast camera mounted at the mid-plane. The camera and its sight lines are shown in blue, and the plasma edge is shown in pink. The dotted blue line shows the tangency plane where the camera line of sight is tangent to the toroidal magnetic field.

ing with a translation and a rotation to map the origin to the location of the camera pupil and align the z-axis with the centre of the camera field of view. Next a perspective projection is applied to scale the x and y coordinates by the inverse of the distance from the camera pupil. An image distortion model is used to account for imperfections in the lens optics, before the *camera matrix*, describing the focal length of the lens and the sensor pixel size, is used to calculate the required image pixel coordinates. Details on the implementation of each of these steps can be found in the OpenCV documentation [195].

Calcam accounts for radial (barrel and pincushion) distortions and tangential (wedge-prism) distortions. Radial distortions arise from differences in the magnification of rays passing through the centre and edges of a lens and can be increased by apertures in the system. Barrel and pincushion distortions correspond to a decrease and an increase in magnification outwards from the centre of the image respectively. Tangential distortion arises from de-centring of optical components from the optical axis and leads to vertical variation in the magnification of the image. More detail on the distortion model can be found in [196]. Working with 1^{st} order radial distortion coefficients has been found to give the most robust results for this application, with rms (route-mean-square) fit errors routinely less than 1 pixel.

The camera registration enables points and paths in real space to be cast onto the

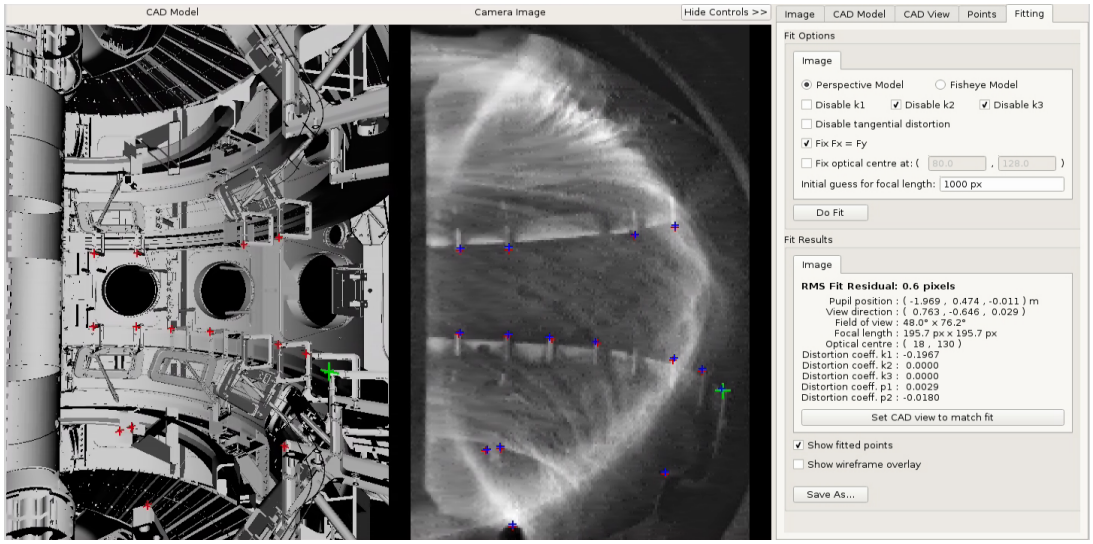

Figure 3.4: Mapping of point pairs between the 3D CAD model and an image from a discharge in the Calcam GUI.

image plane of the camera. In this way, 3D magnetic equilibrium information can be mapped onto the camera field of view in order to identify field aligned structures.

Figure 3.4 shows how the camera calibration is performed by matching clearly identifiable features in camera images with points on the MAST CAD model. Images taken during bright events like sawtooth crashes and disruptions light up the walls well and so typically make good calibration frames. Sharp corners are best for calibration points as it is important to constrain both the radial and toroidal location well. For this reason points on the centre column cannot be used in the calibration due to the uncertainty in their toroidal location. It is generally difficult to produce clearly identifiable point pairs in the top region of the machine. As a result the error in the spatial calibration can be slightly more pronounced at the very top of the image. However, as shall be seen in the next chapter, the analysis focusses on the central portion of the images where the spatial calibration is well constrained.

3.4 Discharges analysed

Data from three MAST discharges with pulse numbers 29852, 29840 and 29023 are analysed in this thesis. For these discharges, the camera was installed on the machine and the data was collected by James Harrison. Overviews of the time evolution of the main plasma parameters for each discharge are shown in Figure 3.5, while Table 3.3 summarises the properties of each discharge in more detail.

The majority of the data analysis uses visual camera data from MAST pulse 29852. This was a double null L-mode discharge with 0.7 MW of Ohmic and 1.7 MW of NBI heating power. The addition of NBI power made this a particularly bright L-mode

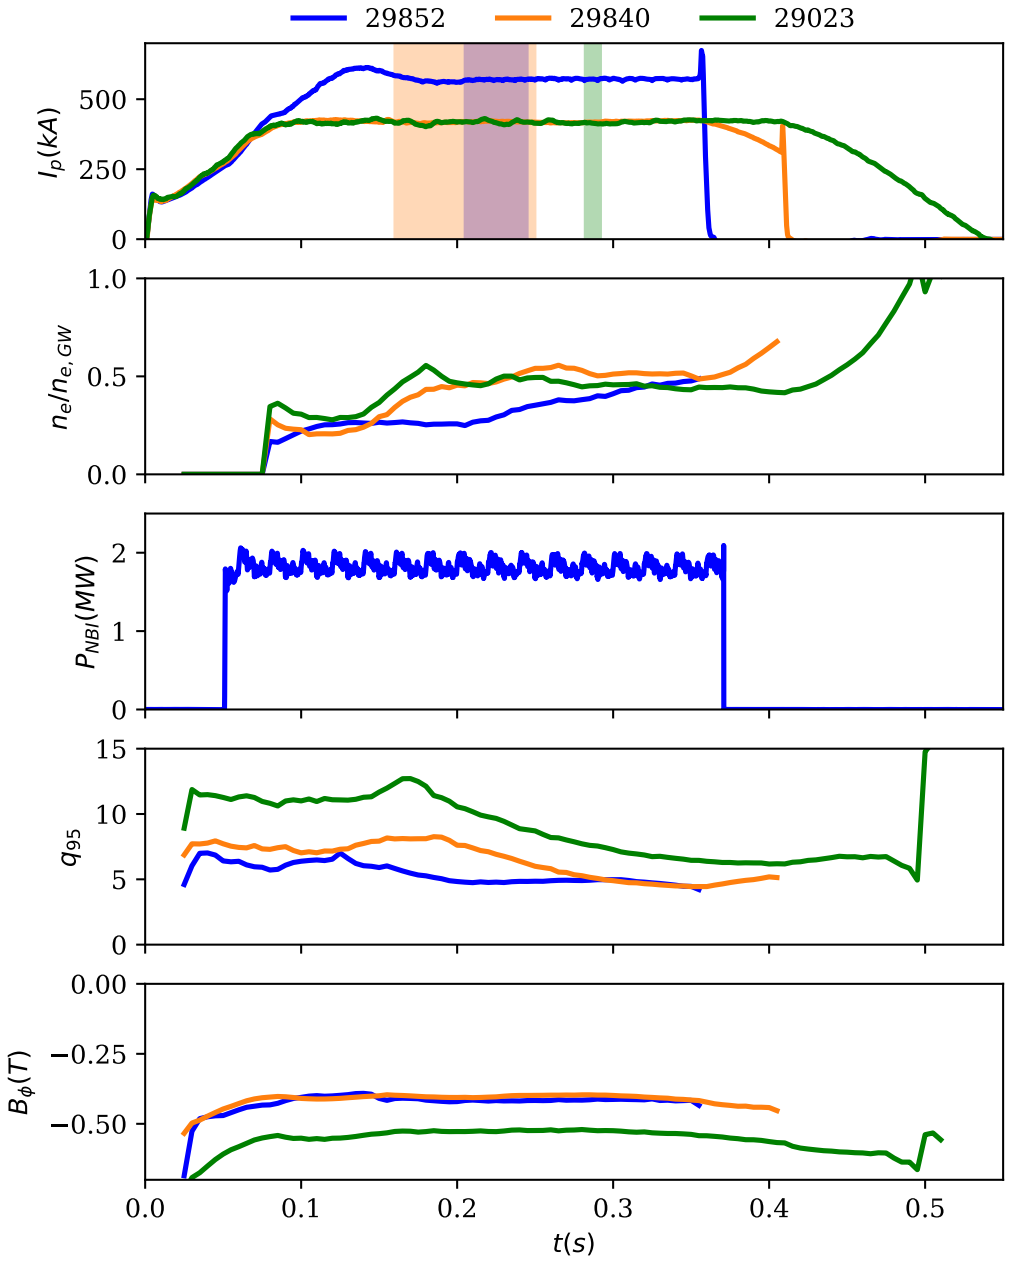

Figure 3.5: Plasma current, I_p, Greenwald density fraction, $n_e/n_{e,GW}$, NBI power, P_{NBI}, edge safety factor, q_{95} and on-axis toroidal magnetic field strength B_ϕ for the MAST discharges analysed in this work. Discharges 29852, 29840 and 29023 are indicated in blue, orange and green respectively. Analysis time windows are indicated by coloured bands of the same colours in the top panel.

Parameter	29852	29840	29023
Operating regime	L-mode	Ohmic	Ohmic
Magnetic configuration	Double null	Double null	Double null
Major radius, R_{axis} [m]	0.9	0.91	0.94
Minor radius, a [m]	0.57	0.61	0.62
Triangularity (average), δ	0.39	0.46	0.4
Plasma current, I_p [kA]	570	420	415
Plasma density (separatrix), n_e [$\times 10^{19}$ m^{-3}]	1.1	0.55	0.55
Plasma temperature (separatrix), T_e [eV]	35	10	15
Collisionality (separatrix), ν_*	13	77	34
NBI heating power, $PNBI$ [MW]	1.7	0	0
Greenwald fraction, n_e/n_G	0.29	0.45	0.45
Molecular density at wall, n_0 [$\times 10^{19}$ m^{-3}]	3.2	3.3	No data
Edge safety factor, q_{95}	4.8	7.6	7.6
Toroidal magnetic field (on-axis), B_T [T]	-0.42	-0.40	-0.52

Table 3.3: Plasma parameters for the analysed discharges.

discharge, in which filaments could be easily resolved and the quality of the camera data was particularly good. The camera analysis focused on the time window 0.20500-0.24500 s, corresponding to movie frame numbers 10500-14500.

During the current flat top the plasma current was $I_p \approx 0.56$ MA, with the EFIT++ magnetic equilibrium reconstruction code locating the outer mid-plane separatrix at $R_{sep} \approx 1.365$ m and giving an edge safety factor, q_{95}, of 4.8. The toroidal field strength was $B_T = -0.42$ T. The magnetic flux surfaces from EFIT reconstructions of this and two other discharges (discussed shortly) are shown in Figure 3.6.

Figure 3.7 shows the edge density and temperature profiles measured by the high resolution Thomson scattering (HRTS) system. These measurements indicate the separatrix electron temperature, T_e, was around 35 eV and the separatrix density, n, was around 1.2×10^{19} m^3. These values correspond to an edge collisionality of $\sim$13, indicating the plasma was weakly in the conduction limited regime. The molecular density at the vessel walls was estimated to be $\sim 3.2 \times 10^{19}$ m^{-3}, while the ionisation mean free path for neutrals in the SOL was estimated to be around $\sim$30 cm [99]. Small sawtooth oscillations were present during the discharge, with a period of $\sim$8 ms, but the analysis of the filaments was performed during quiescent periods between the crashes.

In Chapter 6 a comparison between the camera and conventional Langmuir probe analysis techniques is performed. Unfortunately, due to the limited number of discharges in MAST's M9 campaign for which the fast camera was operating on the machine and delays in the operation of MAST-U, there were no experiments with simultaneous camera and reciprocating probe measurements. As a result the comparison was made between two separate discharges with very similar operating parameters.

The camera analysis was performed on MAST discharge 29840 and the reciprocating

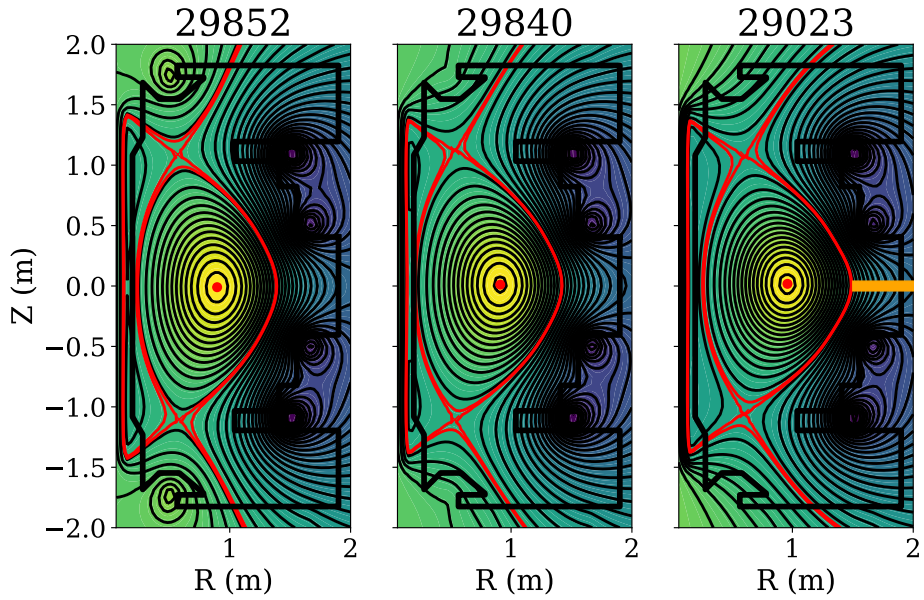

Figure 3.6: Poloidal cross sections of the plasma equilibria during the analysis time windows for the three discharges analysed here (left 29852, centre 29840, right 29023). The location of the reciprocating probe at the outboard mid-plane is indicated by the orange rectangle in the right most panel. Contours show the normalised poloidal flux, ψ_N, which is defined as 0 on the magnetic axis (yellow) and 1 at the separatrix. The separatrices are shown in red.

Langmuir probe analysis on discharge 29023, both of which were ohmic, L-mode, double null plasmas with a flat top plasma current of 420 kA and a line averaged plasma density of 4×10^{19} m^{-3}, corresponding to a Greenwald fraction of 0.50. The Langmuir probe data was collected by Scott Allen.

The magnetic equilibrium reconstruction places the separatrix around $R_{sep} \approx 1.54$ m in both discharges. The location of the reciprocating probe at the outboard mid-plane ($Z = 0$) is indicated in orange in the right most panel of Figure 3.6. The HRTS measurements in Figure 3.7 show the separatrix neutral densities were both $n \approx 0.55 \times 10^{19}$ m^3, while the separatrix electron temperatures were around $T_e \approx$ 10-15 eV. These measurements indicate the edge collisionality was around ν_*∼75 and ν_*∼35 for discharges 29840 and 29023 respectively, indicating these plasmas were also in the conduction limited regime.

The only significant difference between the two discharges was a 30% lower toroidal field of -0.39 T in discharge 29840, although at the times of the analyses, both discharges edge safety factors were 7.6, so that their edge magnetic geometries and thus connection lengths where very similar.

The camera analysis was performed over the time range 0.1600-0.2500 s during the plasma current flat top, corresponding to the movie frame number range 6000-14998. The probe reciprocation occurred over a 350 ms time window also during the plasma

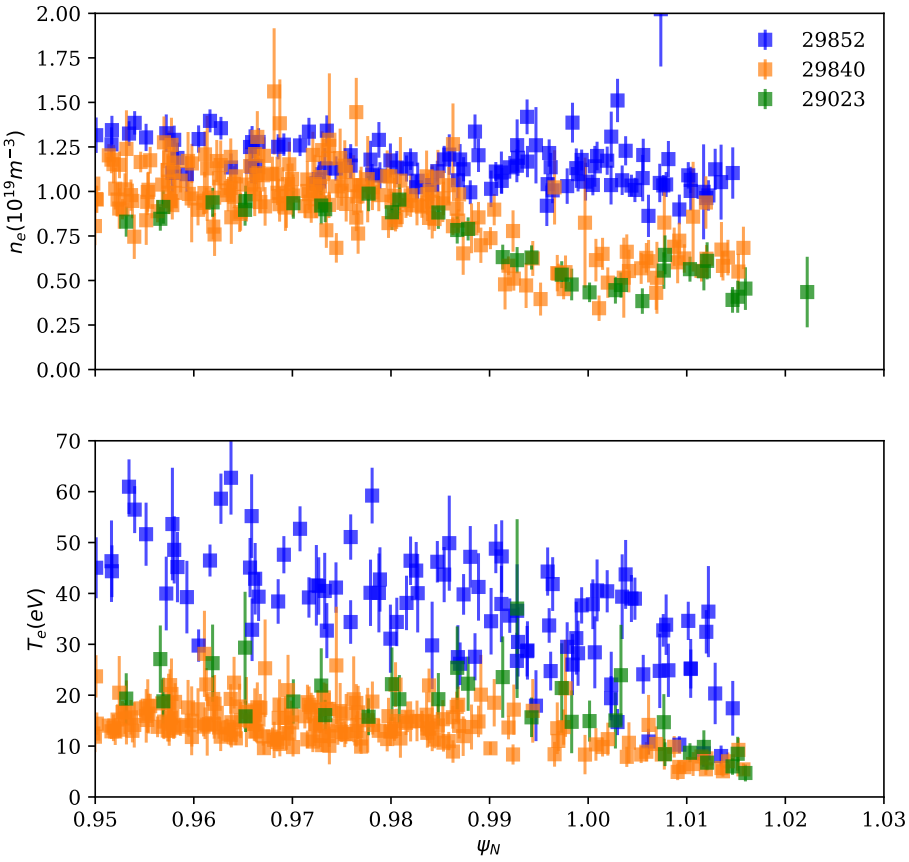

Figure 3.7: Edge Thomson scattering measurements of density (top) and temperature (bottom) as a function of normalised magnetic flux, ψ_N for discharges 29852, 29840 and 29023. The separatrix is located at $\psi_N = 1$.

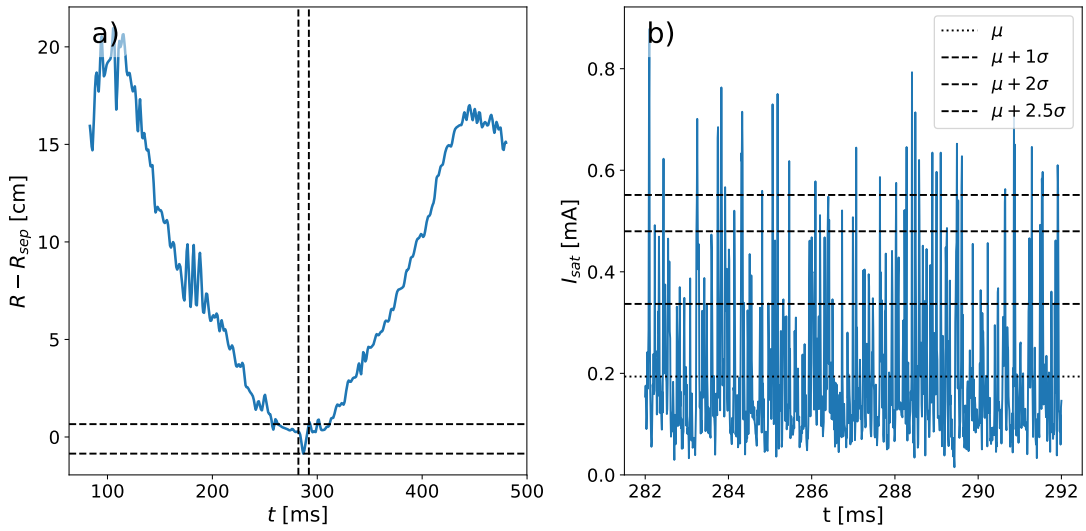

Figure 3.8: a) Position of reciprocating probe head during the reciprocation in discharge 29023. Vertical and horizontal dashed lines indicate the 10 ms analysis time window at the separatrix and the corresponding radial variation of the probe head over this time. b) Ion saturation current signal over the 10 ms time window around the separatrix. The mean and successive multiples of the standard deviation of the signal are indicated by horizontal dashed lines.

flat top, centred around $t = 0.287$ s and traversed to the separatrix from a distance of around 15 cm. A range of analysis was performed over the specific time window 0.282-0.292s for which the probe head was within ∼1 cm of the separatrix. The analysis was performed on the RCP's ion saturation current signal which had a sample rate of 500 kHz. Figure 3.8a) shows the radial location of the probe during its reciprocation and the time window over which measurements within ±1 cm of the separatrix were made. Figure 3.8a) shows the I_{sat} signal over this time window.

Chapter 4

The Elzar camera inversion technique

4.1 Introduction

In this chapter, we describe a new analysis technique that is used to identify and track filaments in fast camera data. The fast camera data contains a great wealth of valuable information about the filaments. However the complexity of the 3D information in the 2D images makes its analysis challenging. First, an overview of the assumptions on which the technique is based is given. This is followed by the procedure for mapping the light emission in the camera images to a field aligned basis set at the mid-plane, starting with the pre-processing steps applied to the images before describing the tomographic inversion procedure. Once the data have been inverted the filaments are located and analysed using the watershed algorithm (to be discussed later) to extract their positions, sizes and amplitudes. Finally, a tracking algorithm is described to follow the filaments between successive frames and measure their radial and toroidal velocities and lifetimes.

4.2 Camera tomographic inversion technique

The purpose of the camera analysis technique presented here is to take as inputs 1) camera images, 2) a camera spatial calibration and 3) a magnetic equilibrium description and return the values at the mid-plane of the positions, widths and amplitudes of filaments within the frames. To do this, we wish to map the complex 3D information in the camera images to a simplified plane of emission in which the filaments can be readily identified and analysed. This class of problems falls under the field of *tomography*.

Tomographic inversion problems are encountered in the analysis of other tokamak diagnostic data such as soft-X-ray imaging [197–199], bolometry [200–202], spectral imaging [203, 204] and linear plasma imaging [205].

Diagnostics that have overlapping views of the emission field from several different directions, such as bolometer diagnostics, can make use of the Radon transform [206].

In applications where only a single field of view is available, as is the case with a single camera view, Abel inversion based techniques are widely used [206].

The Abel inversion can be used to reconstruct axially symmetric emission profiles corresponding to circularly symmetric 2D emission functions $f(r)$, where r is the minor radius. Then the Abel transform, $F(y)$, describing the line integrated 1D projection along a ray with closest distance, y, from the symmetry axis ($r = 0$) is given by [207]:

$$F(y) = 2\int_y^\infty \frac{f(r)r}{\sqrt{r^2 - y^2}} dr. \tag{4.1}$$

Under the assumption $f(r)$ drops to zero more rapidly than $1/r$, the emission function can be recovered from an experimental measurement of $F(y)$ via the inverse Abel transform [207]:

$$f(r) = -\frac{1}{\pi}\int_r^\infty \frac{dF}{dy}\frac{dy}{\sqrt{y^2 - r^2}}. \tag{4.2}$$

By mapping magnetic flux coordinates onto cylindrical coordinates [208], the Abel inversion can be generalised to non-circularly symmetric flux functions in the poloidal plane.

Unfortunately, the 3D emission from plasma filaments that we wish to reconstruct does not meet the symmetry requirements for application of the Abel inversion for two reasons. Firstly, the helical emission from a filament cannot be mapped onto a flux function $f(r)$, parametrised by a single radial coordinate. Secondly, filamentary emission is not toroidally symmetric i.e. the emission in all poloidal planes cannot be described by a single 2D poloidal emission function. Instead we exploit the helical symmetry of the magnetic field that the filaments conform to. As the filaments are highly aligned to the magnetic field due to the efficiency with which the plasma propagates along field lines, the symmetry of the system parallel to the magnetic field can be utilised.

The technique presented here operates by using information about the magnetic geometry to reconstruct the 3D intensity information present in the 2D camera images onto the horizontal R-$R\phi$ plane at the mid-plane ($Z = 0$), where R is major radius, ϕ is the angular toroidal position and Z is the vertical machine coordinate. $R\phi$ is the toroidal arc length at the outer mid-plane and is used in place of ϕ as the toroidal coordinate, so that both dimensions have units of length, from which meaningful filament dimensions can be found. Due to the so called 'barbers pole effect' [209, 210] a parametrisation of uniform field lines in ϕ is also equivalent to a parametrisation in poloidal angle, θ.

The code used to perform the tomographic inversion and identify filaments in the camera images has been named *Elzar*[†].

The rest of this chapter proceeds with an overview of the assumptions on which the analysis technique is based, before describing the implementation of the technique,

[†]The name *Elzar* was chosen in reference to the *spice-weasel* background subtraction algorithm developed by Ben Dudson [47] that it utilises.

starting with pre-processing of the camera images, followed by the inversion of the camera data and ending with the identification of filaments within the inverted data.

4.2.1 Overview of assumptions

Two central assumptions about the nature of filaments are made in order to progress in interpreting the camera images:

1. Filaments are assumed to be well aligned to the background magnetic field which can be calculated via magnetic equilibrium reconstruction.

2. The light emission from a filament is treated as being constant in the direction parallel to the magnetic field.

The first of these assumptions is to be expected given the magnitude of parallel transport relative to perpendicular transport, which quickly spreads blobs of plasma along field lines. Furthermore, it is justified by a large base of experimental data from MAST [97] alongside other tokamaks [69, 125].

The second assumption contains two factors. First, filaments are assumed to be homogeneous along magnetic field lines in the camera field of view. This is justified *a posteriori* by observation of the fact that filaments are always observed to span the entire length of field lines within the camera's view. While some apparent parallel variation in light emission is observed, it is many times weaker than the cross-field variation. This is expected, as parallel transport is very efficient on the relevant time scales and so will tend to smooth out differences along field lines. In the absence of a precise description of the parallel variation of the filaments, treating them as homogeneous is sufficient for us to proceed. The second factor is the assumption that the neutral density in MAST remains homogeneous along the length of the filament. This is motivated by the open vessel design of MAST which leads to homogenous neutral distributions away from the divertor [189]. While this assumption breaks down close to the divertor surface, we are only concerned with the part of the filaments that extends between the X-points, as below the X-point the filaments are heavily distorted by the magnetic shear.

In principle, both of these assumptions could be relaxed through suitable modifications to the technique. However, the complexity this introduces was not deemed to be worthwhile at this stage.

Combining the two assumptions above, the filament structures observed in fast camera images can be formed by a linear superposition of images of uniformly emitting field-lines from the equilibrium magnetic field, cast onto the camera field of view. The matrix of camera pixel intensities forming the camera image, I, can therefore be expressed as a weighted sum of basis images of uniformly emitting field lines, I_{basis}:

$$I = \sum_{j=0}^{N} \epsilon_j I_{basis,j} \tag{4.3}$$

where ϵ_j describes the relative contribution of the j^{th} field line basis image to the camera image. The summation is valid due to the edge plasma being optically thin to visible radiation [206]. Equation 4.3 represents the mapping of the information in the camera images between a pixel basis set and a field line basis set. All the information required to construct the field line basis set is contained in the magnetic equilibrium data and the camera calibration that describes how emission from 3D point sources are mapped to pixels on the camera sensor. We wish to convert the many non-local line integrated measurements from each pixel in the camera image into an estimate of the local plasma emissivity. The expression of the camera information in the field aligned basis set facilitates the identification of filaments, since field aligned plasma structures appear as blob-like clusters of high intensity field lines in 2D plots of ϵ_j. By analysing the regions of enhanced emissivity present in ϵ_j, the filaments' sizes and locations can be deduced in a similar way as for poloidal intensity distributions produced by GPI or BES diagnostics. Therefore, the vector of field line basis weightings, ϵ_j, containing the filament information is the quantity that we seek to calculate.

4.2.2 Image pre-processing

A series of pre-processing steps are applied to the camera images before the data are inverted. The pre-processing technique is motivated by the fact that filaments in the SOL represent positive density perturbations on top of a background, so that perturbations in the light emission are also positive. Therefore, we define the background as the minimum in the signal over time. Thus a background subtraction technique is applied to the movie to remove the low-frequency background component of the light emission and extract just the fluctuating part containing the filaments. This step is necessary to avoid the filaments being washed out by background light emission in the SOL and prevent first wall structures confounding the analysis.

The background calculation, which follows in a similar manner to that described by Dudson [47], is achieved by taking the pixel-wise minimum in the light intensity over a set of N frames that preceded the frame of interest in time,

$$I_{bg,i}^{(n)} = \min\left(\left\{ I_i^{(k)} \,\middle|\, n - N < k \leq n \right\}\right), \tag{4.4}$$

where $I_{bg,i}^{(n)}$ is the background intensity of the ith pixel for frame n and $I_i^{(k)}$ is the intensity of the ith pixel in image k. Typically N is set to 10, corresponding to taking the minimum in intensity over a 0.1 ms time window preceding the frame of interest. This technique has previously been applied in Refs. 47, 51, 96, 97 and discussion of the validity of the technique can be found therein.

Other background calculations were considered by instead taking the background as represented by either the mean, median or the low frequency component of the total emission extracted from a pixel-wise FFT. In the first two cases, this led to non-physical

negative perturbations in the SOL, which could not be reconciled with Langmuir probe measurements [67]. The latter performed similarly to the minimum technique, however the calculation was considerably slower, so the minimum based background subtraction was adopted.

Following background subtraction, a bi-lateral median filter [211] is applied, which is a gradient preserving noise removal method. This filter corresponds to taking the median of each pixel and its neighbours of similar intensity and is implemented in OpenCV. For this application a 5×5 neighbourhood of pixels is typically considered. The similarity of the neighbouring pixels is assessed by a Gaussian 'similarity' function that enables the filter to preserve sharp edges by excluding pixels with very different values, such as occur one side of edge features. The reduced dynamic range resulting from the background subtraction amplifies the relative contribution of the shot noise to the image, making noise removal particularly useful. The noise removal is followed by a Gaussian de-blur [212] which helps to re-sharpen the image after the smoothing of the noise removal. This is implemented through the weighted subtraction of the image convolved with a Gaussian kernel:

$$I' = (1 + w)I - wI * G\,, \tag{4.5}$$

where I' is the de-blured image, w is the weighting of the de-blur operation (typically 0.5 here) and G is a matrix describing a set of 2D Gaussian weights given by

$$G(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right), \tag{4.6}$$

where x and y are pixel coordinates and σ is the symmetric Gaussian width (typically set to the kernel size of $\sigma = 15$). The OpenCV implementation of the Gaussian blur function is used here.

These two filters, while not vital elements of the technique like the background subtraction, have been found to effectively reduce the propagation of noise through the analysis. They are particularly useful for working directly with the partially inverted moment matrix data that does not benefit from regularisation, which will be discussed in the next section. Figure 4.1 shows an example of an image with each of the pre-processing techniques described in this section applied in turn.

4.2.3 Tomographic inversion

The next stage of analysis is an inversion of the signal (as results from the previous three stages of pre-processing) onto real space coordinates. In order to perform the tomographic inversion of the camera data, a set of basis field line images are first produced. These are created by projecting magnetic field line trajectories onto the camera image plane. The field lines are traced using a 4^{th} order Runge-Kutta integrator built into the UKAEA pyEquilibrium python package, with the magnetic field structure

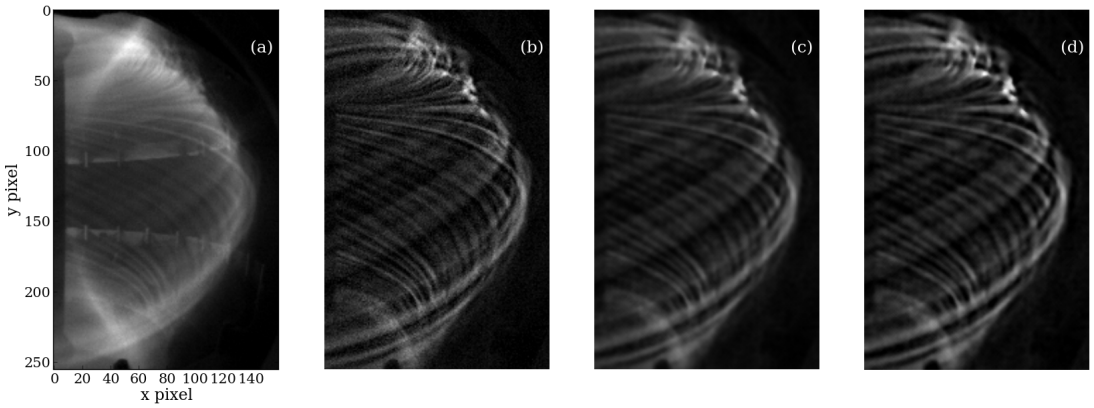

Figure 4.1: a) Raw image from MAST shot 29841 at 0.22424s. b) Background subtracted image using 10 frames prior to the desired frame to construct a background. c) Application of the bi-lateral median filter. d) Application of a Gaussian de-blurr.

provided by an EFIT++ [213, 214] kinetic equilibrium reconstruction. EFIT++ solves the Grad-Shafranov equation for a 2D magnetic equilibrium in an axisymmetric torus, giving the toroidal and poloidal flux functions from which the toroidal and poloidal magnetic fields are calculated. The field lines are parametrised by their launch positions at the mid-plane $(R, R\phi)$, about a central toroidal angle ϕ_0. The field lines are followed for one toroidal rotation, so that they only enter the camera field of view once. The tracing is typically performed in parallel steps of 5 mm, in order to ensure that each pixel along the length of the field line is sampled. The spatial calibration of the camera from Calcam is used to project the 3D coordinates describing the points along the field lines onto 2D pixel coordinates in the camera image plane, according to the camera field of view and distortion model. With the pixels imaging each relevant field line identified, this information is used to produce an image of each field line, accounting for the inverse square reduction in intensity with distance from the camera pupil and volumetric line of sight integration effects. An example of a field line basis image is given in Figure 4.2a). Summing a collection of similar localised field line basis images according to Equation 4.3 results in the image of the uniformly emitting filamentary flux tube in Figure 4.2b). Note that the flux tube is fainter at the top of the image due to fanning of the field lines and consequent stretching of the flux tube's cross section as the field lines get closer to the X-point. This results in the observed drop in emission which is also seen in experimental images.

Field lines are traced for a regular grid of launch positions in the R-$R\phi$ plane, so that the set of points along all the field lines describe a volume of the SOL that is visible in the camera field of view. It is important to note that, while each magnetic line is distinct in 3D, once they are projected onto the camera's 2D image plane, many of the field lines overlap in the camera's field of view so that they are no longer unique (i.e. they share common pixels). Therefore, the field line basis set is *not* orthogonal,

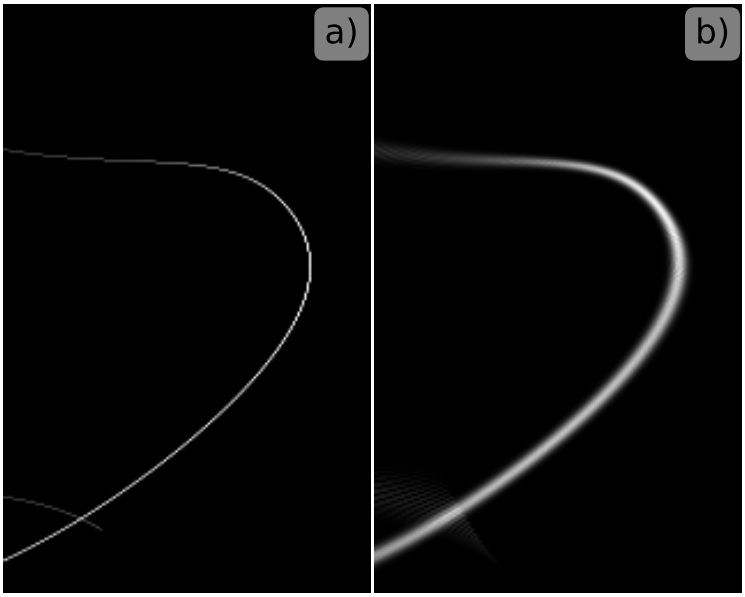

Figure 4.2: a) Example of a basis field line image used to construct the background subtracted camera images. The flattened basis image forms a single column of the geometry matrix $\mathbb{G}_{ij}$ (see Equation 4.7). b) Emission from a bundle of field lines with emission amplitudes, ϵ_j, given by a 2D Gaussian in the toroidal plane.

such that $\langle I_{basis,i} \cdot I_{basis,j} \rangle \neq 0$. This 2D grid determines the resolution of the image of the inverted field aligned emission, ϵ. Figure 4.3 illustrates the relation between image plane and the inverted emission plane onto which the camera data are mapped. Note that due to the narrow radial width of the SOL the toroidal extent of the inversion domain of 120 cm is much greater than the radial extent of 18.4 cm. In real, space the inversion domain takes the form of a curved trapezoidal plane, although for simplicity it is represented as rectangular in Figure 4.3a).

Figure 4.4 illustrates the geometry of the inversion domain further by re-projecting the emission at the mid-plane, obtained from the inversion technique, along field lines onto a series of a) toroidal and b) poloidal planes. The red field lines mark the bounding corners of the inversion domain. The central toroidal plane shows the original inversion at the mid-plane, while the planes above and below show the re-mapped emission. The mappings of emission to poloidal planes are reminiscent of the data produced by GPI and BES diagnostics, although covering a much larger region of the SOL.

The inversion problem can be rigorously described if it is cast into linear algebra. The camera image can be thought of as a matrix of pixel intensities with dimensions $m \times n$. This matrix can be collapsed into a single contiguous image vector, $\mathbf{I}$, with elements I_i, and dimensions $mn \times 1$ ($i = 1, 2, \ldots, mn$). For the images from the SA1.1 camera with which we are concerned, m and n are given by the camera resolution, $m \times n = 256 \times 160$, so that $\mathbf{I}$ is a vector of $mn = 40960$ pixel intensities. Likewise, the

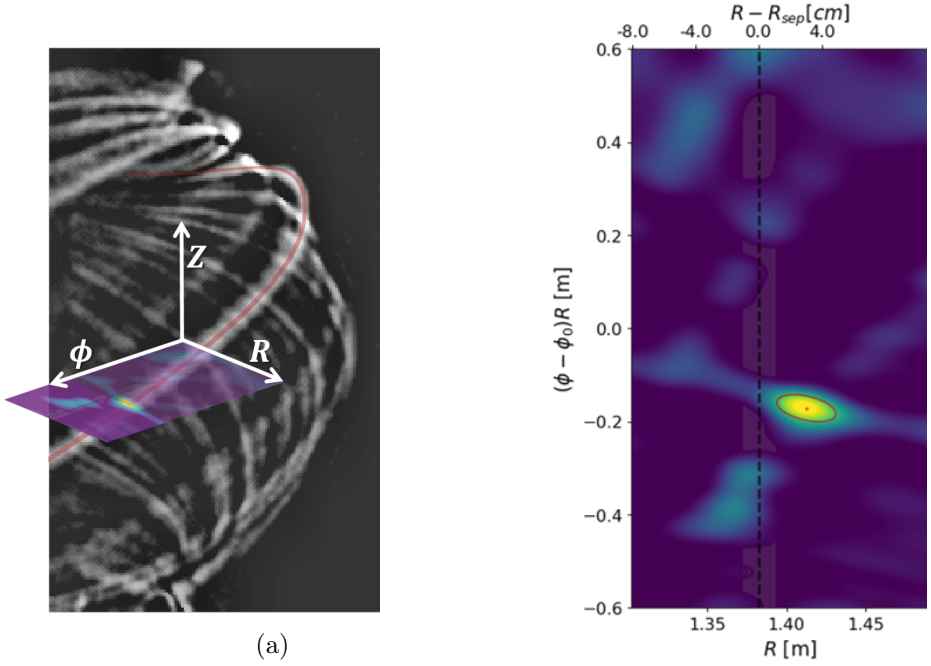

(a) (b)

Figure 4.3: a) An experimental camera frame from MAST shot 29852, with the pre-processing described in §4.2.2 applied. b) Inversion calculated using the frame in a). The inversion in b) is also superimposed on the frame in a) to illustrate the (R, ϕ) mid-plane coordinate system.

result of the inversion describing the field aligned emission in the toroidal plane with dimensions $p \times q$ can be flattened into a vector, $\boldsymbol{\epsilon}$, with elements ϵ_j and dimensions $pq \times 1$, equal to the number of field lines in our emission basis ($j = 1, 2, \ldots, pq$). Then we can compose a *geometry matrix*,

$$\mathbb{G}_{ij} = [\mathbf{I}^{(0)}_{basis},\ \mathbf{I}^{(1)}_{basis}, \ldots,\ \mathbf{I}^{(pq-1)}_{basis},\ \mathbf{I}^{(pq)}_{basis}], \tag{4.7}$$

where the column vectors, $\mathbf{I}^{(j)}_{basis}$, correspond to the flattened basis field line images (see example image in Figure 4.2a)). Here, $\mathbb{G}_{ij}$ represents the element from the ith row and jth column of matrix $\mathbb{G}$. The resulting geometry matrix, $\mathbb{G}_{ij}$, has dimensions $mn \times pq$ and contains all the necessary information about the magnetic equilibrium structure and camera viewing geometry.

It has been found empirically that to be able to represent a mid-plane camera image with a basis of field-lines in this manner requires a well resolved grid, typically with a radial spacing $\lesssim 4$ mm and a toroidal spacing $\lesssim 0.2°$. For the work presented here, the resolution of the inversion is typically of the order $p \times q = 47 \times 123$, corresponding to a total of $pq = 5,781$ distinct field lines in the basis. Further details on the choice of inversion resolution will be discussed later.

Equation 4.3 can now be expressed in matrix format, so that the intensity of an individual pixel in the camera image can be written as the sum of the intensities associated

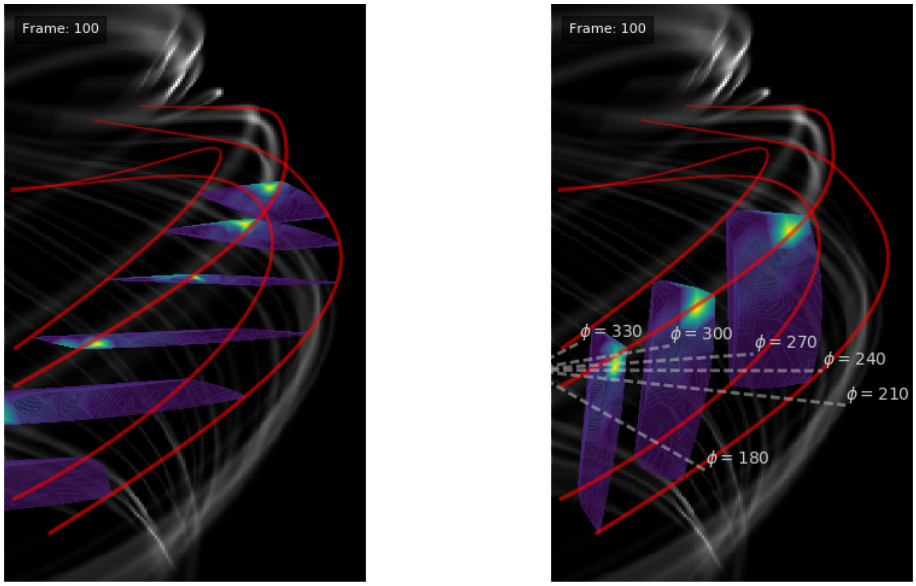

(a) (b)

Figure 4.4: Illustration of the domain of the inversion by mapping intensities along field lines to a series of a) toroidal planes and b) poloidal planes. The red field lines mark the bounding corners of the inversion domain.

with the field line basis vectors composing the columns of $\mathbb{G}$ according to:

$$\underbrace{I_i}_{mn\times 1} = \underbrace{\mathbb{G}_{ij}}_{mn\times pq}\underbrace{\epsilon_j}_{pq\times 1}, \tag{4.8}$$

where the Einstein summation convention is used to represent sums over common indices. The basis field line image in Figure 4.2a) can now be thought of as forming a column of $\mathbb{G}$. The image in 4.2b) results from the product of $\mathbb{G}$ with an assigned (rather than measured) emission vector describing a 2D Gaussian emission profile in the toroidal plane.

Once the matrix $\mathbb{G}$ has been computed it can be applied to describe any number of frames for which the magnetic equilibrium is not changing much. Experimentally, it typically takes more than 500 μs for the separatrix to move by more than $\sim$1 cm. Hence a given geometry matrix can normally be used for at least 50 consecutive frames.

Equation 4.8 represents an overdetermined system of nm linear equations in pq unknowns, $40960 \approx mn \gneqq pq \approx 5781$, so that $\mathbb{G}_{ij}$ contains more information than just that needed to construct ϵ_j. As $\mathbb{G}_{ij}$ is not square, it does not have a true inverse, so we cannot invert $\mathbb{G}_{ij}$ directly to solve Equation 4.8 for ϵ_j. Two factors limit the size of the field line basis set, such that this problem cannot be overcome by simply using a larger basis of 40960 in order to make $\mathbb{G}$ square. First, increasing the size of the geometry matrix increases the computational cost of calculating its inverse. Since we wish to calculate inversions for many thousands of camera frames, this computational cost should not be too large. Secondly, the camera resolution, magnetic geometry

and viewing geometry are such that as the size of $\mathbb{G}$ is increased further, many of the additional basis field line images are either degenerate, or contain little additional information constraining the inversion (i.e. additional field lines are either outside of the camera's field of view or are so close to each other that the field lines cross the same pixels in the camera image). Therefore, a square 40960×40960 geometry would be both large and singular due to the linear dependency of its columns. Therefore, the system has no unique solution and we instead take an ordinary least squares approach to finding the 'best' value of the inverted emission vector $\hat{\boldsymbol{\epsilon}}$ that minimises the sum of the quadratic difference between the true camera image and the field aligned basis description of the image:

$$\hat{\boldsymbol{\epsilon}} = \underset{\boldsymbol{\epsilon}}{\operatorname{argmin}} \left\| \mathbb{G}\boldsymbol{\epsilon} - \mathbf{I} \right\|^2, \tag{4.9}$$

where 'argmin' is the operator returning the value of $\boldsymbol{\epsilon}$ that minimises its argument. This minimization problem has a unique solution, provided none of the field line basis images are degenerate, so that the columns $\mathbf{I}^{(j)}_{basis}$ of $\mathbb{G}$ are linearly independent. The ordinary least squares solution is then given by the one-sided *left inverse* [215]:

$$\begin{aligned} \hat{\boldsymbol{\epsilon}} &= (\mathbb{G}^T\mathbb{G})^{-1}\mathbb{G}^T\mathbf{I} \\ &= \mathbb{H}^{-1}\mathbb{G}^T\mathbf{I} \\ &= \mathbb{H}^{-1}\boldsymbol{\mathcal{E}}, \end{aligned} \tag{4.10}$$

where $\mathbb{G}^T$ represents the transpose of $\mathbb{G}$, $\mathbb{H} = \mathbb{G}^T\mathbb{G}$ is termed the *Gramian matrix* of $\mathbb{G}$ [216] and $\boldsymbol{\mathcal{E}} = \mathbb{G}^T\mathbf{I}$ is the *moment matrix*. This is an example of the More-Penrose *pseudo-inverse* for a matrix of full rank. Given that $\mathbb{H}$ is square, it has a true inverse, $\mathbb{H}^{-1}$, satisfying $\mathbb{H}\mathbb{H}^{-1} = \mathbb{I}$, where $\mathbb{I}$ is the identity matrix (not to be confused with the image vector $\mathbf{I}$). Thus the problem of finding $\hat{\epsilon}$ has been reformulated to involve the inversion of a square matrix, $\mathbb{H}$.

As in all ill-posed problems, the transformation represented by our geometry matrix has a 'smoothing' effect on the 3D emission information, so that the reverse process, represented by the inverse matrix, will tend to have the opposite effect and amplify any high-frequency components in the emission, as described in [217]. Therefore, random noise lacking correlation between neighbouring pixels is prone to amplification, to the extent that it can overwhelm the required solution. Therefore, regularisation schemes are widely applied in these situations that impose further constraints on the solution [206], effectively damping high frequency components of the solution, which is equivalent to truncating the singular value expansion (SVE) of the matrix before the small singular values start to dominate [217]. Regularisation supplements the system described by Equation 4.9 with additional *a priori* information about the smoothness, entropy, or probability distribution of the solution. This additional information is supplied through a regularisation function $O(\boldsymbol{\epsilon})$.

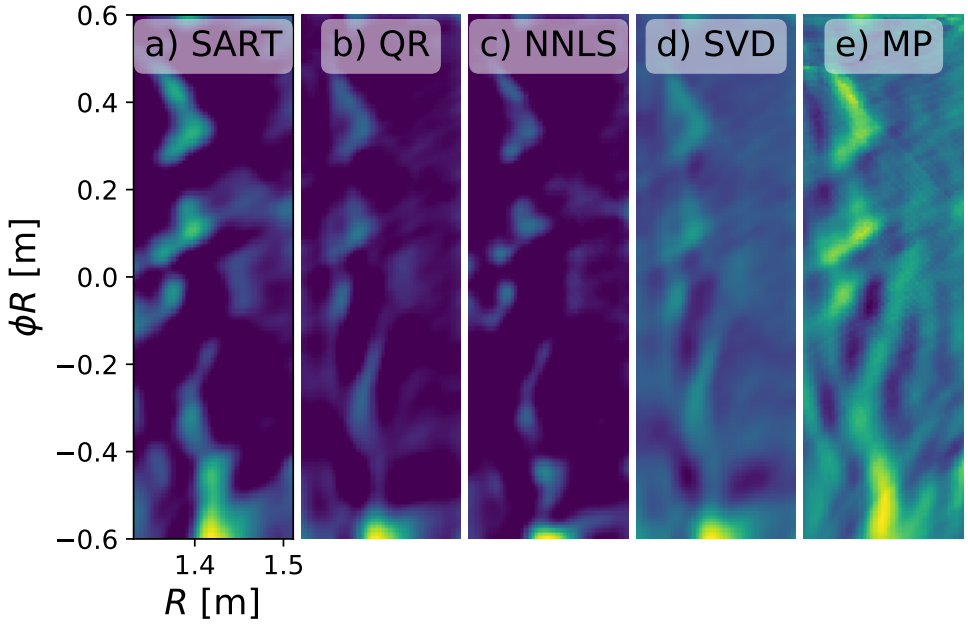

Figure 4.5: Comparison of the inverted emission profiles, $\hat{\boldsymbol{\epsilon}}$, produced using different inversion algorithms to calculate the inverse of $\mathbb{H}$. Each algorithm has been applied to the same experimental camera 6043 from MAST discharge 29852. The first four algorithms have been applied with Laplacian regularisation of equal strength, with the MP algorithm has applied smoothing through truncation of the set of singular values in the SVD of $\mathbb{H}$.

The general expression for the constrained optimisation of a system using Phillips-Tikhonov regularization is

$$\hat{\boldsymbol{\epsilon}} = \underset{\boldsymbol{\epsilon}}{\operatorname{argmin}} \left(\left\| \mathbb{G}\boldsymbol{\epsilon} - \mathbf{I} \right\|^2 + \lambda O(\boldsymbol{\epsilon}) \right), \tag{4.11}$$

where λ is the regularisation parameter that sets the strength of the regularisation condition, and whose optimal value depends on the level of noise in the system. The value of the regularisation parameter is chosen as a compromise between propagation of the errors in the data and over-smoothing of the data.

For $\lambda \to 0$ the solution depends primarily on the data and Equation 4.11 reduces to equation 4.9, while for very noisy systems for which λ is set to large values, the solution depends heavily on the regularisation condition. Many options for $O(\boldsymbol{\epsilon})$ exist, including linear (i.e. norm, gradient, Laplacian), maximum entropy and minimum Fisher information functionals [218].

As the density field in the SOL is a smoothly varying quantity, we apply Laplacian regularisation, with $O = \Delta^T \Delta$, where Δ represents the finite difference representation of the Laplacian operator. This minimises the curvature of the solution ensuring that the resulting emission is smoothly varying and provides the constraint needed to identify a unique solution.

Inversion technique	Frame inversion time [s]	Implementation	Parameters
SART	29	Implementation of Ref. 219 in `Python`	$\lambda = 5 \times 10^{-2}, n_{iter} = 400$
QR	53	`numpy.linalg.qr`	$\lambda = 5 \times 10^{-2}$
MP	118	`numpy.linalg.pinv`	$\lambda = 0, r_{cond} = 3 \times 10^{-2}$
SVD	227	`numpy.linalg.svd`	$\lambda = 5 \times 10^{-2}, r_{cond} = 0$
NNLS	952	`scipy.optimize.nnls`	$\lambda = 5 \times 10^{-2}$

Table 4.1: Timings and implementations of the different inversion algorithms compared in Figure 4.5.

A range of algorithms exist that can be applied to the calculation of the inverse of $\mathbb{H}$. Figure 4.5 compares results from using five different numerical inversion algorithms: a) non-negative simultaneous algebraic reconstruction (SART), b) QR factorisation (QR) [220], c) non-negative least squares (NNLS) [221], d) singular value decomposition (SVD) [222] and e) More-Penrose pseudo-inverse (MP) [223]. The timings and implementations of each algorithm are given in Table 4.1. Each algorithm, with the exception of MP, has been applied with Laplacian regularisation to constrain the solutions, while the MP algorithm applies SVD without regularisation, but instead truncates the singular values in the expansion of $\mathbb{H}$ at a threshold, r_{cond}.

As can be seen from the timings of each algorithm listed in Table 4.1, the SART algorithm is the fastest, whilst producing qualitatively similar results to the QR and NNLS algorithms. Furthermore, a qualitative assessment of the emission structures produced by the algorithms, found the SART algorithm produced inversions from which the filaments in the camera images could most consistently be identified. For these reasons, together with the robustness of the algorithm to noise in the camera data, the non-negative SART algorithm was selected for subsequent analysis.

The constrained SART algorithm used here is a python implementation of the C-SART algorithm described in [219], with the addition of a threshold non-negativity constraint at each iteration of the algorithm. This constraint was added to avoid unphysical negative emission in the inversion solution and is the simplest non-negativity constraint discussed in Ref. 224. The emission of the j^{th} field line on the k^{th} iteration of the SART algorithm is given by [219],

$$\hat{\epsilon}_j^{(k)} = \epsilon_j^{(k-1)} + \frac{\omega}{\mathbb{G}_{\oplus j}} \sum_{i=1}^{mn} \frac{\mathbb{G}_{ij}}{\mathbb{G}_{i\oplus}} \left(I_i - \mathbb{G}_{ij} \epsilon_j^{(k-1)} \right), \tag{4.12}$$

where the relaxation parameter $0 < \omega < 2$ controlling the speed of convergence is typically set to 1, and $\mathbb{G}_{\oplus j} = \sum_{i=0}^{mn} \mathbb{G}_{ij}$ and $\mathbb{G}_{i\oplus} = \sum_{j=0}^{pq} \mathbb{G}_{ij}$ are the column and row sums of $\mathbb{G}$ respectively. The algorithm typically converges after 300-400 iterations. This generally corresponds to less than a minute per frame.

In the absence of regularisation or truncation of the SVD expansion, both approaches

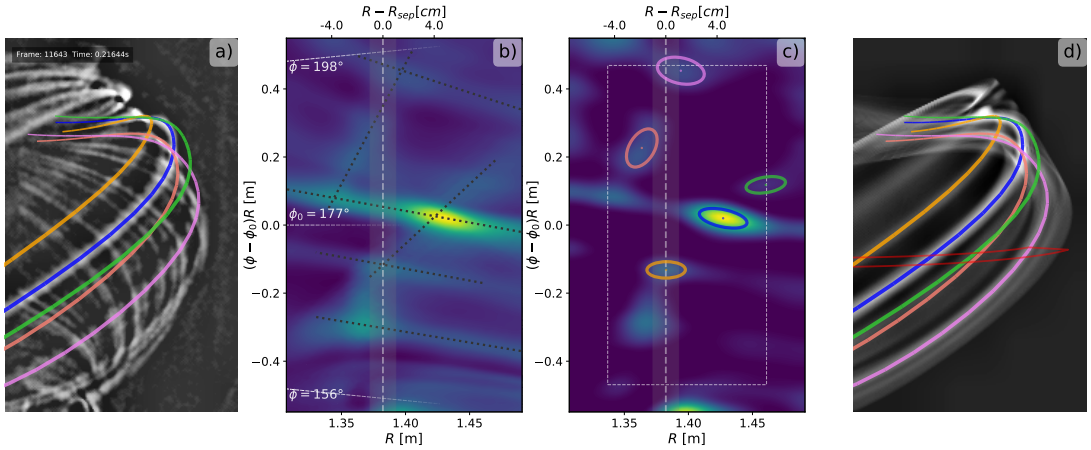

Figure 4.6: a) Pre-processed experimental camera frame from MAST shot 29852 with field lines superimposed marking the centres of the detected filaments. b) Emissivity data from the moment matrix $\boldsymbol{\mathcal{E}}$. The range of toroidal angles covered by the inversion domain is indicated. c) Inverted field aligned emissivity represented by $\hat{\boldsymbol{\epsilon}}$. Coloured ellipses indicate the filaments projected on the frame images. The white dashed rectangle indicates the analysis region which avoids boundary effects and the white vertical dashed line shows the location of the separatrix. d) Re-projection of the emissivity data in c). The red line indicates the inversion plane from which the emission is projected along field lines.

are swamped by high frequency noise and no filamentary structures can be resolved. While the regularisation overcomes this, a consequence is that genuine small scale structures with scale lengths of the order of the analysis grid resolution will be suppressed as well as the noise. Therefore, filaments with widths below ~ 5 mm will not be reliably reproduced. The strength of the regularisation parameter, λ, was tuned empirically so as to avoid spurious sharp banding patterns in $\hat{\epsilon}$ resulting from noise, while avoiding over smoothing of the large scale physical structures in the inversion.

4.2.4 Example application of the inversion algorithm

We now present an example of how the technique is applied to the inversion of experimental camera data from MAST. Figure 4.6 illustrates the two main calculation steps in the inversion process, along with a re-projection of the inversion output back onto the camera field of view. Figure 4.6a) shows an experimental background subtracted camera image from MAST discharge 29852 which is to be inverted. In the next panel, Figure 4.6b) shows the intensities in the moment matrix, $\boldsymbol{\mathcal{E}}$, resulting from the product of the transpose of the geometry matrix with the image vector produced by flattening the image in Figure 4.6a). The range of toroidal angles spanned by the inversion domain are also shown in white. It is worth noting at this point that the moment matrix is very inexpensive to compute, but is itself an approximation of $\hat{\boldsymbol{\epsilon}}$. The structure of the moment matrix is discussed further in the next section. Figure 4.6c) shows the inversion

resulting from the application of the SART algorithm to solve Equation 4.11. A number of bright blobs of field aligned emission are visible corresponding to filaments in the camera image. These structures are identified and measured using the filament identification algorithm described shortly, with the resulting detections indicated by coloured ellipses. Finally, Figure 4.6d) shows the re-projection of the information captured in $\hat{\boldsymbol{\epsilon}}$ back onto the camera field of view. The red line bounds the inversion plane in which the emission is mapped.

The re-projection is produced by taking the product of the geometry matrix with the inverted emission vector. This process will be discussed further in the next chapter in relation to the generation of synthetic camera images. The re-projection is useful for interpreting the inversion information and comparing how effectively the structures in the original camera image are reproduced. It also helpful for understanding the portion of the camera image spanned by the field lines in the inversion domain (also refer back to Figure 4.4). Field lines at the centres of the detected filament ellipses in Figure 4.6c) are over-plotted in Figures 4.6a) and 4.6d) in the same colours. This makes it possible to identify which structures in the camera images are responsible for the structures in the inversion. The identified structures align well with the filaments visible in the camera image, giving a visual indication of the quality of the detections.

The design choice to analyse a subset of the toroidal domain is motivated by the large computational gains that can be made by making the geometry matrix smaller. This reduces the size of the geometry matrix by around an order of magnitude, leading to a significant gain in computation time. Given many of the field lines excluded from the analysis domain are either outside the field of view of the camera, are only partially visible or are in regions of the image where the filaments are poorly resolved, minimal useful information is lost, while information about the most readily observable filaments is retained.

The chosen analysis domain focuses on the foreground filaments. This is because these filaments are closest to the camera so their widths and positions can be best resolved. Filaments in the background of the camera view can be only a few pixels across, so only their toroidal positions can be well constrained. The field lines at the top and bottom of the images are excluded from the analysis domain since only short portions of them are in the field of view of the camera, so that the properties of the filaments there cannot be well constrained either. Furthermore, the regions close to the divertors suffer from further complicated factors such as shearing of the filament cross sections, limited magnetic field pitch, reduced camera registration accuracy and stronger neutral density variation. Therefore, at least at this stage, analysis of filaments whose visible portions are located in the divertor regions is avoided.

While the background filaments do not fall within the analysis domain, they do overlap with the foreground filaments we wish to measure. However, they are near-

perpendicular to the foreground filaments so that their contribution to the intensity is similar for all foreground field lines. Therefore, to first order, the additional intensity from background filaments appears as a background offset in the foreground filament intensities and does not significantly impact the analysis of the foreground filaments. This is discussed further in the next chapter.

A consequence of the reduced inversion domain is that, at its boundaries where there is a sharp cut off in the basis, the SART algorithm struggles to distribute emission correctly. This results in a commonly observed artefact where excess emission is placed at the borders of the inversion domain, particularly in the far SOL. For this reason, filament detection is only applied to structures with centres within the region defined by the dotted white border box in Figure 4.6c).

4.2.5 Artefacts in the reconstructed emission

Insight into artefacts in the reconstructed emission can be gained by considering the structures present in the moment matrix, shown in Figure 4.6b). The moment matrix gives the sum total emission in the camera image along each field line in the basis. Therefore emission from a single pixel in the image contributes equally to each field line that intersects that pixel in the camera field of view. This inconsistent 'double counting' of emission leads to the spreading of emission from the true origin of the emission to all intersecting field lines, leading to distortion artefacts. By contrast, the least squares solution attempts to self-consistently distribute the emission amongst the field lines such that it is only placed on a given field line if no other linear combination of basis field lines better reproduces the emission pattern in the camera image. Therefore, any remaining artefacts from line of sight effects in the inverted emission, originate in stronger forms in the moment matrix. The geometry of the field lines resulting from the smoothly varying magnetic equilibrium means that overlap of field lines in the camera field of view tends to be greatest for field lines that are spatially close to one another. Therefore, the moment matrix captures the main structural features of the inversion solution, but these structures are distorted over smaller length scales by the inter-field line dependencies, resulting in banding and blurring in $\boldsymbol{\mathcal{E}}$. Some examples of banding artefacts in the moment matrix are indicated with dotted lines in Figure 4.6b).

Due to the narrow width of the SOL, the variation in magnetic field and thus pitch angle across the radial extent of the inversion domain is rather small. Therefore, the geometry of the basis field lines varies slowly in the radial direction, with the consequence that there are a significant number of common pixels in radially adjacent basis field line images. The field lines in close radial proximity to a filament can therefore have raised emission along a portion of their length even if the field line as a whole does not properly align with the filament. This explains why many of the banding structures are predominantly radial in nature. We term these bands extending outward from the

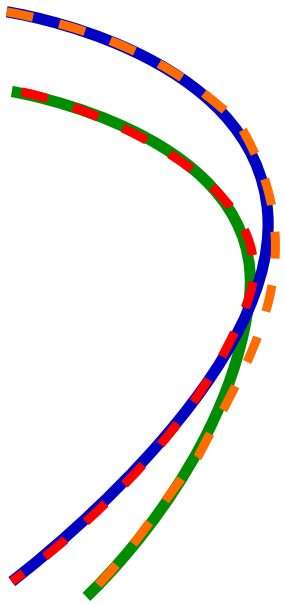

Figure 4.7: Illustration of the origin of 'ghost' filaments. The blue and green lines represent central field lines of true filaments that overlap in the line of sight of the camera. The red and orange dashed lines represent alternative magnetic field lines in the geometry matrix that do not correspond to the filaments, but that when projected onto the camera sensor, overlap with significant portions of the blue and green filaments (i.e. have many common pixels). Therefore, the red and orange dashed lines have high average intensities in camera images whenever filaments are present on the blue and green field lines. Thus, the red and orange dashed lines present as spurious 'ghost' filaments.

filaments the 'shadows' of the filaments, in the sense that the filaments 'cast' emission onto their neighbouring field lines.

While these shadows are heavily attenuated in the inverted emission, some remanent of them remains. Mathematically the shadows correspond to basis columns of the geometry matrix that are close to being degenerate, contributing to the poor conditioning of $\mathbb{G}$. The most direct consequence of the shadows is that they broaden the emission from the filaments, so that their widths will generally be overestimated. However, they also have a second important consequence, arising from the fact that some basis field lines can be weakly approximated by the linear combination of more spatially separated field lines. This is illustrated in Figure 4.7 where the flux tubes represented by the red and orange dashed lines overlap with significant portions of the blue and green flux tubes. Therefore, when filaments lie along the blue and green flux tubes, the average emission along the length of the red and orange flux tubes will also be high, despite there being little genuine localised emission along those field lines. This leads to secondary regions of enhanced emission in the inversion, that have the same appearance as filaments. These spurious filament structures we term 'ghost' filaments. In the emission plots these ghost filaments occur where the shadows of multiple parent filament overlap

to form new local maxima in intensity.

The ghost filaments typically occur at toroidal positions intermediate to the parent filaments, but at lower or higher radial positions, corresponding to positions in the emission plane where the shadows overlap, forming an X-like structure.

Bright spots in the camera images, typically occurring above the mid-plane where sections of field lines are nearly tangential to the camera field of view, can lead to further enlarging of a filament's shadow in both the radial and toroidal directions by raising the average intensities of all field lines that pass through the bright spot. Such filaments are particularly susceptible to producing ghost filaments.

The green filament in Fig. 4.6 is an example of a ghost filament arising from the overlapping of sections of the red and blue filaments in the camera line of sight.

4.2.6 Alternative approach to the camera analysis

In the moment matrix discussed until this point, the emission of each field line is given by the sum of the intensities along the field line and it therefore represents the average properties of the emission along the field line. However, a completely different approach to the analysis of the camera data is possible. Indeed, knowing how a given field line projects on the camera image, we can identify how the emission changes along it by tracking down the intensity on each pixel belonging to the said field line. This produces the possibility to investigate different statistical moments of the intensity distribution along the field lines and therefore acquire additional information. By doing this, we have dropped the assumption that the emission is homogeneous along the field lines, which was essential in the original inversion technique that we have discussed so far. If instead of taking the matrix product of $\mathbb{G}^T$ with $\mathbf{I}$, we take the element-wise product we produce a 'weighted geometry matrix',

$$\mathbb{F}_{ij} = \mathbb{G}_{ij} I_i, \tag{4.13}$$

where $\mathbb{F}_{ij}$ is the intensity of the i^{th} pixel in the camera image multiplied by the corresponding intensity of the i^{th} pixel in the basis image of the j^{th} field line. In Equation 4.13 we are suspending the use of the Einstein summation convention. The familiar moment matrix used above can be recovered by taking the sum over each column of F_{ij},

$$\mathcal{E}_j = \sum_i \mathbb{F}_{ij} \ . \tag{4.14}$$

The primary advantage of taking the sum of the pixel intensities along each field line is that the linear operation preserves the ability to perform the subsequent least squares optimisation to reduce the distortion of the structures in the moment matrix. The use of a linear field line intensity metric also has the benefit of preserving relative filament amplitudes and internal filament structure.

However, a range of 'parallel' operators acting on the columns of weighted geometry matrix, $\mathcal{F}_i(\mathbb{F})$, other than the sum $\mathcal{F}_i = \sum_i$ can quantify the emission along the field lines in different ways that also provide useful information for the analysis of field aligned plasma structures, as expressed by:

$$\xi_j = \mathcal{F}_i(\mathbb{F}_{ij}) \ . \tag{4.15}$$

Many quantitative measures of field line intensity can be conceived of and have been tested in order to most reliably identify filaments. Each have different strengths for constraining different filament properties, but are often accompanied with a range of weaknesses and limitations.

The geometric mean (G.M.), $(\prod_{i=1}^{n} x_i)^{\frac{1}{n}}$, which penalizes small values in the data series more effectively than the arithmetic mean (A.M.), $\frac{1}{n}\sum_{i=1}^{n} x_i$, is effective at differentiating field lines that consistently lie on top of filaments from those that do not. While it is not compatible with a subsequent least-squares optimisation due to the non-linearity of the operation, the G.M. is effective at identifying and measuring the position of filaments with far less computational expenditure. However, it is less appropriate for measuring filament widths and amplitudes. Of all the functional forms investigated, the geometric mean squared (G.M.^{2}) of the intensity along the field lines has been found to have the best all round properties for constraining the identification of filaments. However, other metrics such as the 2^{nd} and 98^{th} percentiles of the distribution of intensities along a field line, or parametrisations of smoothness, have potential for identifying ghost filaments. This is because they are sensitive to the short portions of field lines with particularly low intensities that typically occur where parent filaments intersect and the alignment of the shadow field lines to the parent structures is poorest.

These metrics are of particular interest as multiple channel inputs for convolutional neural network filament identification algorithms. These algorithms' performance can improve dramatically with a broader range of input information.

4.3 Filament identification

To extract filament properties from the inverted emissivity a blob detection algorithm based on the watershed algorithm described in Ref. 225 is used. The implementation of the algorithm used here is from the OpenCV library [195]. 2D Gaussians are fitted to the identified blobs in order to extract the position (R, ϕR), widths (δ_R, $\delta_{\phi R}$), amplitudes (A) and orientation (θ) of the filaments.

4.3.1 The watershed algorithm

The watershed algorithm treats the inverted emission data as a contoured terrain in which filaments appear as hills. The algorithm seeks to locate the the largest perimeter

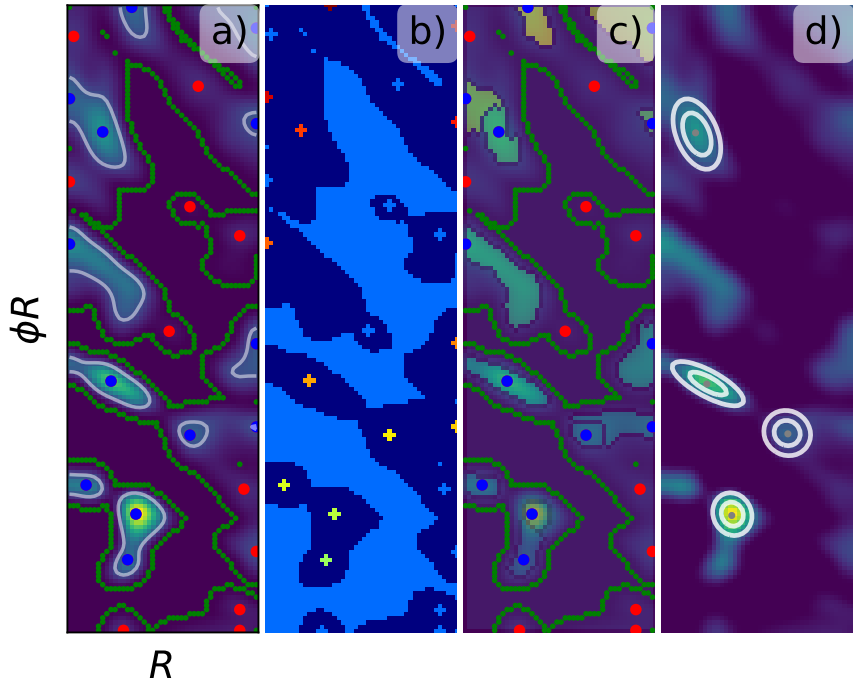

Figure 4.8: a) Emissivity data, with the minimum peak blob intensity threshold indicated as a white contour. Local maxima above and below the threshold are indicated as blue dots and red dots respectively, while local minima and boundaries of zero emission are indicated in green. b) Regions given as input to the watershed algorithm. Light blue is definite background (not a blob), dark blue areas are yet to be classified by the watershed algorithm and all other colours indicate foreground regions (sources for potential blobs). Note each pixel in the image represents a basis field line. c) Coloured regions show contours returned by the watershed algorithm. Coloured points are as in a). d) Contours of 2D Gaussians fitted to the data in the coloured regions shown in c). Some regions have been filtered out using a range of criteria and thus do not have fits.

of terrain that is unique to each individual hill, such that conjoined hills are treated separately.

Figure 4.8 illustrates the steps in the filament identification process using the watershed algorithm. The watershed algorithm requires the initial identification of regions of emission that are definitely inside filament structures (foreground regions) and definitely not inside filaments (background regions). To identify these regions, an empirical field line emissivity amplitude threshold is set, ϵ_{thresh}. This threshold is chosen in such a way that a significant proportion of local peaks in field line emissivity with amplitude below ϵ_{thresh} are found to originate from noise or remnant distortion effects. This threshold is indicated by the white contour in Figure 4.8a). Alternatively, the threshold can be set to a number of standard deviations above the mean inversion intensity across all the frames, akin to many past measurements in the literature [48, 226]. However, the first method has the benefit that all filaments above the noise floor of the technique can po-

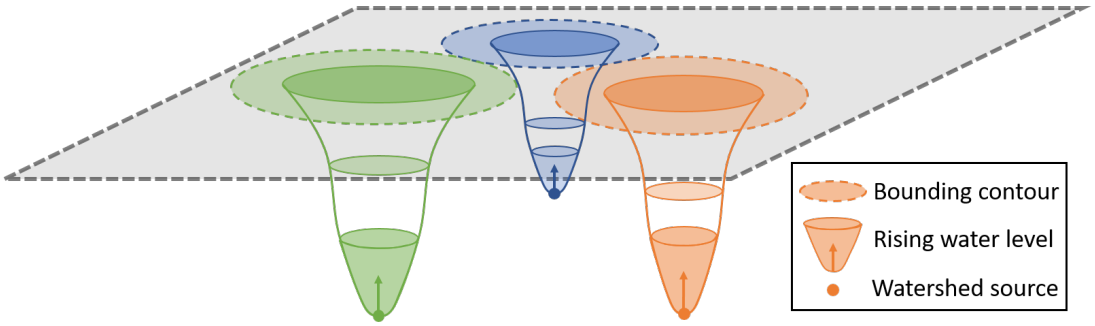

Figure 4.9: Illustration of the watershed algorithm, with different coloured water sources rising until they touch another region, at which point they define a bounding contour for a filament.

tentially be analysed, as opposed to imposing a somewhat less physically justifiable e.g. 2.5σ cut-off in the amplitude of the coherent filamentary structures that are analysed. Further details on the optimal selection of this parameter are given later in §5.4.1.

Clusters of high intensity field lines with a peak amplitude above this intensity threshold, are considered to be filament candidates. All 2D local minima and local maxima with intensities below ϵ_{thresh} are labelled as background (not filaments). These are indicated respectively as green and red dots in Figure 4.8a) and Figure 4.8c). A dilation operation is applied to these sets of background points, which expands the set of pixels to include their nearest neighbours. This dilated set of pixels, along with regions of zero emission, are labelled as the background region for the watershed algorithm, as indicated by the light blue regions in Figure 4.8b). Dilated regions around the 2D local maxima with intensities above ϵ_{thresh} are labelled as foreground (potential filament centres). These are shown as blue dots in Figures 4.8a) and 4.8c) and coloured regions in Figure 4.8b).

With the regions indicated in Figure 4.8b) as inputs, the watershed algorithm sorts the remaining unassigned regions into background [purple regions in Fig. 4.8c)] or distinct foreground regions [coloured regions in Fig. 4.8c)]. The watershed algorithm can be pictured in terms of different coloured pools of water originating at the distinct foreground markers. The water in each pool is allowed to rise and spread according the terrain of the data until it comes into contact with a background marker or water of a different colour. On contact with another region, the area assigned to that colour (filament) is frozen, while the other coloured regions continue to expand. This process is illustrated in Figure 4.9.

The contour regions identified for each filament are filtered according to a number of conditions. First, the contours are assessed on their *solidity*, defined as the ratio of the area of the contour to the area of its convex hull, which is a measure of how uniformly convex the contour is and thus how elliptical it is in shape. Contours with a low solidity (typically $< 60\%$), are rejected as their irregular shape indicates they cannot

be described well by a fitted 2D Gaussian and they may be the results of multiple interacting filaments which the watershed algorithm was unable to separate. Next, contours are rejected if they extend to the edge of the analysis grid and have significant intensities there (typically $> 50\%$ of the peak contour amplitude) as these contours extend outside of the analysis grid and thus are missing information required for an accurate Gaussian fit. Finally, any excessively large regions that are too large to be a filament ($\delta_R > 15$ cm, $\delta_{R\phi} > 20$ cm) are rejected. Note that Figure 4.8 displays the data in its true aspect ratio, so that the resolution of the field line basis grid can be readily seen.

4.3.2 Fitting 2D Gaussians

The contours passing the required section criteria are fitted with tilted 2D Gaussians described by

$$\epsilon(R,\phi R) = A\exp\left(-\left(a\Delta_R^2 + 2b\Delta_R\Delta_{\phi R} + c\Delta_{\phi R}^2\right)\right) + A_{bg} \tag{4.16}$$

where

$$a = \frac{\cos(\theta)^2}{2\sigma_R^2} + \frac{\sin(\theta)^2}{2\sigma_{\phi R}^2} \tag{4.17a}$$

$$b = -\frac{\sin(2\theta)}{4\sigma_R^2} + \frac{\sin(2\theta)}{4\sigma_{\phi R}^2} \tag{4.17b}$$

$$c = \frac{\sin(\theta)^2}{2\sigma_R^2} + \frac{\cos(\theta)^2}{2\sigma_{\phi R}^2} \tag{4.17c}$$

$$\Delta_R = R - R_0 \tag{4.17d}$$

$$\Delta_{\phi R} = \phi R - \phi_0 R_0. \tag{4.17e}$$

Equation 4.16 describes a 2D Gaussian with amplitude, A, centred at $(R,\phi R) = (R_0,\phi_0 R_0)$ (note here ϕ_0 is the toroidal angle of the centre of the filament, not the centre of the inversion grid), with respective radial and toroidal widths σ_R and $\sigma_{\phi R}$, tilted at an angle θ anti-clockwise with respect to the (R, ϕR) axes. The background intensity level, A_{bg} is set to the average intensity of the local minima in the inversion domain, which is typically very close to zero. Thus the Gaussian fits yield the filaments' central positions, 1σ major and minor axes widths, amplitudes and inclinations. The inclinations of the filaments are small and are heavily influenced by the shadows of the filaments and so are currently treated as fitting parameters, rather than physical parameters of interest. Detected filaments are indicated in Figure 4.8d) by white ellipses showing the fitted 1σ and 2σ contours of the filaments.

A Gaussian description of the filaments is chosen because it is the simplest fitting function that produces a representative width and amplitude for the observed blob-like data. The Gaussian description is also consistent with the parametrisation of filament cross sections used in a variety of numerical and analytic filament studies [99, 115, 169]. The combination of line of sight integration effects and limited spatial resolution mean it is not possible to reliably discern any more complex internal filament structure.

4.3.3 Discussion of the identification algorithm

The watershed algorithm has a number of important advantages over previously tested algorithms that relied upon contours generated from thresholding the intensity data at a suitable level identified for each peak. Firstly, it handles interacting filaments very effectively, preventing them from being merged into a single larger blob and avoiding problems from multi-modal emissivity structures producing multiple contours for a given filament. Secondly, it efficiently maximises the amount of relevant intensity data assigned to each filament, with each filament assigned to the largest possible surrounding region, out to the closest local minima in intensity or the closest neighbouring filament (whichever is closer). When combined with the fitting of 2D Gaussians to the contours, this means that large secluded filaments can utilise intensity information spanning the whole structure, leading to particularity high accuracy fits. In the case of small, interacting filaments, this makes the detection of the filaments possible by relying on the upper sections of the amplitude distributions, which would not have been correctly identifiable by the previously tested algorithms, due to the absence of the required contour level (e.g. 50%) within the filament's enclosing contour.

In Fig. 4.6c) the detected filaments were indicated by coloured ellipses showing the 1σ contours of the fitted Gaussians. The corresponding central field lines for each fitted filament are indicated by coloured lines of the same colour in Fig. 4.6a) and 4.6d).

4.4 Filament tracking

4.4.1 Introduction

As discussed in Chapter 2, the radial velocity of the filaments is an important quantity in determining the transport of particles in the SOL. It is therefore desirable to perform direct measurements of the filaments dynamics in order to validate theory and better inform the choice of inputs to the analytic framework.

While a great deal of attention in the literature has been paid to understanding the radial motion of filaments, relatively little is understood about their toroidal motion. This is a result of several factors. Firstly, the toroidal motion of the filaments is typically of less interest than the radial motion, as the toroidal motion has been shown to not influence the radial cross-field transport of particles [168]. Secondly, the main driver of drifts at the outboard mid-plane are bad-curvature effects, that are directed radially rather than toroidally. Therefore, toroidal dynamics must arise from weaker secondary effects, or may be associated with flows in the background plasma. Finally, it is often difficult to perform accurate experimental measurements of the bulk toroidal plasma rotation, which makes it challenging to produce reliable measurements of the intrinsic toroidal rotation of the filaments in order to inform or verify theoretical models. Bulk plasma rotation occurs due to the addition of momentum imparted by neutral beams or

radio-wave heating [16], in addition to intrinsic rotation driven by turbulent momentum transport due to up-down asymmetries [227] and the edge radial electric field.

With the positions of the filaments in each frame identified with the detection technique described in the previous section, we now wish to identify which detections are multiple observations of the same filament over its lifetime.

Given the camera records a frame every 10 μs and filament lifetimes are of the order of a few tens of microseconds, filaments are only observed for a few frames. Therefore, the time resolution of our measurements are only marginally sufficient to resolve the motion of the filaments and many filaments with short lifetimes will only be observed once, so that no assessment of their velocity can be made. However, a sufficient proportion of the filaments persist for multiple frames that valuable velocity information can be extracted.

Object, particle or 'blob' tracking is performed in a range of disciplines, in particular computer vision [228] and measurements of fluids [229]. However, as shall become apparent, the problem of tracking filaments is particularly challenging, owing to the low time resolution of the data with respect to the phenomena being observed and the lack of strong assumptions that can be made about the object being tracked.

4.4.2 Overview of assumptions

Many tracking algorithms applied in the field of computer vision deal with tracking of rigid and articulated objects [230] which assume distinctive features are preserved between frames. These techniques often use the spatial gradients of the images to identify displacements of edge features between frames. However, filaments are diffuse objects with no well defined edges or features.

The very limited number of time point measurements also means there is typically insufficient information (often only one or two points) available to form an accurate prediction of a blob's future position and state. This makes it difficult to apply kernel based techniques such as Kalman tracking [228].

A particularly simple and effective method of blob tracking is to simply match objects that overlap in consecutive frames. Such an approach is for instance applicable in the tracking filaments in turbulence simulations [231], where the time resolution of the simulation can be set sufficiently high that filaments' motions cover short distances with respect to their widths over one time step. Similarly, in many other imaging applications looking at particle motion in fluids [232], the motions of the objects that must be tracked are sufficiently slow that the camera frame rate can easily be increased to ensure the overlap of objects between consecutive frames. However, in the case of filaments, their velocities are sufficiently extreme ($\sim$km/s) that the frame rate is limited by the hardware capabilities of the fast camera. Therefore, the frame rate of the measurements is insufficient to ensure overlap of the filaments between consecutive frames. Figure

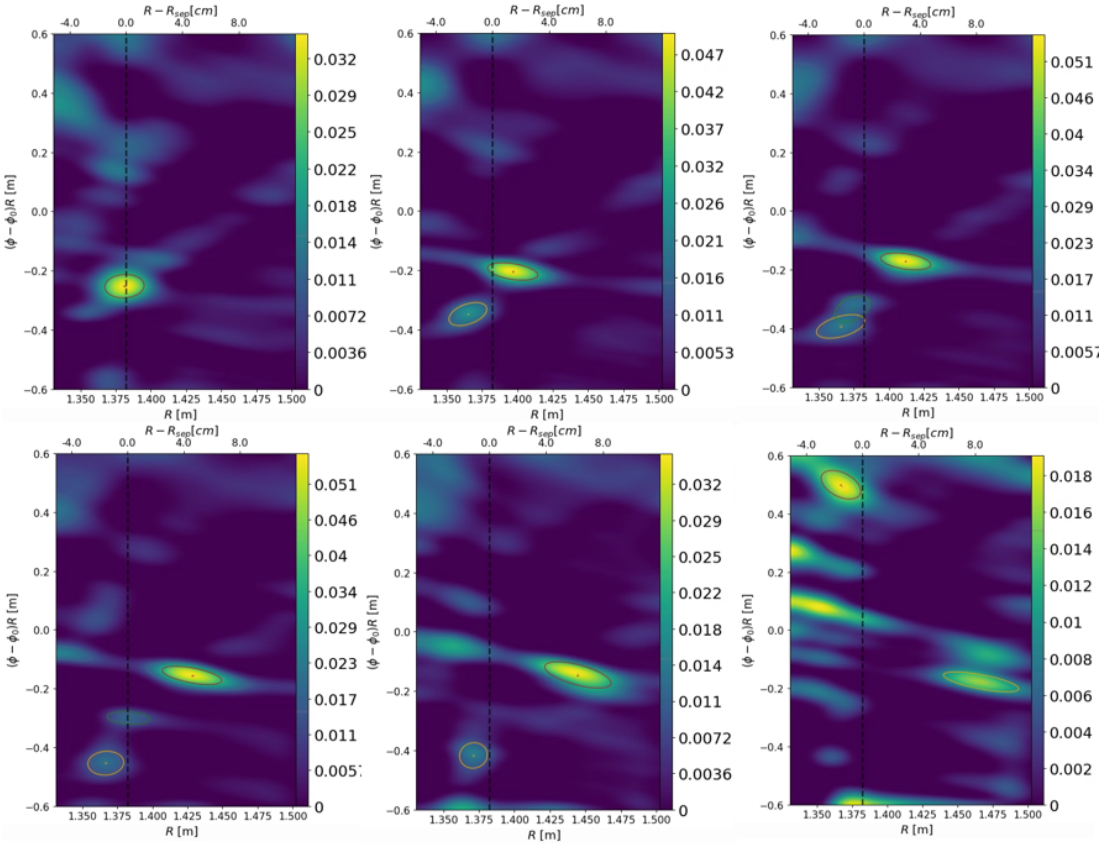

Figure 4.10: Series of camera inversions from MAST discharge 29852 over the interval $t = 0.21641$-0.21647 s, progressing in time from left to right, top to bottom. A large filament highlighted with an ellipse in each inversion is seen detaching from the separatrix and travelling radially across the SOL. The separatrix is shown by a black dashed line.

4.10 shows a series of consecutive camera inversions illustrating the displacement of filaments between frames. Therefore, common particle image velocimetry (PIV) [232] and optical flow [229, 233, 234] based techniques are not applicable to the tracking problem encountered here.

This application is also made challenging by the further lack of many conserved quantities that are useful in other applications. In particular, conservation of filament number is not applicable, as filaments are born and die over the course of the measurements. Furthermore, if a filament persists in a following frame but its intensity drops below the required threshold for detection, ϵ_{thresh}, it is not recorded. Size and shape are also not conserved, as filaments change shape over their lifetimes, both due to internal restructuring of the filaments and distortion effects arising from the inversion as discussed previously. In general their widths should increase with time as the filament density spreads across field lines and the vorticity of the filament breaks it into lobes. However, the drainage of the filaments parallel to the magnetic field will tend to lower the amplitude of the filaments over time, potentially decreasing the size of their foot-

print in the inversions. Conservation of intensity (or colour) is also not applicable here. Filaments' amplitudes should be expected to in general decrease with time. However, variation in the neutral density close to the separatrix or internal restructuring of the filament both have the potential to produce rising amplitudes in some cases.

As a result of the lack of conserved quantities and the limited *a priori* information available to constrain the system owing to the sparse experimental and theoretic understanding of filament birth and death in the literature (see discussion in §2.3.1), a relatively simple tracking algorithm has been employed, based on the radial and toroidal displacements of the filaments between frames.

The technique makes use of the following assumptions, in large part justified *a posteriori* by experimental observations with the fast cameras:

1. The radial motion of a filament is primarily driven by the $\mathbf{E} \times \mathbf{B}$ drive mechanism discussed in Chapter 2, so that the dominant radial motion of the filaments is radially outwards.

2. Blob splitting and mergers are rare events so that a filament only ever appears as one 'blob' in the 2D plot of a given inversion.

3. Filaments are born at low major radii in the vicinity of the separatrix (i.e. near the left hand side of the images), but can die anywhere depending on the rate of parallel losses.

   (a) Therefore filaments detected at large major radii (i.e. towards the right hand side of the images) must have propagated there from lower major radii (i.e. have previously been further to the left in previous images).

   (b) Filaments that are observed for the first time at large major radius are a consequence of missed detections in previous frames and are rare compared to the many filaments born at low major radii.

4. The primary component of a filament's toroidal velocity in the lab frame is due to the bulk toroidal rotation of the plasma.

   (a) The bulk plasma rotation is approximately constant over a filament's lifetime.

   (b) The bulk plasma rotation is slowly varying across the SOL.

As a result of assumption (1), the toroidal motion of all filaments is in the direction of the bulk plasma rotation and the variation in the toroidal velocities is small, both between filaments and over a filament's lifetime. This means that the toroidal motion of the filaments is relatively simple, so that the range of possible locations for each filament in the following frame can be estimated.

Assumption (2) provides a similar simplifying constraint on the radial velocity so that range of possible radial locations in the subsequent frame can be estimated. Assumption (3) significantly simplifies the implementation of the tracking algorithm by only allowing one pair of filaments to be matched between frames.

4.4.3 The filament tracking algorithm

With the basic assumptions of the technique established, we now describe the implementation of the technique. The filament tracking algorithm proceeds as follows:

1. For each filament with index i in frame n, the radial displacement ΔR_{ij} and toroidal displacement $\Delta(\phi R)_{ij}$ from each filament of index j in frame $n+1$ is calculated.

2. If both the radial and toroidal displacements are consistent with filament motion adhering to the assumptions discussed above, the jth filament in frame $n+1$ is assigned to the set of preceding filament detections (of which the ith filament in frame n is the previous member).

3. If more than one filament in frame $n+1$ meets the matching criteria, the following secondary criteria are used to select a single match:

   (a) The average toroidal velocity of all other identified chains of filaments is calculated for a window of 30 frames around the frame of interest. The candidate filament is selected that corresponds to the toroidal velocity that is most similar to its neighbours and consistent with the previous elements of the chain (given their toroidal motion is dominated by the same bulk plasma rotation).

   (b) If no clear candidate is found from condition (a), the candidate is selected that has an amplitude most consistent with the preceding observations of the filament.

4. Steps 1-3 are repeated for all frames, so that all filament detections are uniquely assigned to a set of one or more detections corresponding to the same filament.

Steps 1 and 2 of this procedure are illustrated in Figure 4.11, while Figure 4.12 illustrates step (3a), in which more than one candidate is found for the continuation of a filament.

The bounds on the allowed displacements in step (2) are based on past experimental observations [51] that show L-mode filaments in MAST have toroidal velocities in the range 1-9 km/s and radial velocities in the range $0 - 2.5$ km/s. Given the interval between frames is 10 μs and allowing for errors in the calculated displacements of twice the grid resolution, the allowed toroidal displacement of the filaments between frames is

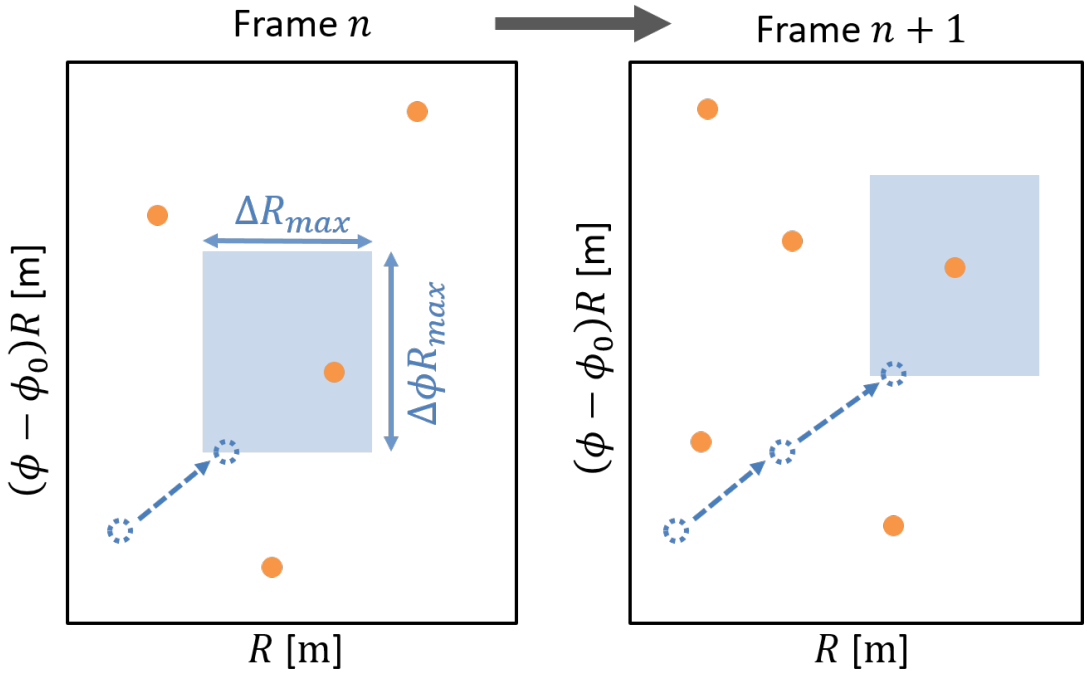

Figure 4.11: Illustration of the tracking algorithm. The black rectangles represent the domain of the inverted emission for two consecutive frames. Orange dots represent filaments detections in the current frame. Blue open dashed circles represent chains of detections from preceding frames. Blue shaded regions represent the bounds in which a filament must fall to be matched to the previous filament in the chain, based on possible filament velocities.

$8-92$ cm and their permitted radial displacements are -1−25 cm. Thus, small negative radial displacements are allowed.

While steps 1 and 2 are automated and do not require user intervention, step 3 is currently performed manually. This involves comparing the multiple match candidate using the criteria in steps 3a) and 3b) and deleting the incorrect/duplicate chain indices. However, situations where more than one candidate match the tracking criteria are found to be rare, occurring in around $\lesssim 0.5\%$ of frames, so little user intervention is required.

The near complete automation of the tracking procedure has two key benefits. Firstly, it leads to a tremendous increase in speed of the procedure and thus the number of filaments whose dynamic properties can be studied. This leads to larger sample sizes and improved measurement statistics. Secondly, it provides a robust, reproducible and self-consistent means of tracking the filaments, free of the subjectivity and personal biases inherent in challenging manual interpretation of the filament data.

Visual inspection of the tracked filaments in the camera images indicates the algorithm does a good job of identifying chains of detections that would be identified manually. With the time evolution of the filaments identified, the average velocities of

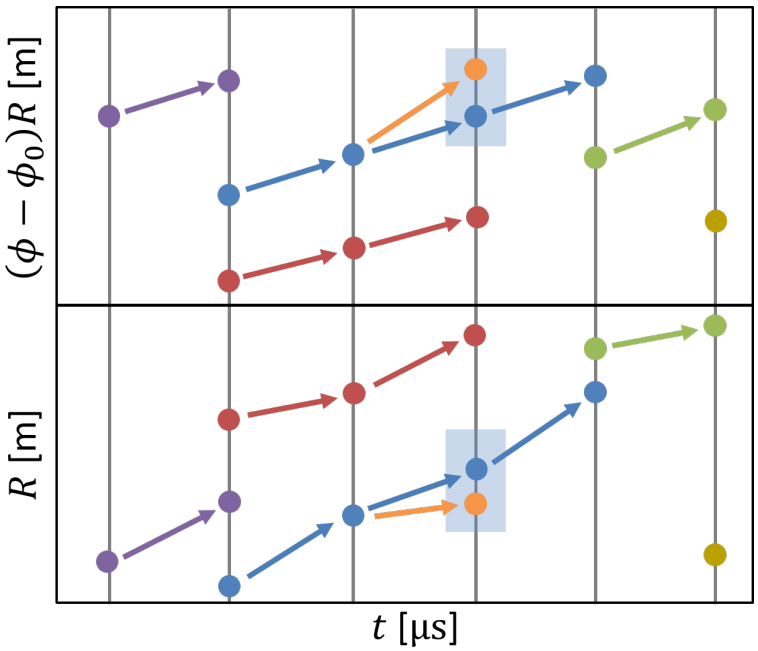

Figure 4.12: Illustration of multiple candidate filaments meeting the tracking criteria. In this situation the toroidal motion of each option is compared to that of the preceding and neighbouring chains. The candidate is picked that produces the most consistent toroidal motion.

the filaments between each set of observations is readily calculated.

Figure 4.13 shows an example application of the tracking algorithm to experimental data from MAST discharge 29852. Each chain of detections assigned to the evolutional of a single filament is coloured separately. Note that distinct chains of filament detections that have very similar radial positions are well separated toroidally, so that they do not overlap in real space.

4.5 Implementation of the Elzar code

4.5.1 Introduction

The camera inversion and filament identification and tracking techniques described in the previous sections has been implemented in an analysis code named *Elzar*. The code is written in `python` and can be run through a combination of script, command line and graphical user interface (GUI) tools. The code has been developed with application to analysis of large volumes data in mind. The code can be run in parallel or segmented serial batch modes. It is not currently open source, so requests for access to the code should be directed to the author[†].

[†]External readers may contact the author at tom.farley@ukaea.uk to request a copy of the code. Readers internal to CCFE can access the code at https://git.ccfe.ac.uk/tfarley/elzar

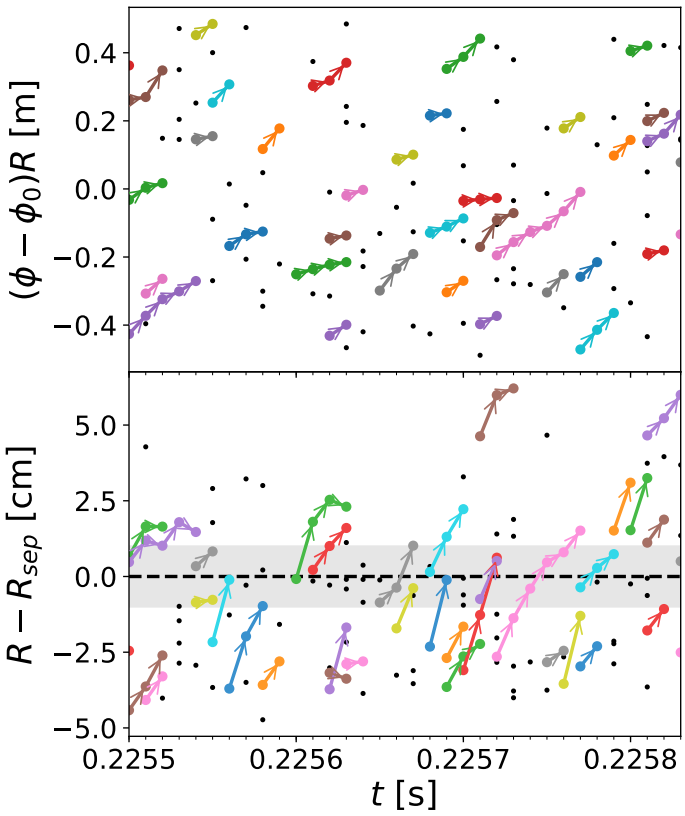

Figure 4.13: Toroidal (top) and radial (bottom) locations of filament detections as a function of time in MAST discharge 29852. Filament detections assigned to the same chain are given the same colour. Arrows indicate the sequence of the detections. Isolated filaments that are only observed for one frame are shown in black. The position of the separatrix and it's uncertainty are shown by the black dashed line and shaded region respectively.

4.5.2 Code dependencies

The code makes use of a number of widely used python packages. In particular `numpy`, `xarray` and `pandas` are used extensively for the storage, manipulation and representation of data structures. The `openCV` graphics library is used for a range of image specific operations, in particular the image pre-processing steps and the watershed algorithm. The `matplotlib` package is used to visualise data and the `pyQt` library is used to produce the GUI.

In addition to these common packages dependencies, the code makes use of several other packages developed at CCFE. The `pyIpx` library is used to load the movie data in the MAST specific `ipx` file format. The `Calcam` library is used to perform camera spatial calibrations and the `pyEquilibrium` packaged is used to load and manipulate magnetic equilibrium data stored in the `gfile` file format.

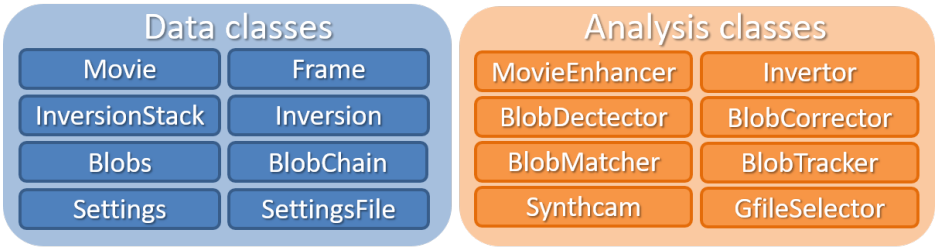

Figure 4.14: Summary of important classes used in the Elzar code. Data classes hold and manipulate different forms of data. Analysis class operate on the data classes to output new data.

4.5.3 Elzar analysis workflow

The Elzar code is written in a modular, object orientated fashion in order to facilitate the swapping in and out of different analysis components. Figure 4.14 lists some of the most important classes used in the code. The data classes are used to hold and manipulate the key data structures encountered in the workflow, while the analysis classes implement operations that typically convert data from one format to another.

The Elzar analysis workflow is summarised in Figure 4.15. Blue boxes describe the primary data input at each step of the analysis. Orange arrows and ellipses describe the important analysis steps. The processed data can be saved and loaded at each of these intermediate steps.

Each class has an associated settings file describing the various input parameters and paths that it requires. The settings from the files are collated into a `Settings` object allowing centralised editing and recording of the analysis settings. All output files are saved with a hash id uniquely tying the data to the specific set of settings used to produce it. In addition to recording the settings for data provenance purposes, this system enables pre-processed data at each step of the analysis to be readily located and loaded if it already exists, or calculated and saved if it is likely to be reused. This makes managing a large database of analysed data much more practical.

The raw movie data is read from file using the `Movie` class. When loading data from `ipx` or `mraw` movie files, or loading movies from stacks of `jpg` files, the CCFE `pyIpx` class is utilised. When synthetic movie data is used (see Chapter 5) it is loaded from `netcdf`, `pickle` or `numpy-zipped` file formats. A range of image enhancement steps can be applied using the `MovieEnhancer` class, most importantly the background subtraction step.

The inversion of the data is performed using children of the `Invertor` base class. Typically the `InvertorSartLapReg` class is used to invert the camera data using the C-SART algorithm with Laplacian regularisation. However, other classes such as the `InvertorMomentMatrix` class can be used to process the data using alternative methods. The `Invertor` class is also used to generate the geometry matrix for a given magnetic

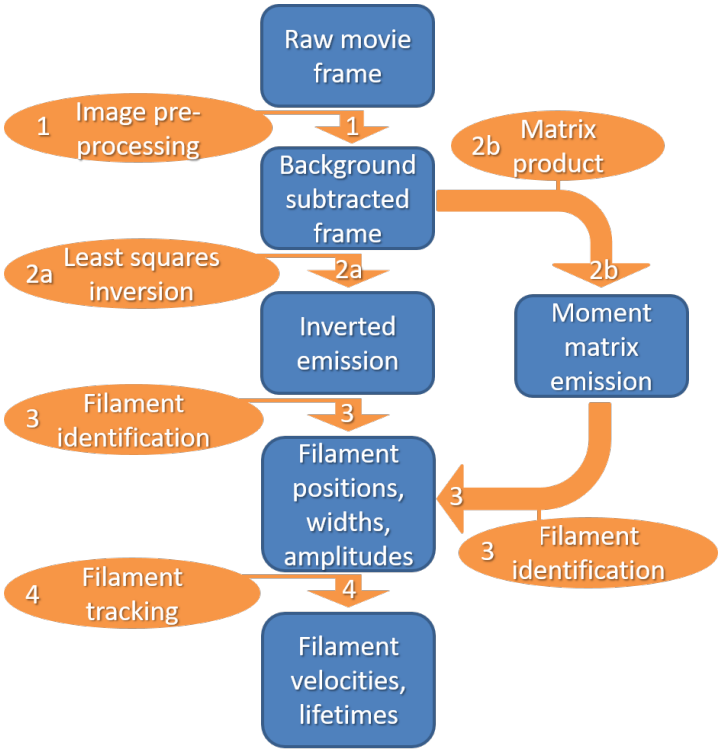

Figure 4.15: Block diagram of the Elzar analysis workflow. Blue boxes describe the primary data input at each step of the analysis. Orange arrows and ellipses describe the important analysis steps.

equilibrium, camera calibration and inversion grid resolution settings. The inverted camera data is stored in an `InversionStack` object that is similar to the `Movie` class.

Children of the `BlobDetector` class take the inverted emission and return a `Blobs` object containing the parameters of all the detected filaments. As with the `Invertor` class, different child classes can be used at this stage to apply different detection algorithms.

Finally, the `BlockTracker` class takes the detected filaments and identifies chains of detections, from which filament velocities and lifetimes are calculated.

A range of other classes also exist for processing, treating and visualising the data at various stages of the analysis. Of particular importance is the `BlobMatcher` class which is used to match synthetic and detected filaments in the benchmarking exercise described in Chapter 5.

4.5.4 The Elzar GUI

The Elzar code has a graphical user interface (GUI), implemented using `pyQt`, to aid analysis and interpretation of the data. The GUI shown in Figure 4.16 facilitates the easy use, with functions including loading and playing movie files, applying image enhancements and performing and viewing camera inversions. A particularly powerful

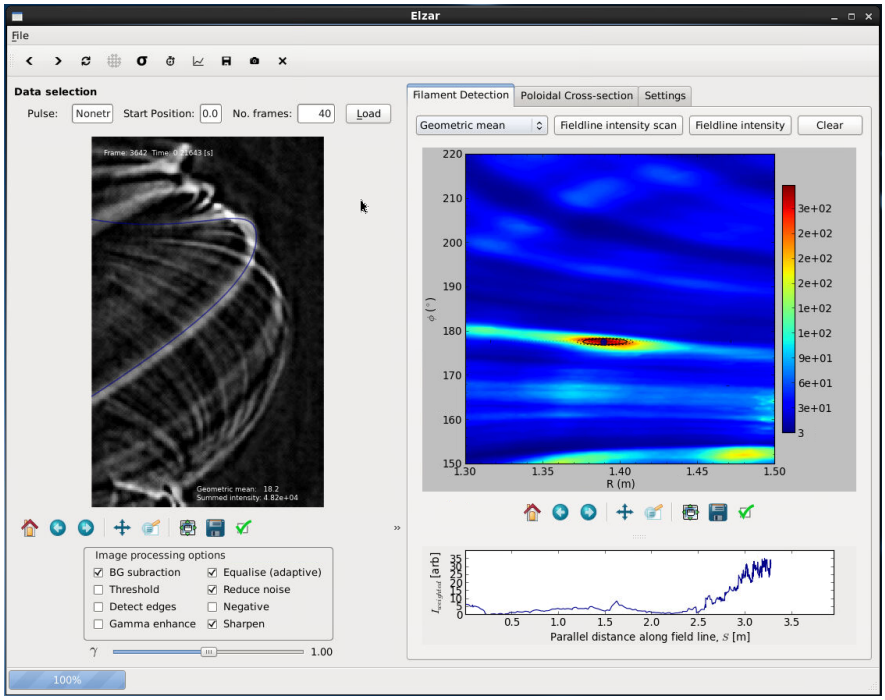

Figure 4.16: The Elzar graphical user interface. The enhanced movie is shown on the left, while the corresponding inversion is seen on the right.

feature is the ability to click on regions of the inverted data to project the corresponding field lines onto the camera image, in order to gain grater insight into the origin of structures in the inversion.

4.6 Conclusions

The tomographic inversion and filament detection technique implemented in the `Elzar` code discussed in this chapter is able to successfully identify filamentary field aligned structures in visible fast camera images and measure their positions, widths, amplitudes and velocities. However, a number of complicating factors arise from the fact that upon projection onto the camera's 2D field of view, many basis field lines are close to being linearly dependent. This results in multiple viable inverted emissivity solutions for the same camera image, leading to spurious false detections from 'ghost' filaments. Therefore, it is important to benchmark the technique thoroughly and assess its accuracy, limitations and errors. This is addressed in the next chapter.

Chapter 5

Benchmarking of the Elzar technique

5.1 Introduction

As discussed in the previous chapter, the results of the inversion process are subject to artefacts that can broaden the footprint of filaments and produce spurious 'ghost' filament detections. Therefore, it is of great importance to thoroughly benchmark the performance of the Elzar camera inversion and filament detection algorithm, so that experimental results can be correctly interpreted and strong conclusions drawn.

Ideally it would be possible to benchmark the diagnostic through a direct comparison with another well established diagnostic that is capable of performing the same measurements. However, due to the unique capabilities of this direct visible imaging technique, in particular its large spatial coverage, such a diagnostic is not available. The next most informative diagnostic on MAST for filament measurements was the reciprocating probe diagnostic. However, due to the reduced dimensionality of the Langmuir probe measurements with respect to that of the camera measurements, only a statistical comparison can be performed, which is the subject of §6.1. In the absence of an alternative source of trusted experimental measurements of filaments with which to perform a direct comparison, the most insight into the performance of the technique can be gained with artificial camera data containing filaments with known, well defined properties. A synthetic camera diagnostic has been developed to produce artificial camera data for this benchmarking exercise. This chapter discusses the application of the synthetic diagnostic data to characterise the technique and quantify its errors.

5.2 Synthetic camera diagnostic

A synthetic camera diagnostic has been created to aid in the development and benchmarking of the identification technique, as well as facilitating the forward modelling of simulation outputs for comparison with experimental measurements (as applied in [99]

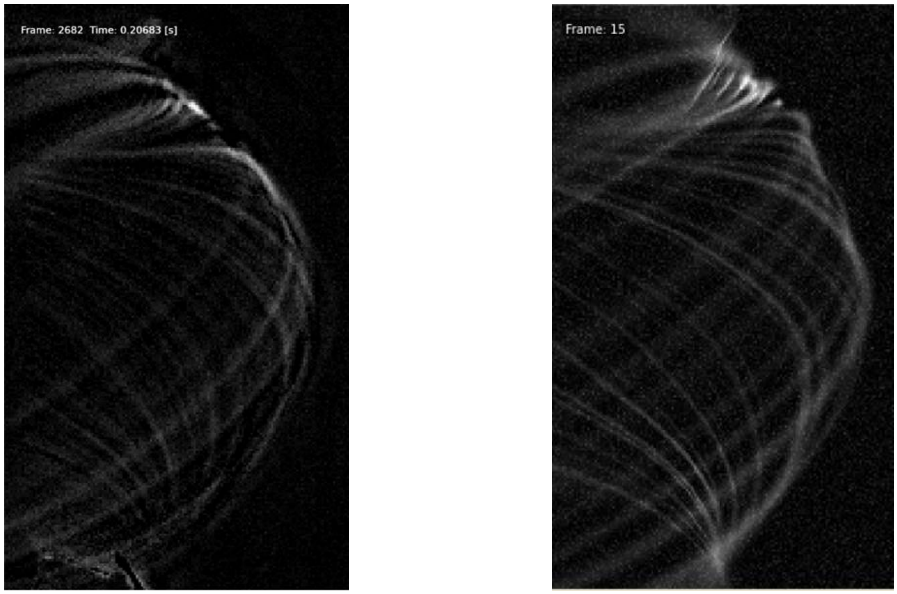

(a) (b)

Figure 5.1: a) Experimental camera frame from MAST shot 29852, with the preprocessing described in Section 4.2.2 applied. b) Frame from the reference synthetic data set produced using the synthetic diagnostic, containing 40 synthetic filaments, with positive definite Gaussian noise applied.

and discussed in Chapter 6). The principle used is the converse of the inversion process discussed in the previous chapter: Equation 4.8 is applied to produce a synthetic camera image $\mathbf{I}_{synth}$ for a predefined emission profile $\boldsymbol{\epsilon}_{synth}$ and geometry matrix $\mathbb{G}$. The resulting image contains emission that is equivalent to that in background subtracted experimental camera frames.

In principle, the input emission vector can be any 2D pattern. If an experimentally derived emission vector is used as input, we produce a synthetic re-projection of the inversion data, demonstrating the visual form of the structures that are captured in the emission vector, as was shown in Figure 4.6d). The input emission vector can also be provided by numerical simulations [99] and in this application, the technique provides forward modelling. A third option is when entirely artificial filament images are produced for the purposes of benchmarking and technique optimisation and each filament's emission profile at the mid-plane is set to a 2D Gaussian.

In the latter case, images of synthetic filaments can be combined additively to produce full synthetic frames. In order to closely match experimental data, positive definite Gaussian noise can be added to the frames to simulate noise from the camera sensor. Random noise of this type is added to all synthetic frames discussed in this chapter.

As the synthetic diagnostic composes images through a linear combination of field line basis images, it can only produce images of field-aligned structures and no variations in emission along field lines is possible in its current form. This limitation could be

relaxed by including further basis functions of increasing complexity in $\mathbb{G}$, although for the reasons outlined in §4.2.1, the emission of interest is strongly field aligned and such an increase in complexity is not warranted at this point.

Figure 5.1 compares output from the synthetic diagnostic to experimental data, showing a good visual qualitative similarity, further justifying this choice. Some differences are seen in the upper and lower divertor regions due to poorer constraint of the camera calibration in those regions (as discussed in §3.3). However, they are outside of the analysis region that is inverted and so do not impact the results below.

5.2.1 Synthetic Datasets

A reference data set, A, of 5000 synthetic fast camera frames was generated for this study with 40 synthetic filaments per frame, corresponding to a total of 200,000 filaments with random positions and characteristics. These synthetic filaments were distributed uniformly around the full toroidal extent of the machine. A range of numbers of filaments per frame was explored and 40 filaments per frame were chosen as it corresponds to a commonly observed toroidal quasi-mode number seen in MAST [51] and produced synthetic frames which most closely visually resembled experimental camera data. The toroidal quasi-mode number describes the average number of filaments observed in a full toroidal circuit of the machine. The chosen value (40) lies at the upper end of literature values for quasi-toroidal mode number [47, 51]. This is appropriate, given that literature values are likely to have missed the very faintest filaments which are challenging to measure, but which are important to include in the synthetic data. Furthermore, the potential overlap of filaments means the apparent quasi-mode number is likely to appear lower than the input mode number.

Positive Gaussian noise with standard deviation equal 5% of the image dynamic range was applied to all synthetic frames. This amplitude of noise is representative of that found in experimental camera data.

For this application, the geometry matrix is calculated with a higher grid resolution of $\Delta R = \Delta(\phi R) = 2.5$ mm in order to avoid aliasing effects in the resulting images.

The synthetic filaments are produced with Gaussian cross-sections and their properties (positions, widths and amplitudes) are drawn from realistic probability distributions based on experimental findings. The radial distribution of the filaments follows a log-normal distribution [84, 87] starting at the separatrix ($R = 1.36$ m) and peaking 3 cm into the SOL. As shall be discussed in Chapter 6, analysis of experimental camera data with the Elzar technique has shown an exponential distribution of toroidal filament separations, indicating that filaments are randomly distributed toroidally around the machine with no mode number (hence the use of the term quasi-mode number to refer to the average number of filaments). Therefore, a uniform distribution has been adopted for the toroidal locations. The radial and toroidal widths are drawn from log-

Table 5.1: Distribution information for filament parameters in synthetic data sets A and A_1. The uniform distribution is represented by $\mathcal{U}$, the log-normal distribution by $\mathcal{L}(x;\sigma,\mu) = \frac{1}{\sigma x\sqrt{2\pi}}\exp\left(-\frac{(\ln(x)-\mu)^2}{2\sigma^2}\right)$ and the exponential distribution by $\mathcal{E}(x;\beta) = \frac{1}{\beta}\exp\left(-\frac{x}{\beta}\right)$. Due to the parametrisation of the field lines in ϕR space, the inversion domain is slightly wedge shaped. To account for this filaments are not generated in a 5° toroidal wedge. The centre of the distribution is set such that this wedge is on the far side of the machine where the field lines are not in the camera field of view.

Property		Distribution	Parameters	Refs.
Position	R	$\lambda\mathcal{L}+\nu$	$\mu=0$, $\sigma=0.5$ $\lambda=0.04$, $\nu=1.36$	[84, 87]
	ϕ	$\lambda\mathcal{U}+\nu$	$\lambda=355$, $\nu=2$	[168]
Width	δ_R	$\lambda\mathcal{L}$	$\mu=0$, $\sigma=0.4$ $\lambda=0.01$	[51, 97]
	$\delta_{\phi R}$	$\lambda\mathcal{L}$	$\mu=0$, $\sigma=0.4$ $\lambda=0.8$	[51, 97]
Amplitude	A	$\mathcal{E}$	$\mu=0$, $\sigma=0$	[61, 68, 90]

normal distributions. Due to the difficulty of experimental measurements of filament dimensions there is less information in the literature to inform the choice of a specific distribution for the width. However, width measurements on MAST [51, 97] and NSTX [84] show the width distributions to be asymmetric and positively skewed. The log-normal distribution commonly describes many small finite positive definite quantities such as the width [235]. This choice of distribution is further justified by experimental width measurements in Chapter 6 being well described by log-normal distributions. Filament amplitudes in the synthetic data are distributed exponentially as commonly observed [61, 68, 90].

The specific parameters of the distributions are given in Table 5.1. In this and subsequent tables the uniform distribution is represented by $\mathcal{U}$, the log-normal distribution by

$$\mathcal{L}(x;\sigma,\mu) = \frac{1}{\sigma x\sqrt{2\pi}}\exp\left(-\frac{(\ln(x)-\mu)^2}{2\sigma^2}\right), \tag{5.1}$$

the exponential distribution by

$$\mathcal{E}(x;\beta) = \frac{1}{\beta}\exp\left(-\frac{x}{\beta}\right). \tag{5.2}$$

The parameters μ and σ describe the means and widths of the distributions, while λ and ν are used as scaling and offset parameters respectively.

The camera calibration and magnetic equilibrium used to produce the synthetic data are for MAST discharge 29852. The radial extent of the analysis region was $R = 1.290$-1.474 m, spanning the outer mid-plane SOL and including several centimetres inside the separatrix, given the MAST mid-plane separatrix was located at $R \approx 1.36$ m for this discharge. The analysis region extended toroidally from $\phi R = -0.55$ m to $\phi R =$

Quantity	Value
Minimum major radius, R_{min} [m]	1.290
Maximum major radius, R_{max} [m]	1.474
Radial extent, $R_{max} - R_{min}$ [m]	0.184
Radial separatrix position, R_{sep} [m]	1.36
Radial grid resolution, dR [m]	0.004
Minimum toroidal position, $(\phi R)_{min}$ [m]	-0.550
Maximum toroidal position, $(\phi R)_{max}$ [m]	+0.548
Toroidal extent, $(\phi R)_{max} - (\phi R)_{min}$ [m]	1.098
Central toroidal angle, ϕ_0 [°]	177
Toroidal grid resolution, $d(\phi R)$ [m]	0.009
Points in radial dimension, N_R	47
Points in toroidal dimension, $N_{\phi R}$	123
Total number of grid points, $N_R \times N_{\phi R}$	5781
Area in toroidal plane, A_{grid} [m^2]	~ 0.20
Reference discharge	29852

Table 5.2: Resolution of the inverted emission grid used for algorithmic benchmarking with synthetic data.

+0.548 m, corresponding to an angular range in machine coordinates of $\phi = 153°$-$215°$ centred about $\phi_0 = 177°$ (c.f. camera tangency plane at $\phi = 215°$).

The analysis region therefore focuses on foreground filaments (i.e. those that intersect the mid-plane in the foreground). Foreground filaments are chosen due to their proximity to the camera increasing their size in the camera field of view and thus maximising the accuracy of their width measurements, as discussed in the previous chapter. Of the 200,000 filaments in the data set, 20,243 (10.1%) fall within the analysis domain indicated by the white dashed box in Figure 4.6c), corresponding to an average of 4.05 filaments in the analysis domain per frame.

Toroidal positions and widths are generated as angles (ϕ and δ_ϕ) and converted to lengths ($R\phi$ and $\delta_{R\phi}$). The analysis grid had a resolution of 4 mm in R and 9 mm in $R\phi$, corresponding to 47 radial points and 123 toroidal points, giving a total of 5,781 distinct field lines. This is the same resolution used to invert experimental camera data. This resolution was chosen so as to best constrain the radial and toroidal widths of the filaments, while keeping the computational burden manageable. Table 5.2 summarises the parameters of the inversion domain.

Given the decaying exponential amplitude distribution peaks at zero, many of the synthetic filaments are very faint, as may also be true in the experimental case, and therefore cannot reasonably be expected to be detected against background noise and inversion artefacts. However, inclusion of these very faint background filaments, while not currently a focus of measurement, is important for reproducing the complexity of the likely distributions of the experimental data.

As the sensitivity of camera sensor is not absolutely calibrated, the absolute ampli-

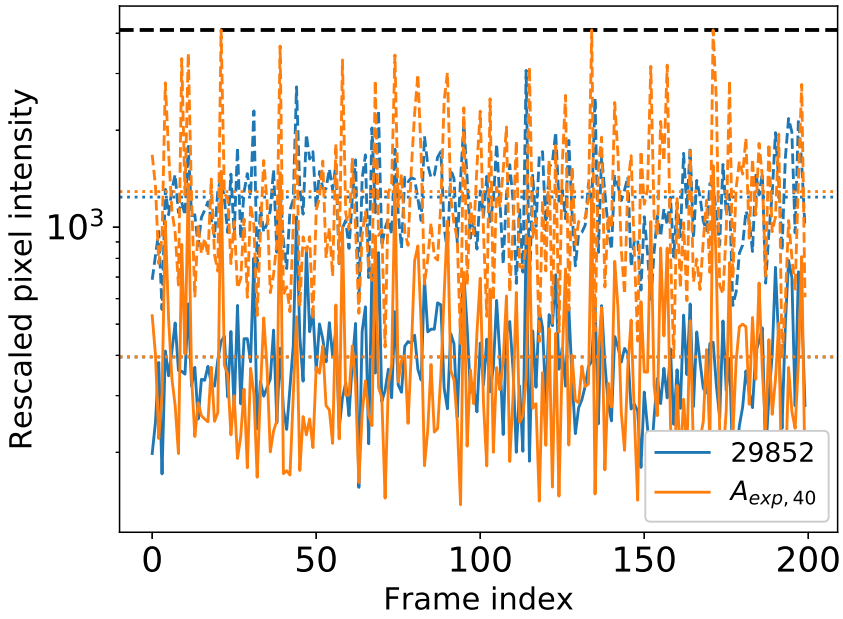

Figure 5.2: Comparison of intensity variation across a subset of 200 experimental (blue) and synthetic (orange) frames. Solid and dashed lines show the mean and maximum intensity of each frame respectively. The amplitude of the synthetic data has been scaled by a factor of 242 in order scale it to the same intensity range as the experimental data. Dotted lines show the averages of each line across the whole datasets, demonstrating that after the rescaling the average frame intensities of the two datasets are the same. The black dashed line indicates the saturation intensity of the camera ($2^{12} = 4095$).

tudes of the synthetically generated filaments are arbitrary. In order for the amplitudes of the synthetic filaments to be consistent with those in experimental data, the synthetic images are scaled so that histograms of the pixel intensities in the synthetic images are consistent with those from experimental frames. This calibration is required for results from synthetic data at a given filament amplitude threshold to be translated to experimental measurements. The distributions of pixel intensities in the synthetic and experimental datasets are compared over the central portions of the images, excluding the divertor regions where the filaments are not analysed. The top 80% of the amplitude range of the images is compared in order to match the intensity distribution of the high amplitude structures constituting filaments, while ignoring the intensity distribution in the dark areas of the images whose structures are not of interest and are mainly a result of the background subtraction algorithm. The intensities of the synthetic frames is then rescaled so that the frame by frame average intensities of the the experimental and synthetic data agree. Figure 5.2 demonstrates this process for calibrating synthetic frame intensities to a set of frames from MAST discharge 29852, for which a calibration factor of 242 is required.

In addition to the primary reference synthetic dataset discussed so far, analysis of

Table 5.3: Distribution information for filament parameters in synthetic data sets A_δ and $A_{\delta,1}$. $\mathcal{L}$ and $\mathcal{U}$ are defined as in Table 5.1, while $\delta(x;\mu) = \delta(x-\mu)$ represents the Dirac delta distribution, centred at μ.

Property		Distribution	Parameters
Position	R	$\lambda\mathcal{L}+\nu$	$\mu = 0$, $\sigma = 0.5$ $\lambda = 0.04$, $\nu = 1.36$
	ϕ	$\lambda\mathcal{U}+\nu$	$\lambda = 355$, $\nu = 2$
Width	δ_R	δ	$\mu = 1.5$
	$\delta_{\phi R}$	δ	$\mu = 1.5$
Amplitude	A	δ	$\mu = 0.02$

three further synthetic datasets will be discussed. Each of these similarly consist of 5000 frames, generated with the same resolution of inversion grid, but contain varying numbers of filaments and varying filament parameter distributions. The properties of each synthetic dataset and their primary application is briefly described below, while a more detailed discussion is given with the results of their analysis in §5.4.6.

Reference Synthetic Data Set The reference synthetic data set is the primary synthetic dataset used in this work. The distributions for each filament parameter are chosen so as to be best representative of experimental data [see Table 5.1]. Each frame contains 40 filament, giving good resemblance to experimental data. The reference dataset shall be abbreviated to data set A.

Single filament physical data set The A_1 data set was produced with the same filament parameter distributions as the reference synthetic data set, but with only one filament in each frame, each of which was ensured to be within the inversion domain. Having only a single filament in each inversion removes the effect of filament overlap and interaction.

40 filament simplified data set The A_δ data set was produced with the same spacial distribution of filaments as in the reference synthetic data set, but the amplitude and width distributions were replaced by delta functions. Therefore all the filaments are in different positions, but have identical properties. See Table 5.3 for the parameters describing this dataset. Here, $\delta(x;\mu) = \delta(x-\mu)$ represents the Dirac delta distribution, centred at μ.

Single filament simplified data set The $A_{\delta,1}$ data set was produced with the same filament parameter distributions as the A_δ data set, but with only one filament in each frame.

5.3 Benchmarking techniques

As previously discussed, geometrical distortions that occur in the inversion process due to the non-orthogonality of the field line basis set can lead to errors and false detections of shadow filaments, complicating the interpretation of the outputs of the technique. Benchmarking is required to quantify and understand these effects in order to draw strong conclusions from experimental applications of the technique. By analysing large quantities of synthetic data and matching detected filaments returned from the filament detection algorithm to input synthetic filaments and comparing their properties, the accuracy and reliability of the technique can be assessed.

5.3.1 Matching synthetic input and detected filaments

In order to quantitatively assess the quality of our inversion technique, we need to classify the detected filaments as either 'true positives' that can be assigned to corresponding true input synthetic filament or spurious 'false positives' that cannot be matched to a true input filament. Before we proceed, we define the *detected filaments* as the output from the filament identification technique and the *synthetic filaments* as the input filaments used to generate the synthetic camera images.

The method by which the set of detected filaments, $\{D_{\text{all}}\}$ and the set of synthetic filaments $\{S_{\text{all}}\}$ are compared and matched is as follows:

1. As discussed earlier, since the synthetic filaments are produced homogeneously around the whole machine, only some of them will fall in the toroidally localized analysis domain. These will form the subset $\{S_{\text{domain}}\}$ of filaments inside the analysis region (see white dashed box in Fig. 4.6c)). Of the filaments in $\{S_{\text{domain}}\}$, those that 1) lie within the analysis domain and 2) have amplitudes above the detection threshold, ϵ_{thresh}, (defined in §4.3) will form an even smaller subset that we call $\{S_{\text{analyse}}\}$.
2. Each detected filament is compared to each unmatched synthetic filament in $\{S_{\text{domain}}\}$. The detected filaments are examined in order of decreasing amplitude.
   (a) If a single synthetic filament centre lies within 1σ of the Gaussian fitted to the detected filament, the input and detected filaments are added to the sets of matched filaments $\{S_{\text{matched}}\}$ and $\{D_{\text{matched}}\}$ respectively.
   (b) If multiple synthetic filaments are within 1σ, the highest amplitude is taken as the 'match' and the rest are considered to 'overlap' with the 'matched' filament as they are too close to it to be distinguished. Overlapping filaments can be generated due to the statistical nature of the synthetic frame generation. However, we do not expect them to occur very frequently in actual experiments.

3. Synthetic filaments in $\{S_{\text{analyse}}\}$ that are neither matched nor overlapping are labelled as 'missed' and added to $\{S_{\text{missed}}\}$. Detected filaments that are inside the analysis region and are not matched are labelled as 'false' detections and added to set $\{D_{\text{false}}\}$.

To summarize, we hence have that

$$\{S_{\text{all}}\} > \{S_{\text{domain}}\} > \{S_{\text{analyse}}\} \approx \{S_{\text{matched}}\} + \{S_{\text{missed}}\} + \{S_{\text{overlap}}\} \tag{5.3}$$

and

$$\{D_{\text{all}}\} = \{D_{\text{matched}}\} + \{D_{\text{false}}\}. \tag{5.4}$$

During this matching procedure a number of additional cases are also handled which result in the approximate equality in Equation 5.3 and influence the figures of merit which are discussed shortly. If a detected filament within the analysis region is matched to a synthetic filament outside the analysis set (i.e. a synthetic filament with an amplitude slightly below ϵ_{thresh}, or located just outside the edge of the analysis region), then the synthetic filament is added to S_{analyse} and the filaments are considered matched as above. If a synthetic filament in S_{analyse} is matched to a detected filament outside the analysis region (i.e. in the border region subject to basis boundary effects) the synthetic filament is removed from S_{analyse} as it does not manifest in the analysis domain and thus would not be observed experimentally.

Finally, with all detected filaments assessed for synthetic matches, we define the set

$$\{S_{\text{isolated}}\} = \{S_{\text{analyse}}\} - \{S_{\text{overlap}}\} = \{S_{\text{matched}}\} + \{S_{\text{missed}}\}, \tag{5.5}$$

which is the set of filaments that 1) are in the analysis domain, 2) have amplitudes above ϵ_{thresh} and 3) can be resolved individually. This is the set of filaments we wish the technique to identify as accurately and reliably as possible (i.e. we want to minimize $\{S_{\text{missed}}\}$).

In Figure 5.3 we picked a particular frame that displays each of the categories of synthetic and detected filament (typical synthetic frames do not generally display this complexity). Synthetic filaments are indicated by dashed ellipses and detected filaments by solid ellipses. Matched filaments are shown in green (belonging to $\{S_{\text{matched}}\}$ or $\{D_{\text{matched}}\}$) and false ($\{D_{\text{false}}\}$) or missed ($\{S_{\text{missed}}\}$) detections are shown in red. Overlapping synthetic filaments ($\{S_{\text{overlap}}\}$) are shown in blue, while synthetic filaments outside the analysis domain or with amplitudes below the threshold (filaments not in $\{S_{\text{analysis}}\}$) are shown in pink.

5.3.2 Benchmarking Figures of Merit

In order to quantify the performance of the inversion technique, a number of numerical figures of merit (FoM) have been developed. These are instrumental to an objective

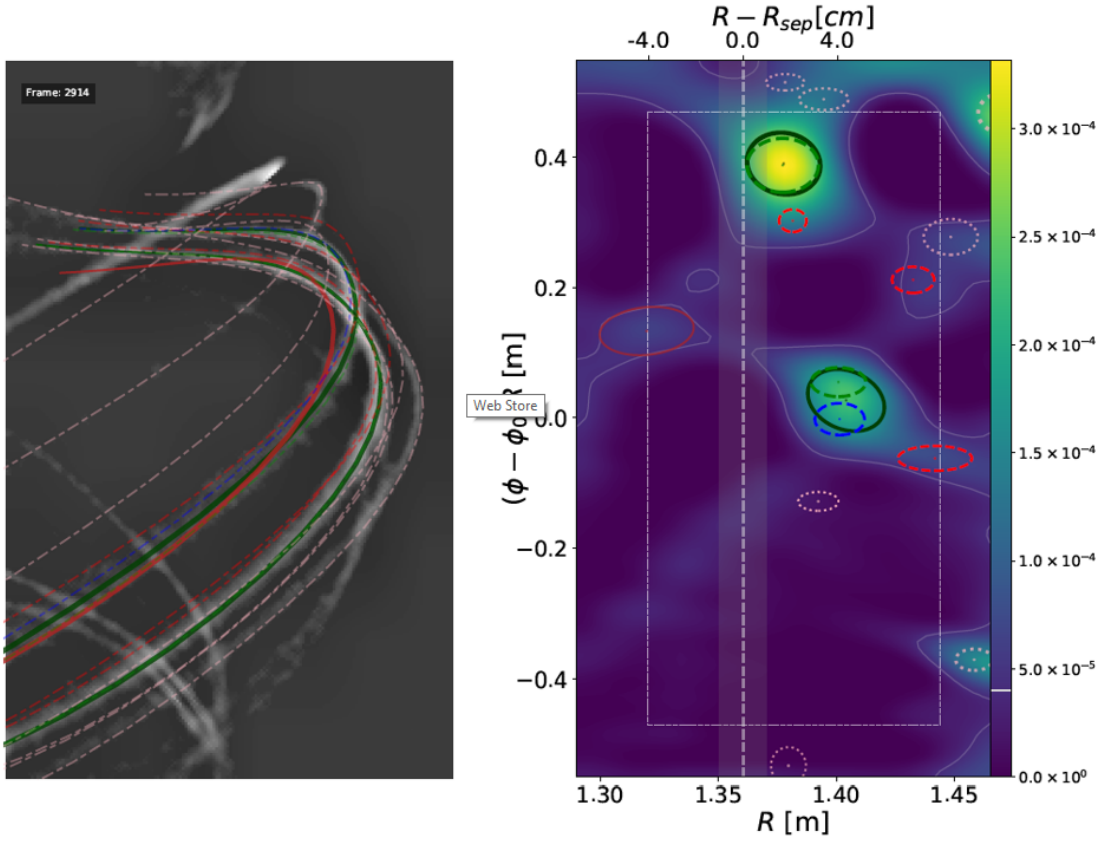

Figure 5.3: a) Synthetic camera frame and b) corresponding inversions with detected (synthetic) filaments represented by ellipses with solid (dashed) lines. Matched filaments are shown in green, false (missed) filaments in red, overlapping filaments in blue and filaments outside the analysis region (white dashed rectangle) or with low amplitudes in pale pink. The white contour and colour bar indicate ϵ_{thresh}.

assessment of the technique, informing further development, as well as giving metrics for benchmarking overall success. Each type of FoM is described below. Here, the notation $|S_{\text{matched}}|$, represents the number of elements in the set $\{S_{\text{matched}}\}$ and the same convention will be used for all the other sets.

Sensitivity We define *sensitivity* (also known as *recall*) as:

$$f_{\text{sens}} = \frac{|S_{\text{matched}}|}{|S_{\text{domain}}|} \times 100. \tag{5.6}$$

This gives the proportion of all filaments in the analysis domain that are detected and is expressed as a percentage. We also define:

$$f_{\text{sens,isolated}} = \frac{|S_{\text{matched}}|}{|S_{\text{matched}}| + |S_{\text{missed}}|} \times 100. \tag{5.7}$$

which gives the proportion of isolated analysis filaments that are detected.

Precision We define *precision* as:

$$f_{\text{prec}} = \frac{|D_{\text{matched}}|}{|D_{\text{matched}}| + |D_{\text{false}}|} \times 100. \tag{5.8}$$

This is the proportion of all detections that are true positives.

These metrics are commonly used in machine learning [236] and related fields [237]. We term these *identification* FoMs as they describe how effectively filaments are identified by the technique.

Parameter fractional errors By calculating the average absolute fractional differences between the measured and synthetic filament parameters the technique can be tuned to individually minimise errors in the position, width and amplitude of the filaments. The percentage error in a parameter p can be expressed as:

$$f_{\text{p}} = \left\langle \frac{|p_{\text{measured}} - p_{\text{input}}|}{p_{\text{input}}} \right\rangle_{\text{matched}} \times 100. \tag{5.9}$$

where $\{p_{\text{measured}}\}$ and $\{p_{\text{input}}\}$ are the sets of detected and synthetic filament parameter values respectively and the angular bracket represents the ensemble average over all matched filaments. Parameter fractional errors are calculated for filament position, widths and amplitudes.

The FoMs calculated using Equation (5.9) are termed *measurement* FoMs as they concern the accuracy with which each parameter is measured.

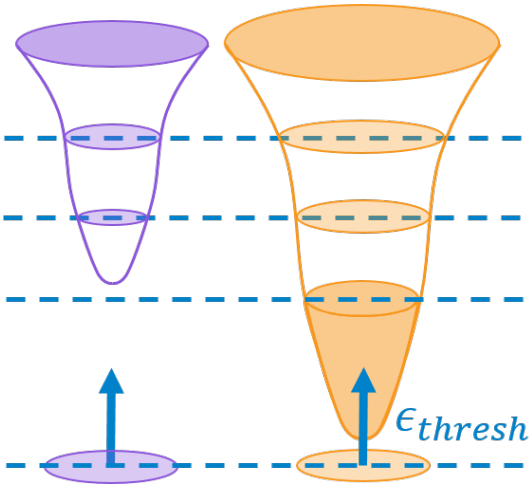

Figure 5.4: Illustration of how successively smaller peaks are included as filament sources in the watershed algorithm as the value of ϵ_{thresh} is lowered. Note the vertical axis is illustrated inverted in analogy with water pooling in low lying regions of the data contours.

5.4 Accuracy and error quantification

5.4.1 Filament amplitude detection threshold

The analysis of synthetic images provides an objective means of setting the threshold amplitude for filament identification described in §4.3. The parameter ϵ_{thresh}, is chosen so as to maximise the precision and sensitivity figures of merit, while sampling a sufficiently large portion of the total filament population. Figure 5.4 illustrates how successively smaller peaks are included as filament sources in the watershed algorithm as the value of ϵ_{thresh} is lowered.

The first function of ϵ_{thresh} is to ensure the exclusion of small local maxima below the noise floor. However, further increases in ϵ_{thresh} yield improvements in the FoMs due to the exclusion of spurious ghost filaments.

At large ϵ_{thresh} only the largest filaments in the population remain in the analysis set, leading to increased sensitivity to filaments in $\{S_{\text{analyse}}\}$, since the detectable objects are brighter and hence more distinct. The precision also increases since only spurious structures from interactions between the very strongest filaments remain, leading to fewer false positives.

Figure 5.5 shows how much of the total filament population (f_{sens}) is excluded for successive gains in precision (f_{prec}) and sensitivity to the targeted filament population ($f_{\text{sens,isolated}}$) as ϵ_{thresh} is increased. The optimal value, $\epsilon_{thresh} = 1.5 \times 10^{-2}$, is chosen, as above this value there are relatively modest gains in precision and isolated sensitivity. At this level, 50% of the total filament population lies above the amplitude threshold. The technique detects 36% of the total filament population and 74% of isolated analysis filaments. This is achieved with a 98.8% true positive detection rate. These FoM values

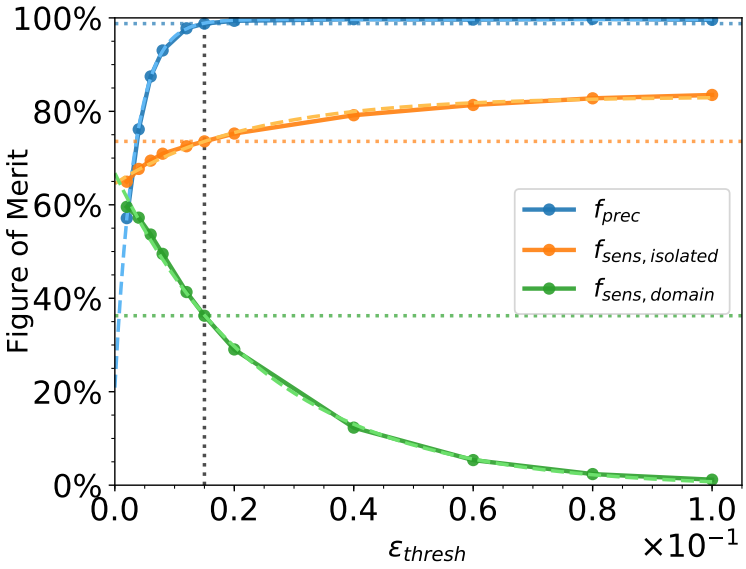

Figure 5.5: Dependence of the precision (blue), analysis sensitivity (orange), and total sensitivity (green) FoMs on filament amplitude detection threshold, $\epsilon_{\mathrm{thresh}}$. Dashed lines show the FoM values achieved at the value of $\epsilon_{thresh} = 1.5 \times 10^{-2}$ used in the analysis.

Table 5.4: Breakdown of the numbers of synthetic filaments in each analysis category for the reference synthetic data set, analysed with $\epsilon_{thresh} = 1.5 \times 10^{-2}$.

Set S_{x}	$\lvert S_{\mathrm{x}}\rvert$	$\frac{\lvert S_{\mathrm{x}}\rvert}{\lvert S_{\mathrm{domain}}\rvert}$	$\frac{\lvert S_{\mathrm{x}}\rvert}{\lvert S_{\mathrm{analysis}}\rvert}$	$\frac{\lvert S_{\mathrm{x}}\rvert}{\lvert S_{\mathrm{isolated}}\rvert}$
S_{domain}	20243	100%	-	-
S_{analysis}	10134	50%	100%	-
S_{isolated}	9978	49%	98%	100%
S_{matched}	7341	36%	72%	74%
S_{missed}	2637	13%	26%	26%
S_{overlap}	156	0.8%	1.5%	-

are summarised in Table 5.6.

The dependence of $f_{\mathrm{sens,\ domain}}$ on ϵ_{thresh} is broadly exponential, rolling over slightly at low amplitudes as the noise floor of the technique is approached. The distribution peaks at 59%, indicating that, for a decaying exponential amplitude distribution, over half of the total filament population is potentially detectable (i.e. above the noise floor of this technique), with the drawback that at this amplitude threshold there are almost equally as many false positives ($f_{\mathrm{prec}} \approx 57\%$). The precision saturates around $\epsilon_{thresh} = 2.0 \times 10^{-2}$, above which nearly all detections are true positives.

5.4.2 Breakdown of filament detections

With the figures of merit in place and the threshold set, we can now quantify the performance of the inversion technique. Table 5.4 gives a breakdown of the number of filaments that are assigned to each category by the matching algorithm for the refer-

Table 5.5: Breakdown of the numbers of detected filaments in each analysis category for the reference synthetic data set, analysed with $\epsilon_{thresh} = 1.5 \times 10^{-2}$.

Set D_{x}	$\lvert D_{\mathrm{x}}\rvert$	$\frac{\lvert D_{\mathrm{x}}\rvert}{\lvert D_{\mathrm{domain}}\rvert}$
D_{domain}	7432	100.0%
D_{matched}	7341	98.8%
D_{false}	91	1.2%

Table 5.6: Summary of the figures of merit from analysing the reference synthetic data set with $\epsilon_{thresh} = 1.5 \times 10^{-2}$. The first three FoMs should be maximised and the remainder minimised. Values prior to the systematic offset correction are given in brackets. $f_{R-R_{sep}}$ has been normalised relative to the 5 cm SOL width between the LCFS and the first limiting structure.

Quantity, p	Error		FoM
	μ_p	σ_p	f_p
f_{prec}	-	-	98.8%
$f_{\mathrm{sens,isolated}}$	-	-	74%
$f_{\mathrm{sens,domain}}$	-	-	36%
$R - R_{sep}$ [cm]	0.32	0.14	4% (7%)
$R\phi$ [cm]	-0.25	0.18	15% (22%)
δ_R [cm]	0.56	0.27	33% (70%)
$\delta_{R\phi}$ [cm]	1.1	0.73	26% (67%)

ence synthetic data set, analysed with $\epsilon_{thresh} = 1.5 \times 10^{-2}$. Around 10% of all the synthetic filaments spawned around the full toroidal extent of the machine fall within the window in toroidal angle of 46° under analysis and within the borders of analysis domain. 50% of those filaments have amplitudes above ϵ_{thresh}, while 49% both have amplitudes above ϵ_{thresh} and do not overlap with higher amplitude filaments. Around 1% of filaments within the analysis domain are obscured by higher amplitude filaments, which present themselves together as a single local maxima in intensity. While these underlying filaments, obscured by larger filaments cannot be detected, they raise the amplitude of the larger detected filaments. However, we do not expect these filament interactions to occur routinely in experiment.

Table 5.5 gives a breakdown of the detected filaments. Of the 7441 detections only 91 (1.2%) were false positives that did not match to a synthetic filament.

5.4.3 Error Distributions and Corrections

After quantifying the detection capabilities, we now address the problem of assessing the quality of the measurements. For each detected filament, it is possible to calculate the absolute error on a parameter p, as defined by $\xi_p = p_{detected} - p_{synthetic}$. Since different filaments can have different errors, the latter are statistically distributed. Assessing the average errors, $\mu_{\xi p}$ identifies systematic errors arising from distortions in the inversion which can be corrected for through offsets in future measurements. Assessing

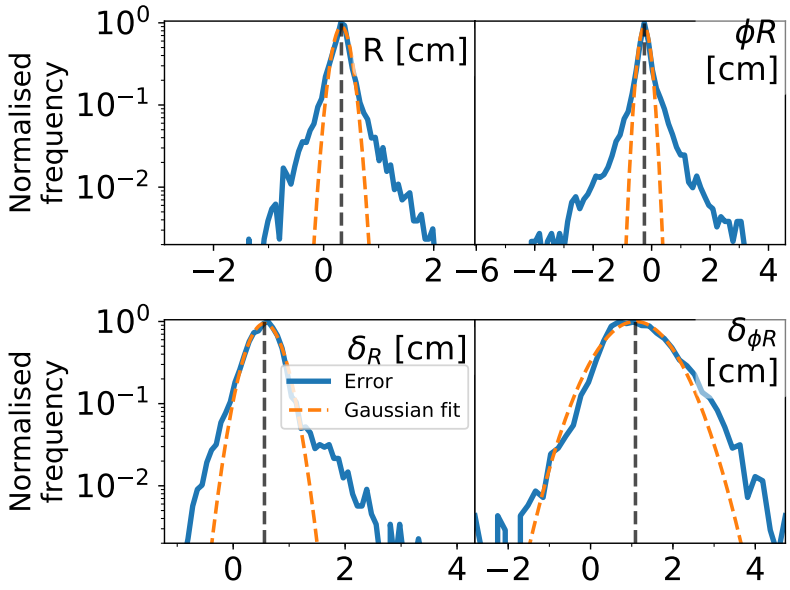

Figure 5.6: Distribution of errors in detected filament parameters for radial position, toroidal position, radial width and toroidal width (left to right, top to bottom respectively). The solid blue lines shows the distributions of errors, while the dashed orange lines shows Gaussian fits to the distributions.

the standard deviation in the errors, $\sigma_{\xi p}$, quantifies the random errors that remain after the systematic errors are corrected for.

Figure 5.6 shows the resultant distributions of absolute errors, ξ_p, on the position and width measurements for all matched output ($\{S_{\text{matched}}\}$) filaments from the reference synthetic data set. Each parameter shows a broadly Gaussian distribution of errors over at least an order of magnitude in frequency, although long low frequency tails are present. Therefore, we fit Gaussians to the distributions in order to parametrise their widths, without suggesting that the underlying distributions are truly Gaussian. The mean and standard deviation of the fitted Gaussians are listed in Table 5.6. The table also gives the f_{p} values given by Eq. (5.9) for each parameter (note $f_{R-R_{sep}}$ has been normalised relative to the 5 cm SOL width between the LCFS and the first limiting structure).

The average errors in the R and ϕR positions are small, at 0.3 cm and -0.2 cm respectively. The widths of the fitted Gaussian are also small, each around 0.2 cm, although the raised wings and long, largely symmetric tails with Fisher kurtosis values of 8.9 and 12.3 respectively, indicate that rare large errors of greater than 1 cm can occur with greater frequency than suggested by the Gaussian widths. Almost 3% of R positions having errors greater than 1 cm (5σ) and around 13% of ϕR positions having errors greater than 1 cm (5σ). The high accuracy of the position measurements is important for accurate filament velocity measurements, corresponding to velocity errors of the order 0.3 km/s in both radial and toroidal directions. For a typical filament

travelling radially at 0.7 km/s and toroidally at 4 km/s, this corresponds to percentage errors of 40% and 7% respectively. The systematic errors in the width δ_R and $\delta_{\phi R}$ are more significant, at 0.6 cm and 1.8 cm respectively. These over estimations of widths arise from a number of factors, most importantly, the finite analysis grid resolution, overlap of filaments and distortions during the inversion process. The error distributions are described well by the fitted Gaussian, with widths of 0.3 cm and 0.7 cm respectively, although again the tails of the distributions are slightly more pronounced, with kurtosis values of 7.9 and 1.1 respectively.

With the systematic errors in these parameters, $\mu_{\xi p}$, identified, they are applied as offset corrections in the analysis that follow, so that measurements should only be subject to random measurement errors.

5.4.4 Spacial dependence of systematic errors

In addition to exploring the overall statistics of the detected and synthetic filaments, the spatial variation of the detection errors can be compared throughout the analysis domain. Given the prevalence of distortive effects in different areas of the inversions depends on the complex geometrical line of sight interactions of the different basis field line images, there is a systematic spatial dependence of the errors on the measured filament parameters. By measuring the errors, ξ_p, on individual synthetic filaments at different positions in the analysis domain, these spatial error dependencies can be measured.

Figure 5.7 shows how the errors on individual filament positions and widths vary with radial and toroidal position. Note that these systematic errors are calculated for filaments with a single size and amplitude, so we emphasize the spatial structure of the distributions rather than the absolute values of the errors. Significant toroidal structure in the systematic error distributions is observed while the radial dependence is weak. This is to be expected given the relatively small radial variations in the magnetic geometry across the radial extent of the analysis domain, compared to that of the toroidal variation. There is a particularly pronounced increase in the errors in the widths and radial positions at low values of ϕR. These measurements motivated the choice of analysis domain discussed above, with a central toroidal angle of $\phi_0 = 177°$, rather than the value of $\phi_0 = 170°$ used in Figure 5.7, in order to avoid the region of low ϕR, where the systematic errors are largest.

5.4.5 Position and Width Distribution Measurements

The ability to accurately measure the statistical properties of filaments is essential to build reliable first principles models and to assess their intermittent loads on the plasma facing components. We now wish to understand how well our technique is able to reproduce the distributions of the synthetic filaments.

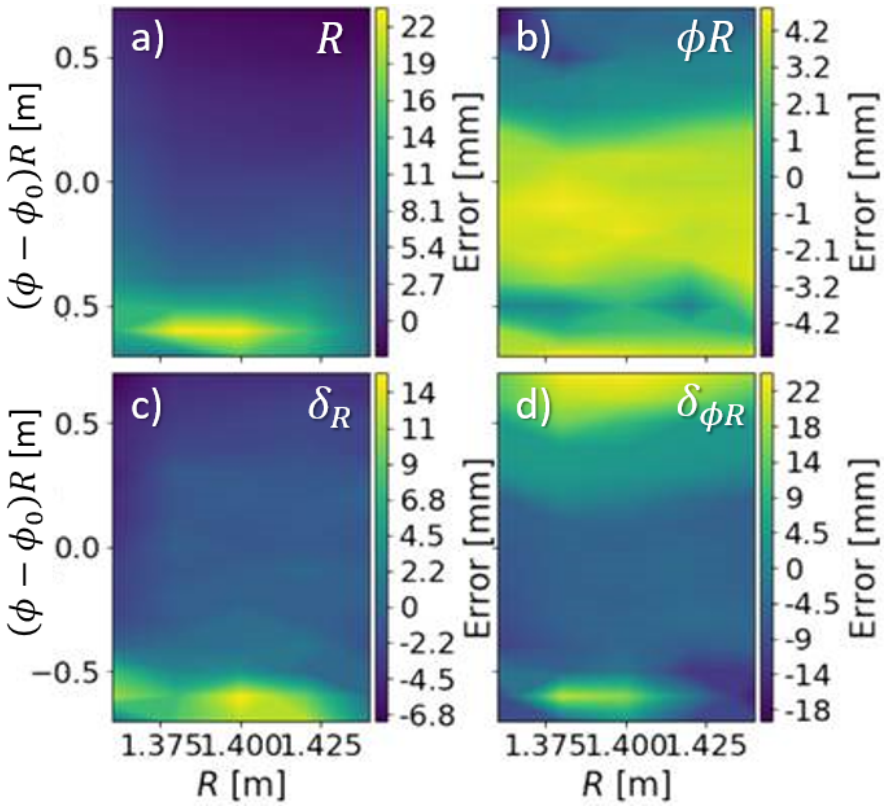

Figure 5.7: Spatial dependence of measurement errors in a) radial position, b) toroidal position, c) radial width and d) toroidal width. These distributions are measured by assessing the systematic error on individual synthetic filaments located at different positions in the analysis domain.

Figure 5.8 compares the frequency density distribution functions of filament positions and widths for the original synthetic filaments ($\{S_{\text{domain}}\}$, orange dashed lines) and the detected output filaments ($\{D_{\text{domain}}\}$, black solid lines). The detected values have been corrected by the systematic offsets identified in the previous section.

The distribution functions of radial and toroidal positions are very well recovered, closely following the synthetic distributions. The detected width distributions preserve the overall shape of the originals, although the distribution of δ_R is more peaked and the distribution of $\delta_{\phi R}$ has its peak shifted to the larger widths by around 0.6 cm. This gives us confidence that detected distributions of filament positions and widths accurately reflect the forms of the true underlying physical distributions.

It is useful to analyse these results in more detail, in particular, the partial distributions of specific subsets in order to determine trends. These partial distributions are displayed and compared in Figure 5.9, which we will now discuss. The distributions of matched ($\{S_{\text{matched}}\}$) and missed ($\{S_{\text{missed}}\}$) filaments collapse onto the overall distribution of synthetic filaments ($\{S_{\text{domain}}\}$) showing that filaments are not selectively detected and thus the full extent of the distributions are evenly sampled. This is strong

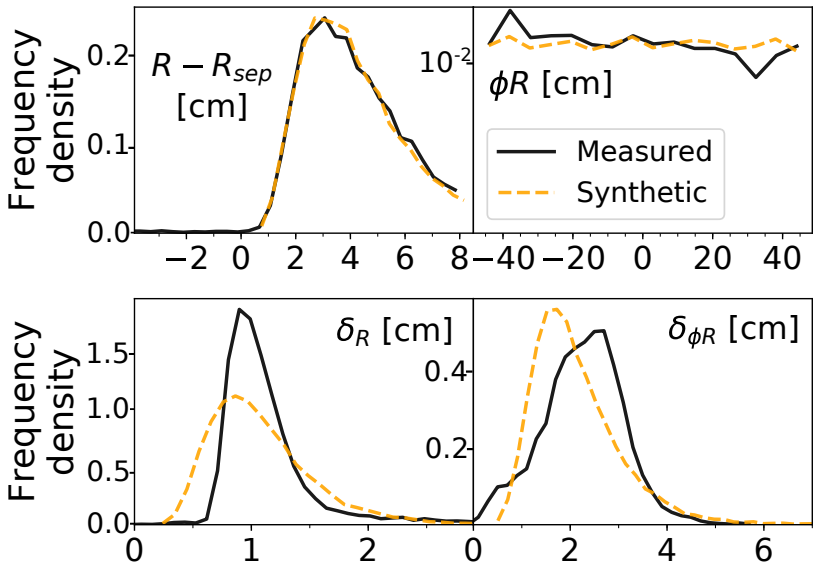

Figure 5.8: Comparison of synthetic input and measured output frequency density distributions of filament positions and widths for the reference synthetic data set. Distributions of synthetic input filaments are given by orange dashed lines and detected distributions by black solid lines.

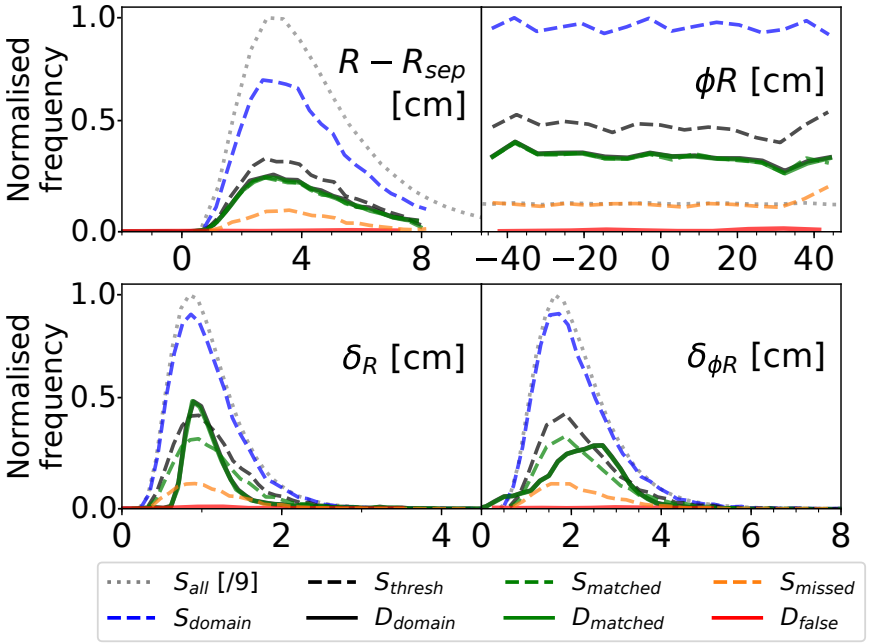

Figure 5.9: Comparison of input and output normalised frequency distributions of filament positions and widths for the reference synthetic data set. Dashed and dotted lines indicate sets of synthetic filaments, while solid lines indicate sets of detected filaments recovered from the analysis procedure. The frequencies for the full set of synthetic filaments (dotted grey line) have been divided by 9 to fit on the same axes.

evidence that filaments have the observed properties and are not, for instance, all a single size. Similarly, the distribution of matched detections ($\{D_{\text{matched}}\}$) collapses onto the overall distributions of detections.

The distribution of false positive detections ($\{D_{\text{false}}\}$), however, shows separate structure that cannot be resolved in Figure 5.9 due to their low frequency. Indeed, false detections are over twice as likely to occur at low R values, which is anticipated since line of sight effects mean ghost filaments typically occur at lower R values than the two true filaments which overlap to produce them. The distribution of false positive detections has two peaks in ϕR located around ϕR = -20 cm and +30 cm, indicating locations where line of sight effects are particularly prone to producing ghost filaments. Both width distributions have higher relative false positive rates for larger width filaments, especially in the case of $\delta_{\phi R}$. Dashed and dotted lines show the original distributions, while solid lines show the measured distributions. The dotted grey lines show the distributions for all 200,000 synthetic filaments in the data set distributed around the full toroidal extent of the machine (divided by 9 to fit on the same axes), while all other lines refer to sets of filaments within the analysis domain. Here, the proportion of filaments in each matching category is clear, demonstrating how few filaments are missed or false detections.

5.4.6 Amplitude Distribution Measurements

In order to compare the synthetic and detected filament amplitude distributions the synthetic amplitudes must be mapped to corresponding amplitudes in the inversion data. Figure 5.10 shows the inversion amplitude measured at the true centre of each synthetic filament in the reference synthetic data set as a function of the uncalibrated synthetic input amplitude. While there is some scatter resulting from the overlap of filaments and distortions in the inversion, there is a strong linear dependence with a Pearson correlation coefficient of 0.79. Excluding high amplitude outliers, the gradient of a linear fit to the data yields the calibration factor required to compare input amplitudes and inversion amplitudes directly.

Figure 5.11 compares the distributions of detected filament amplitudes and calibrated synthetic amplitudes, using line styles consistent with Figure 5.9. The vertical grey line indicates the amplitude threshold $\epsilon_{thresh} = 1.5 \times 10^{-2}$, below which detections are not attempted. The output distribution successfully captures the exponential nature of the input distribution, but for a roll over at low amplitudes, arising from from the increased proportion of missed filaments at low amplitude. The majority of false positive detections also occur at low amplitudes. The matched synthetic filaments curve continues below ϵ_{thresh} due to the subset of synthetic filaments with amplitudes below ϵ_{thresh} that manifest in the inversion with amplitudes above ϵ_{thresh} and are detected.

In order to assess the fidelity of the amplitude measurements and ensure that mea-

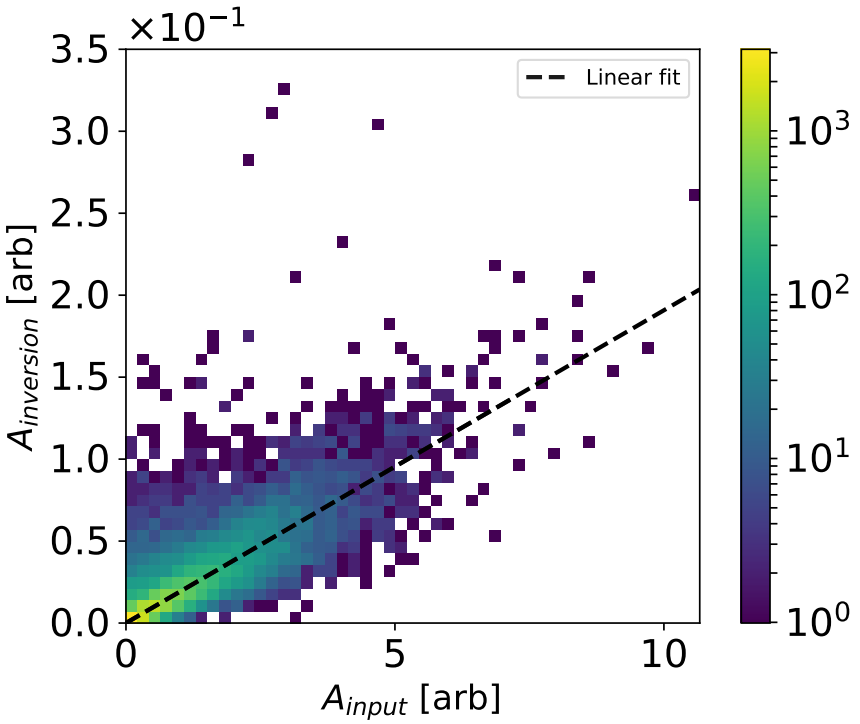

Figure 5.10: Relationship between synthetic input field line emission amplitude and resulting inversion field line emission amplitude.

sured exponential amplitude distributions are not an artefact of the measurement technique, other data sets with different amplitude distributions have been analysed, each containing 5000 frames.

Data set A_1 shares the same experimentally representative parameter distributions as the reference synthetic data set, A, but contains only one filament per frame. This data set is used to assess the effect of filament overlap on results from the reference dataset.

Data set $A_{\delta,1}$ contains a single randomly positioned filament per frame, each with the same fixed amplitude and filament size. This heavily simplified data set collapses the input amplitude and width distributions onto delta functions, so that the width of the recovered amplitude distribution is solely due to the intrinsic inaccuracies in the analysis technique. The output amplitude distribution then provides an estimate of the intrinsic random error on an amplitude measurement.

Data set A_δ contains the same distributions as $A_{\delta,1}$, but with 40 filaments per frame, so that the effects of filament overlap can be readily identified.

Figure 5.12 shows the calibrated synthetic amplitude distributions for each of these data sets (dashed lines) and the inversion amplitudes measured at the true centre of each filament in the analysis domain (solid lines). In the case of the simple $A_{\delta,1}$ data set, the input delta function in amplitude is dispersed but maintains the peaked, broadly sym-

metrical shape. Given there is only one filament per frame, this data set demonstrates the intrinsic dispersion in amplitudes resulting from the inversion technique without any filament overlap. The Gaussian width of the measured distribution is around 15% of the mean, indicating that this is the typical size of error that can be expected on amplitude measurements, without filament overlap effects.

When multiple filaments are introduced as in the A_δ data set, the highly peaked distribution shape is retained, but the interactions between filaments lead to further broadening of the distribution to a Gaussian width around 25% of the mean, a shift to higher amplitudes and an enlarged high amplitude tail. It should be noted that while the bulk of these effects will primarily be the result of the overlap of foreground and background filaments and are thus relevant to experimental measurements, the high amplitude tail and a degree of the broadening will be due to overlap of nearby filaments which are not anticipated experimentally.

In the case of the single filament, physically representative data set, A_1, the inversion amplitudes closely follow the input, indicating that in the absence of filament overlap, amplitude distributions are recovered reliably. The addition of multiple filaments per frame in A introduces a constant positive offset due to the overlap of foreground and background filaments raising the average intensity of each projected field line.

As the filaments are uniformly distributed around the machine and there are many possible points of intersection between different field lines in the field of view of the camera, the average effect of the filaments outside the analysis domain is to provide an effectively uniform background of emission. However, variation in field line intersection in different parts of the image due to line of sight effects will lead to some dispersion in the background offset. Hence, a roll over is observed in the measured amplitudes. Therefore, the measured amplitude distribution is representative of the true distribution at all but the lowest amplitudes, except for an approximately constant shift to higher amplitudes due to filament overlap. Given the roll over occurs below the optimal blob detection amplitude threshold, $\epsilon_{thresh} = 1.5 \times 10^{-2}$, and the constant offset does not affect the distribution shape, this should not affect experimental conclusions about measured amplitude distributions.

5.5 Discussion

In this chapter we have investigated the accuracy and reliability of the Elzar fast camera data processing technique. A number of further factors should be considered that affect the precision and reliability of the technique.

The grid resolution of field lines used in the geometry matrix plays a critical role in the precision of technique's output. This must be sufficiently high to resolve the smallest filaments of interest, which are of the order of several millimetres across and in general a few ion Larmor radii. However, the distortion and smoothing in the inversion

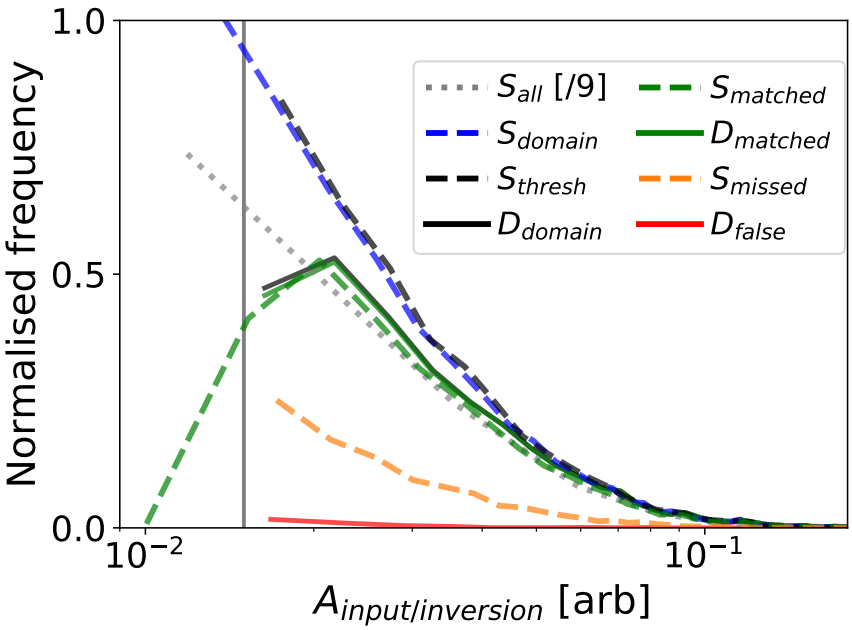

Figure 5.11: Synthetic and detected amplitude distributions. Line styles are consistent with Figure 5.9. The vertical grey line indicates the amplitude threshold $\epsilon_{thresh} = 1.5 \times 10^{-2}$.

and regularisation process enlarges structures in R-$R\phi$ space, such that, provided these effects are accounted for as in the systematic error correction, the grid resolution can be slightly coarser than the structures being measured.

The resolution that can practically be used is limited by the computational time per analysis filament and the extent of the analysis region. It is desirable that the analysis region be large for two reasons. Firstly, analysing a large region makes best use of the available data and produces more filament detections, so that stronger statistical conclusions can be drawn from the larger sample sizes. Secondly, a large analysis domain maximises the proportion of data that is not affected by boundary effects (which scales as the product of the grid dimensions). The grid resolution used in Section 5.3 was chosen to achieve an equitable balance, such that the necessary features could be resolved, without increasing the computation cost to an extent that it would prohibitively limit the amount of data that could be analysed.

The camera spatial resolution, limited by its pixel resolution and viewing perspective, also limits the accuracy of the inversion. Filament widths in the camera field of view are often only of the order of a few pixels across over much of their length, leading to pronounced discretisation of the intensity distribution across their width. This leads to sharp discontinuities in pixel intensity along field lines, particularly in areas where the field lines pass at 45° to the image axes. This can result in artificial inhomogeneities along field lines and striations in inversion intensity, which complicate identification.

While the camera sensor has a 12 bit bit-depth (0-4095), the OpenCV image pro-

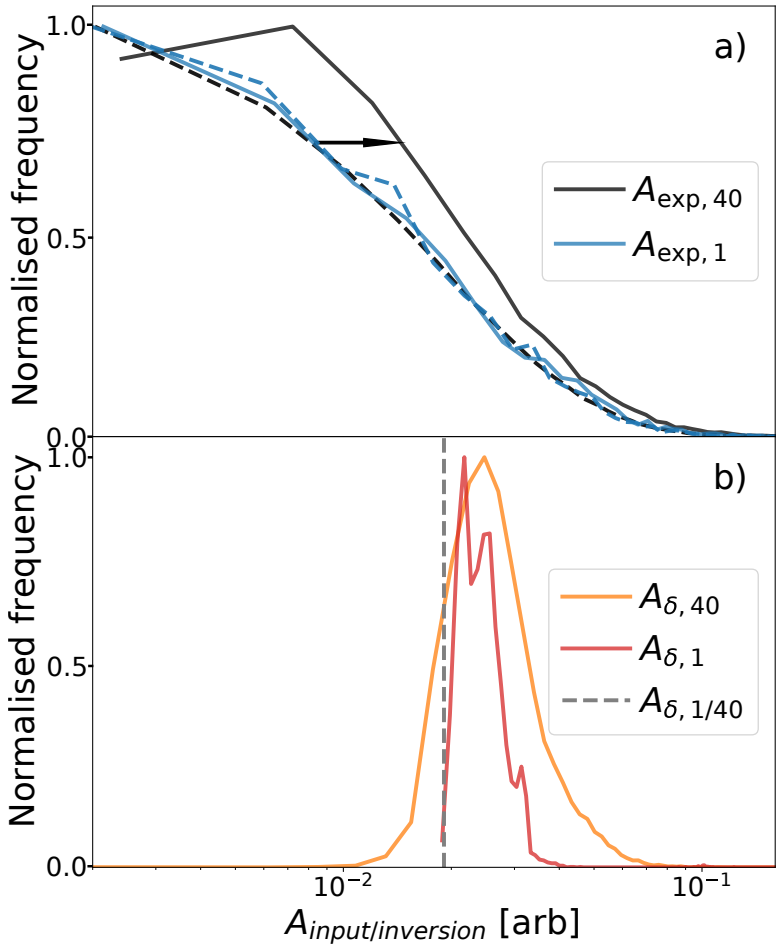

Figure 5.12: Calibrated original and measured inversion amplitude distributions for a) datasets A and A_1 with exponential amplitude distributions and b) datasets A_δ and $A_{\delta,1}$ with delta function amplitude distributions. Dashed lines are the calibrated input amplitude distributions, while solid lines show the distributions of inversion amplitudes. The vertical dashed grey line indicates the delta function in input amplitudes for the the A_δ and $A_{\delta,1}$ datasets.

cessing library used in the prepossessing stages of the analysis requires the images to be converted to 8 bit bit-depth (0-255). This reduction in dynamic range propagates through to the inversion and reduces the fidelity of filament contour selection. The smallest, faintest filaments are only a few pixels across and contain pixels with intensities at the lower end of the 8-bit dynamic range (0-255). Therefore, the range of the amplitudes in the contour can be small leading to significant âĂŸterracingâĂŹ of the intensities across the contour. When fitting the 2D Gaussians across only a few points, this can have noticeable effects for fitting to the smallest, faintest filaments. Therefore, avoiding the conversion to 8-bit data for the OpenCV operations would preferable in future.

At 100 kHz the camera's integration time is 3.3 μs, during which fast filaments [48] with total velocities of ~ 4 km/s can be expected to move up to 1.3 cm. Kirk *et al.* [51] has previously investigated the effects of exposure time on filament measurements under similar conditions and has found it not to be significant.

A source of error that has not been quantified in this benchmarking exercise is that arising from inaccuracies in the magnetic equilibrium and camera calibration, which will never be perfect descriptions of the physical systems. A non-quantitative assessment of these errors has shown they can be significant for poor quality camera calibrations and equilibrium reconstructions. However, if care is taken in producing these inputs their effects on the benchmarking conclusions should not be large. While quantification of these sources of error is an important area of future study, it is beyond the scope of this thesis.

5.5.1 Conclusions

The Elzar camera inversion and filament detection technique has been benchmarked using synthetic camera data, containing filaments with known properties. This has enabled the precision, reliability and limitations of the technique to be assessed, in order to properly inform experimental applications and further development of the technique. The technique successfully identified 74% of the filaments in the analysis region and with sufficient amplitude, with a true positive detection rate of 98.8%. Standard errors on filament's radial and toroidal positions are around ± 2 mm, while standard errors on radial and toroidal widths are around ± 3 mm and ± 7 mm respectively.

Given a sufficiently large sample, the measured filament data has been shown to qualitatively reproduce features of the input distributions, so that the technique can be used to provide important constraining information for analytical interpretation of SOL profiles.

Chapter 6

Application of the Elzar filament analysis technique to MAST data

In this Chapter, we will introduce a first but rather complete experimental application of the inversion technique and detection algorithm that we have discussed in the previous chapters. We start with a comparison of measurements from the visual camera with those from a reciprocating probe (RCP) diagnostic, which is one of the most reliable workhorses in experimental SOL physics. Combined with the synthetic benchmarking in the previous chapter, this establishes the comparative and unique capabilities of the technique relative to established techniques.

We then proceed to show how the detections of filaments using the 2D inversion technique relate to the 0D measurements used in the comparison with the RCP diagnostic. With this established, we describe how the analysis of a small number of filaments with the inversion technique has enabled the benchmarking and cross-validation of a number of turbulence codes. By utilising the synthetic camera diagnostic to forward model the simulation data, the camera inversion technique provided a unique means of benchmarking the turbulence codes, while also gaining additional insight into the experimental measurements from the simulation results.

Next we turn to the application of the inversion technique to large quantities of experimental camera data in order to measure the distributions of filament properties that are required by analytic models of SOL transport. PDFs of positions, widths amplitudes and velocities are presented and insight from the synthetic benchmarking exercise is applied to their interpretation.

6.1 Comparison of imaging and Langmuir probe measurements

6.1.1 Introduction

The comparison of the new camera inversion technique with well established reciprocating probe analysis techniques, puts its capabilities in context with a large body of

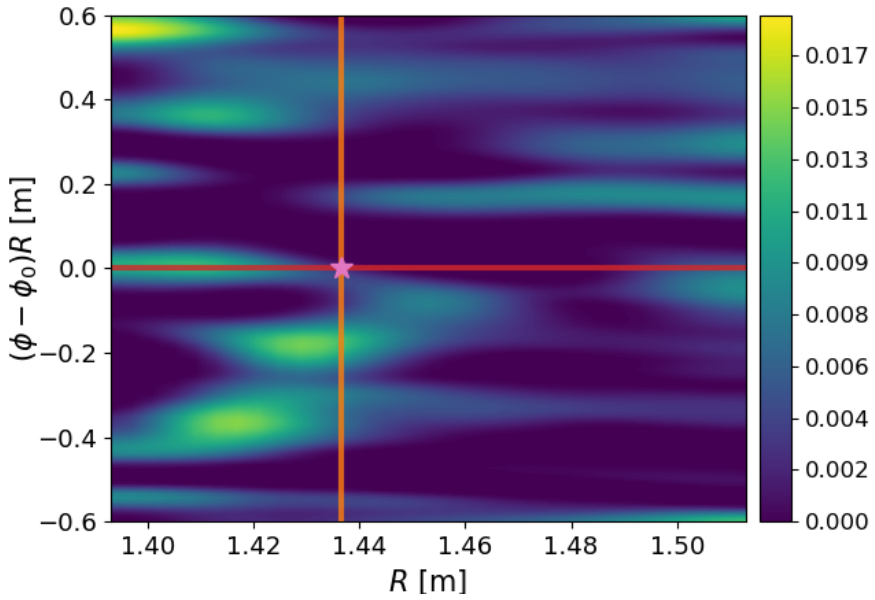

Figure 6.1: Illustration of taking reduced dimensionality slices through the 2D camera inversion data for comparison with Langmuir probe data. The time average and higher order moments of the data indicated by the red radial slice through the inversion can be compared to time equivalent radial profiles of the ion saturation current from RCP measurements. Similarly, the statistics of the time varying signal of an individual field line such as the 'probe field line' indicated by the mauve star can be compared to that of RCP I_{sat} signals. The extensive toroidal slice through the data shown in orange has no equivalent in measurements from other diagnostics and uniquely facilitates the measurement of the toroidal separation between filaments. This inversion is for camera data from discharge 29840 at time $t = 0.16011$ s.

existing measurements in the literature. By identifying commonalities and differences in results from the two diagnostics, results from each can be better understood and interpreted going forward.

MAST's reciprocating probe was moving inwards towards the plasma core on a timescale that was much longer than the typical filament residence time. Therefore, at any given point, its signal approximated that of a probe fixed in space. By considering the signals from given radial and toroidal positions in the inverted camera map, time series are easily produced that are similar to those from a Langmuir probe diagnostic. Figure 6.1 illustrates a radial slice through the camera inversion data in red and a toroidal slice through the data in orange. An individual 'probe field line' is signified by a mauve star. The time varying intensity of such a probe field line shows spikes in intensity due to passing filaments in the same manner as the filamentary spikes seen in RCP ion saturation current signals. Indeed, under the assumption that filaments are homogenous along field lines, a probe field line passing through the location of the RCP tip should measure the same fluctuation. Past measurements on MAST have indeed

shown spikes in I_{sat} to occur at the same time as increased D_α emission from iteration of the filament with the probe head [238]. With the camera data treated in this way, Langmuir probe measurement techniques can be used to perform, for the first time, a like-for-like comparison of inverted visible camera measurements and Langmuir probe measurements.

The emission intensity observed in visible camera images and the ion saturation current measured by Langmuir probes both have a primary dependence on plasma density and thus can be considered proxy measurements for this quantity. However, the secondary dependencies on electron temperature and neutral gas density differ between the two diagnostics (see §1.6), preventing a fully quantitative comparison of the two measurement techniques. Therefore we should expect the two diagnostics to qualitatively reproduce the same trends in the filament statistics.

In the absence of a single MAST pulse with both fast camera and Langmuir probe data, two similar pulses were studied. The camera analysis was performed on MAST discharge 29840 and the reciprocating Langmuir probe (RCP) analysis on discharge 29023, both of which were ohmic, L-mode, double null plasmas (refer back to §3.4 for further details). The probe reciprocation occurred over a 350 ms time window during the plasma flat top, traversing to the separatrix from a distance of around 15 cm. Given the RCP had a sample rate of 500 kHz while the camera frame rate was 100 kHz, the RCP data was down sampled to 100 kHz so that a direct statistical comparison between the two signals could be made. The statistical properties of the down sampled RCP data are very similar, although not identical to those of the full signal. Therefore, the down sampled RCP data is used throughout the analysis.

The camera analysis was performed over a time window lasting 90 ms, enabling large samples of filaments to be analysed. The analysis for each diagnostic was performed at the same distance from separatrix, but at two different toroidal locations, owing to the differing locations of the diagnostics within the machine (see Chapter 3). Therefore, due to the toroidal symmetry of the machine, both diagnostics should measure the same statistical properties of the filaments.

6.1.2 Langmuir probe pre-processing

The Langmuir probe I_{sat} signal for discharge 29023 was subject to electromagnetic interference from a poloidal field coil which impacted the statistical properties of the signal and had to be removed. This noise is evident in the frequency bands present at all times in the spectrogram in Figure 6.2a). The frequencies of the noise were isolated by considering the start of the I_{sat} signal prior to the reciprocation for which the probe was far enough away from the plasma ($R - R_{sep} = 15$ cm) that the majority of the signal could be attributed to the field coil pickup. The spectrum of this signal was used to identify high frequency sources of noise which were filtered with notch filters.

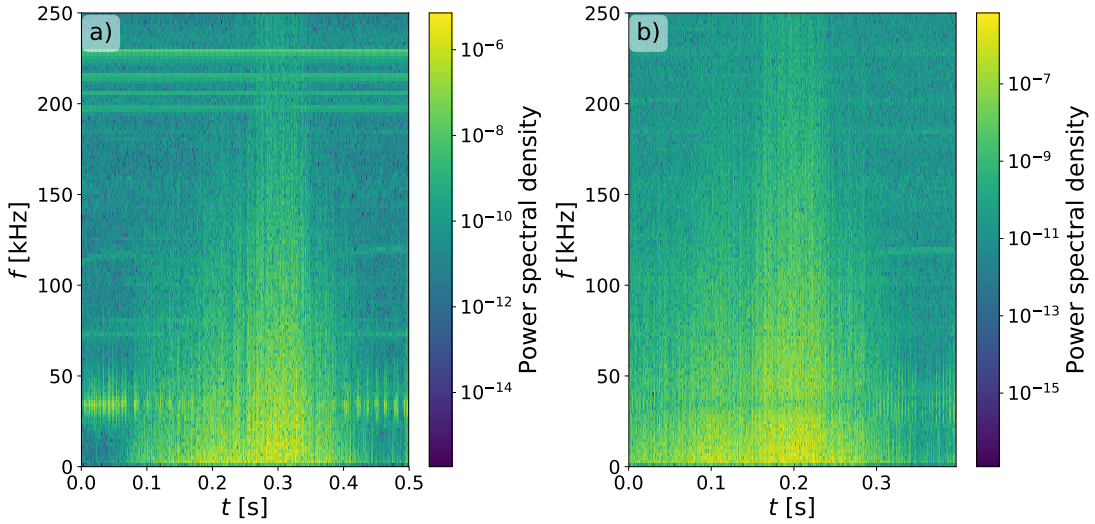

Figure 6.2: a) Spectrogram of the RCP ion saturation current signal for discharge 29023. Electromagnetic interference from a nearby poloidal field coil can be seen in the form of frequency bands present at all times. b) Spectrogram after application of the frequency filter in Figure 6.3, showing the removal of the erroneous bands.

At lower frequencies, around 40 kHz, there was a broadband source of noise coinciding with frequencies present in the filamentary signal spectrum. Therefore these frequencies were suppressed in proportion to their relative noise component. This approach removed the noise without impacting on the statistical properties of the filamentary signal, and thus affecting the comparison. Figure 6.3 shows the filter used to suppress the noise, containing several narrow notch filters and a broader structured filter.

6.1.3 Time series

Figure 6.4 compares 5 ms time series from both the RCP and the camera data taken at the separatrix. The top panel shows the time varying intensity, I_{fl}, of the probe field line intersecting the mid-plane at the separatrix ($R - R_{sep} = 0$ cm) and located toroidally at $\phi = 177°$. The bottom panel shows the ion saturation current drawn by the probe tip at the separatrix.

There is a strong resemblance between the two signals, with both showing similar bursty waveforms as filaments pass the respective probe field line and probe tip. However, several differences can be noted.

First, a distinct flat background in intensity is observed in the camera data. This results primarily from the camera background subtraction algorithm, which results in areas of the camera images having close to zero intensity and is further contributed to by the non-negative SART inversion algorithm, which enforces a sharp cut off at zero to prevent negative intensity values. As the RCP data does not have an equivalent background subtraction step, the I_{sat} signal shows the filamentary fluctuations on top

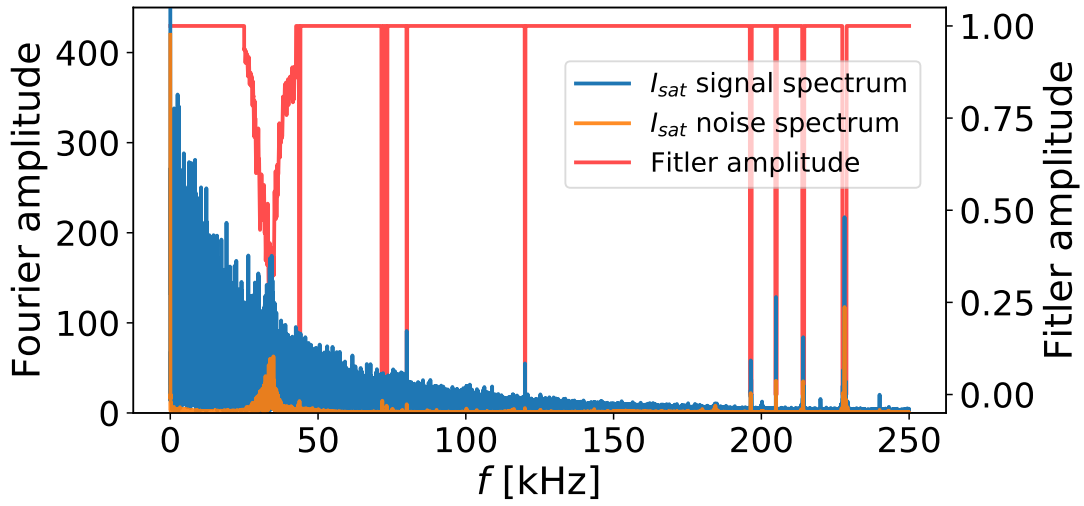

Figure 6.3: The Fourier amplitudes of the portions of the I_{sat} signal with and without filament data are shown in blue and orange respectively. The orange signal was measured before the probe reciprocation far from the plasma, while the blue signal was measured during the probe reciprocation. The red line shows the amplitude of the filter used to suppress the noise.

of the background density.

The reduction in the mean of the camera signal has the consequence that the relative magnitudes of the mean and standard deviation differ between the two signals. Therefore, a given amplitude threshold, a number of standard deviations above the mean ($\mu + n\sigma$), will capture a different proportion of events in each signal. A number of Langmuir probe analysis techniques examine only those fluctuations exceeding such a threshold. Therefore, in order to compare samples of like events measured with the two diagnostics, the threshold applied to the camera data was adjusted accordingly. The threshold was varied so as to capture the same number of events per unit time as the $\mu + 2.5\sigma$ threshold applied to the RCP data. This ensures that both diagnostics subsequent analyses consider the equivalent high amplitude events. The required threshold for the camera data was found to be $\mu + 2.0\sigma$.

Dashed green lines in Figure 6.4 indicate intensity thresholds used for each diagnostic. Peaks in intensity crossing these thresholds are indicated with crosses.

6.1.4 Conditional Averaging

As discussed in §1.6.1, conditional averaging is often used to examine the average waveforms of I_{sat} fluctuations. Figure 6.5 compares the conditionally averaged waveforms for the two diagnostics, produced by averaging time windows around peaks detected in each signal. In each case, the data is de-trended by subtracting the mean of the signal to aid comparison of the waveform shapes. Error bars show the standard deviation in the averaged intensities at each time point, in order to indicate the level of variation

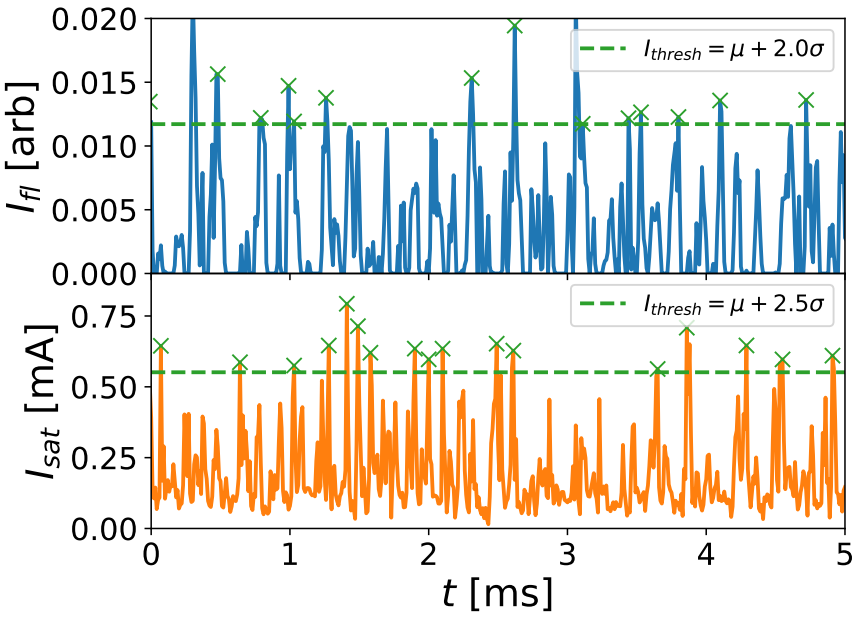

Figure 6.4: Time series of average field line intensity (blue) and RCP ion saturation current (orange) over a 5 ms time window, taken at the separatrix. Horizontal dashed lines indicate $\mu + n\sigma$ thresholds used for conditional averaging, while crosses of the same colour indicate filamentary peaks identified above these thresholds.

in the ensemble of waveforms contributing to the average waveform. The conditional averaging was performed on the full 10 ms of RCP data for which the probe was within ± 1 cm of the separatrix and on 90 ms of camera data.

The camera and RCP signals contained 29 and 262 peaks respectively over these time windows. The greater number of peaks in the camera data highlights a major advantage of the camera inversion technique. The sample size of the RCP analysis is limited by the short dwelling time of the probe at a position in the SOL. By contrast the camera simultaneously observes points over a large region of the SOL, throughout the discharge. Therefore, the length of the time window analysed with the camera is limited only by the evolution of the plasma discharge and the computational expense of inverting the data. Thus, the camera technique can analyse much longer time windows in order to build up better sample statistics for the analysis.

Both diagnostics show very similar filament waveforms that are described well by a double exponential function, a functional shape used also in several theoretical works [45, 131, 167, 168]. Exponential fits to the leading and trailing edges of the waveforms are shown as dotted lines in Figure 6.5. Both diagnostics show very symmetrical waveforms, with the RCP data exhibiting exponential rise and fall widths of 7.5 μs and 8.0 μs respectively and the camera data having rise and fall widths of 14 μs and 16 μs respectively.

The larger widths measured with the camera data are likely due to two factors.

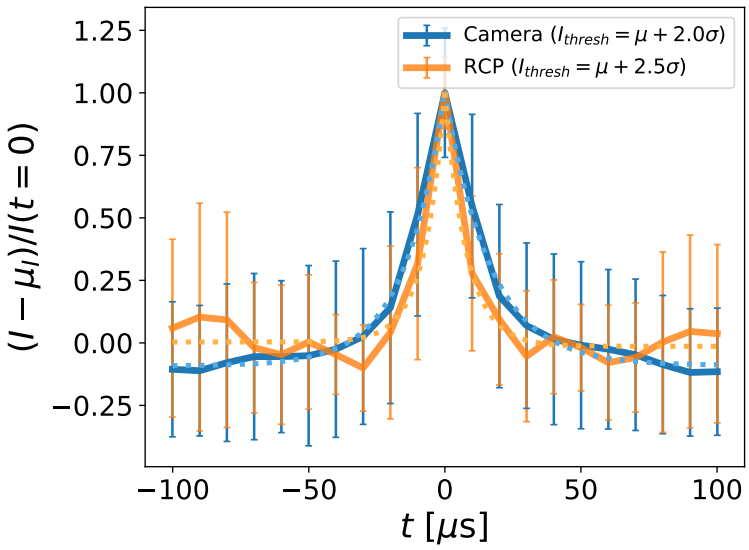

Figure 6.5: Comparison of the normalised conditionally averaged waveforms from the two diagnostics. Waveforms from the camera data are shown for thresholds of both 2.5σ and 0.8σ above the mean. Exponential fits to the rising and falling edges of the waveforms are shown with dotted lines.

First, while the RCP measurements indicate the fluctuating plasma density at a *point*, the camera measurement describes the fluctuating plasma density averaged along a *field line* passing through an equivalent point in the SOL. Thus, while the RCP will respond to high frequency local fluctuations (which may or may not arise due to filamentary structures), the camera measurement will only pick up large scale structures encompassing large parallel extents of field lines, which are thus *necessarily* filamentary in nature. Any high frequency non-filamentary fluctuations will thus tend to narrow the RCP waveform. Second, any remanent 'shadow' of the filaments from the inversion precess will tend to broaden the waveform of the camera measurements. However, the similarity of the conditionally averaged waveforms from the two techniques shows that these effects are not too large.

The strong symmetry of the waveforms observed with both diagnostics is striking. In most machines, the rise time is found to be much shorter than the fall time, corresponding to a sharp rising edge, followed by a trailing wake [48, 69, 134] (refer back to Figure 1.17). The absence of a sharp rising edge and trailing wake in conditionally averaged waveforms has previously been observed in RCP data on MAST [67, 69]. Visual inspection of the time series shows that the symmetry is a property of the underlying peaks and is not a result of averaging sharp edged filament peaks skewed in both directions. The behaviour is attributed to the stronger collisional diffusion in MAST than in other tokamaks. This means that collisional diffusion perpendicular to the magnetic field smooths the filaments and prevents the formation of sharp gradients on the leading edge. This explanation has been supported by numerical modelling in

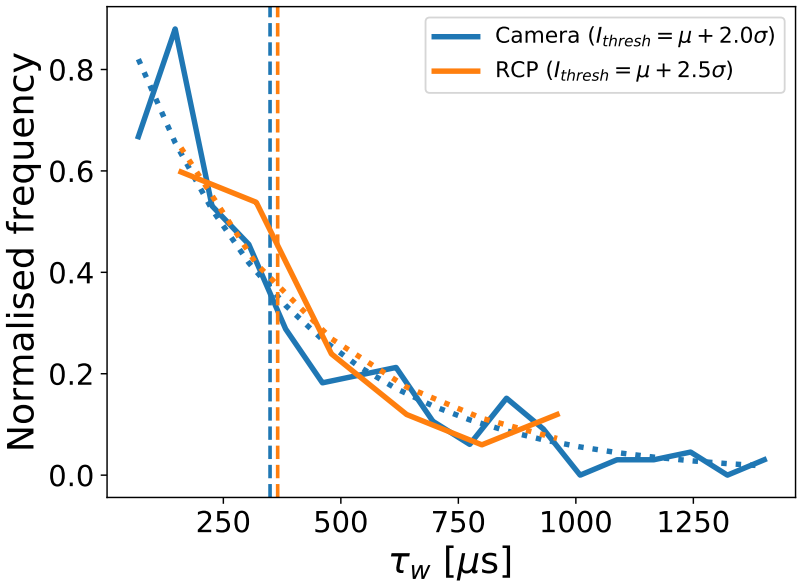

Figure 6.6: Comparison of waiting time distributions for the fast camera (blue) and RCP (orange) diagnostics. Exponential fits to the distributions are shown as dotted lines and the averages of each distribution are shown by vertical dashed lines.

ESEL [67], and now also seems to be confirmed by the camera measurements. While strong rotation of the filaments could also result in symmetrical waveforms by mixing the leading and training edges, this is not believed to be the cause. This is because parallel transport of heat is typically sufficiently strong that temperature perturbations are rarely large enough relative to the density perturbation for the filaments to be in a spinning regime [99].

6.1.5 Waiting times

Figure 6.6 shows the distribution of waiting times, τ_{wait}, between filament events observed with each diagnostic. As is typically observed [61, 160, 165, 239], they are well described by exponential distributions (dotted lines) indicating that filament generation is a Poisson process, the basic assumption of several theoretical models [45, 131, 167].

For $I_{thresh} = \mu + 2.5\sigma$, the RCP waiting time distribution gives an average waiting time of $\tau_{wait} = 365\ \mu s$. Given that the camera threshold 2.0σ above the mean is chosen so as to approximately give the same average waiting time as the RCP, the fit to the camera data gives a similar average waiting time of $\tau_{wait} = 350\ \mu$s.

6.1.6 Fluctuation statistics

The temporal scales present in the fluctuating signals can be quantified using the autocorrelation function shown in Figure 6.7. The autocorrelation time, τ_{ac}, can be extracted

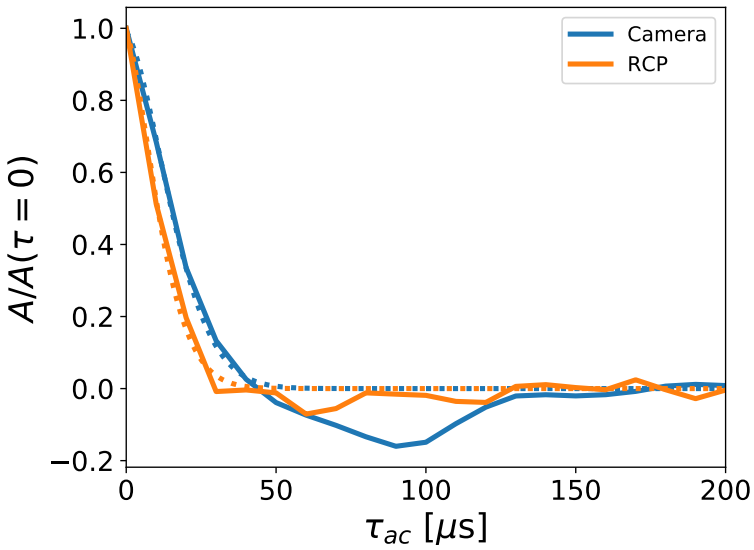

Figure 6.7: Autocorrelation functions for the intensity time series from the camera and RCP diagnostics. Dotted lines show the fits of Equation 6.1 used to extract the autocorrelation time.

by fitting the autocorrelation function with the expression [67]

$$A_{fit}(t) = \exp\left[-\left(\frac{t}{\tau_{ac}}\right)^{\beta-1}\right], \tag{6.1}$$

where the cascade index, β, describes the shape of the autocorrelation function.

As seen with conditional averaging, the camera signal has a longer auto-correlation time of $\tau_{ac} = 20$ μs compared to 14 μs for the RCP. This can be attributed to the same causes given for the broadened conditionally averaged waveform.

Figure 6.8 gives the dependence of the excess kurtosis, $K_{ex}(I)$, as a function of the skewness, $S(I)$ (see Equations 1.19 and 1.20 for definitions of S and K). As described in §2.3.4, systems governed by Poisson statistics exhibit a quadratic relation between K and S. Furthermore, it was shown in [131] that the fluctuating density signal produced by filaments whose occurrences are Poisson distributed and amplitudes are exponentially distributed should follow the specific relation given by Equation 2.17 where the y-intercept is equal to 0 (for Fisher excess kurtosis) and the coefficient of S^2 is 3/2. Figure 6.8 shows both the reciprocating probe and the camera data are well described by this relation, again supporting the assertion that filaments are independent entities, described by Poisson statistics.

6.1.7 Radial profiles

Radial profiles of the mean, variance and skewness of the camera field line intensity and RCP ion saturation current are presented in Figure 6.9, showing general good agreement between the two. The profiles from the camera data are toroidally averaged, while the

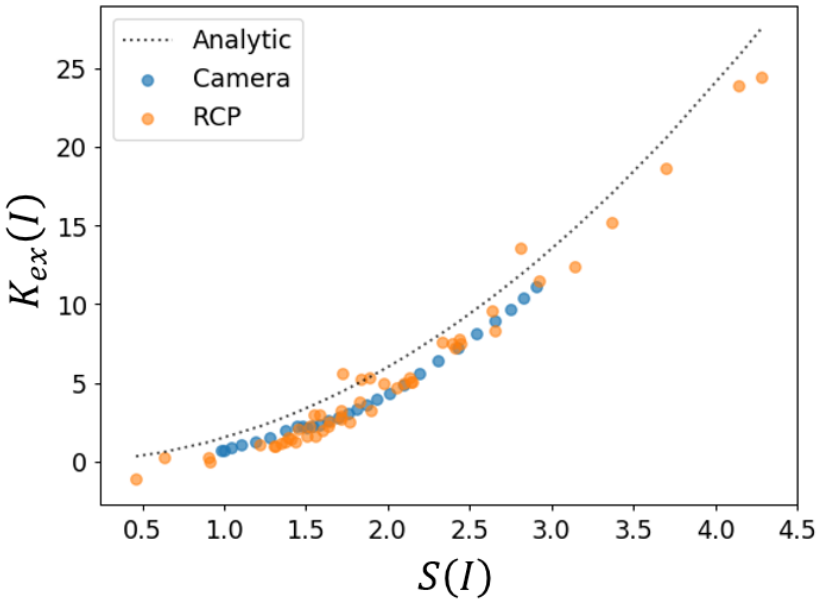

Figure 6.8: Dependence of the excess kurtosis of the fluctuating diagnostic signals as a function of their skewness. Camera data is shown in blue, while RCP data is shown in orange. The data is described well by the analytic relation $K_{ex} = \frac{3}{2}S^2$ derived in [131] for Poisson based bursts with exponential amplitudes and waiting times.

statistical moments from the RCP data are calculated over a rolling radial window of 500 sample points, corresponding a 10 ms window in time. The full sample rate RCP signal was used for this application in order to maximise the accuracy of higher order statistical moments. The radial profiles of the mean and standard deviation are normalised to their separatrix values for each diagnostic, due to the differences arising from the background subtraction of the camera data.

Both diagnostics show the typical exponential-like fall off in the mean with slight flattening in far SOL. Similarly, both show a drop in the standard deviation of the signal with radius. Both diagnostics show similar levels of skewness around ∼2 in the SOL. The commonly observed rise in skewness with increasing distance from the separatrix, corresponding to the increased relative amplitude of the filaments in the far SOL [66] is also observed in both cases. However, the RCP data is quite noisy and shows a drop in skewness beyond 3 cm. This is attributed to insufficient sample statistics in the RCP data to accurately calculate this high order statistical moment [240], due to the speed of the probe reciprocation. The drop in skewness in the far SOL could also be a consequence of interactions with poloidal field coil or the occurrence of another local plasma event during that stage of the reciprocation.

This highlights another strength of the camera technique. Each section of the radial profile from the RCP is measured at a different time in the discharge and may be subject to different local time dependent phenomena (e.g. sawtooth crashes), that cannot be avoided or filtered out of the data. By contrast the camera data can measure instan-

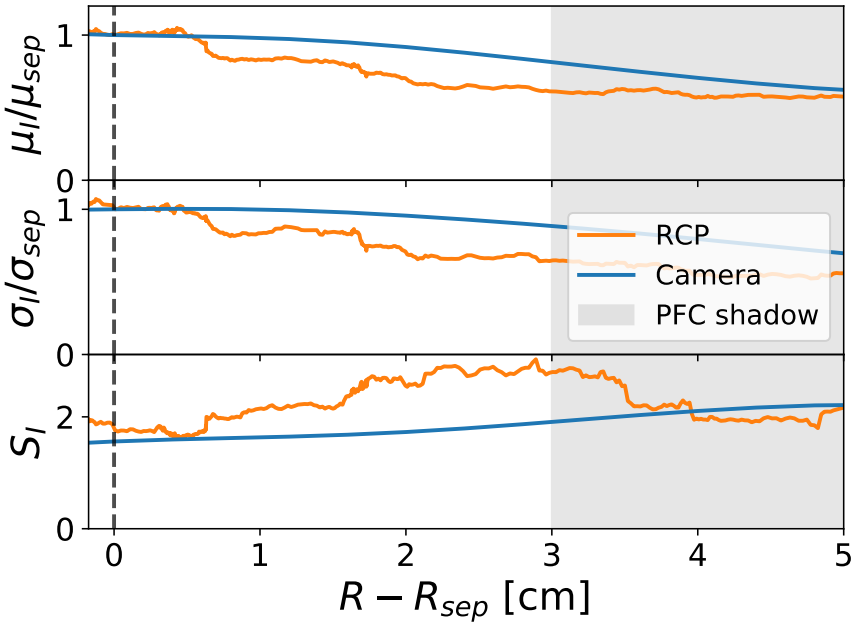

Figure 6.9: Radial profiles of the mean, variance and skewness of intensity fluctuations.

taneous radial profiles, so that all radial measurements are subject to the same large scale plasma phenomena. Furthermore, any spatially localised confounding phenomena that may influence the RCP signal, can be minimised in the camera signal by toroidal averaging.

6.1.8 Toroidal filament separation distribution

The fast camera diagnostic has the unique capability to simultaneously observe filaments over a broad range of toroidal positions. This provides the ability to measure the distribution of the toroidal separations between adjacent filaments. This is an important measurement as it provides a means of testing a fundamental assumption of analytic filament frameworks discussed in §2.3.4, which model filament emission as a Poisson process.

For filament emission to be a Poisson process, each filament must be emitted independently and uniformly around the torus, implying that filaments are not emitted with a preferential toroidal mode number corresponding to a particular toroidal separation between the filaments. Indeed, independent filaments should be produced uniformly around the toroidal extent of the machine, resulting in an exponential distribution of the toroidal filament separations. This is analogous to the uniform probability of filament emission in time resulting in the exponential distribution of filament waiting times.

The analysis proceeds in a similar manner as for the waiting times. Peaks in emission along toroidal slices through the inverted camera data shown by the orange line in

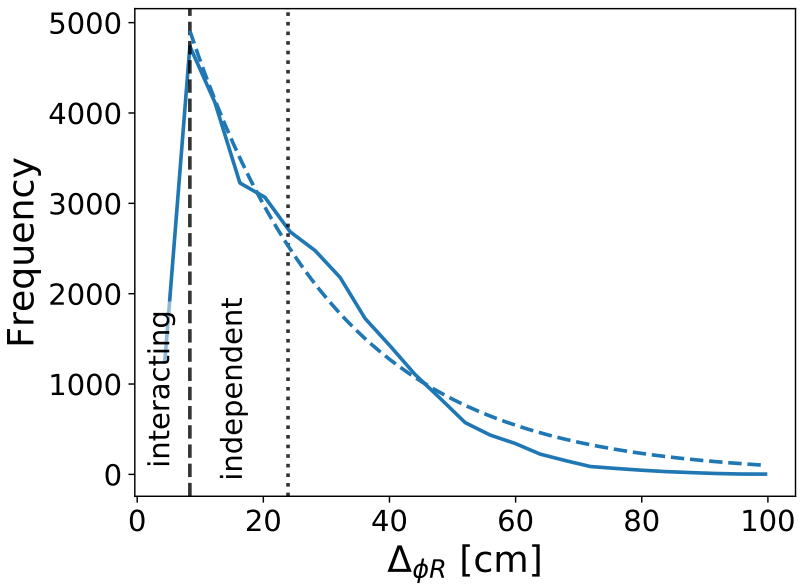

Figure 6.10: PDF of toroidal filament separations, $\Delta_{\phi R}$. For separations greater than ∼8 cm (vertical dashed line) the distribution is well described by an exponential distribution (dashed blue line) with an e-folding length of $\lambda_{\Delta\phi R} = 24$ cm. The mode and inflection point in the distribution at $\Delta_{\phi R} = 8$ cm is indicated by a vertical dashed line, while the mean of the distribution at $\Delta_{\phi R} = 24$ cm is shown by a vertical dotted line. For short separations ($\lesssim 8$ cm) there is a roll over in frequency, indicating interactions between filaments. This result was also published in [168].

Figure 6.1 are identified and the distribution of the separations, $\Delta_{\phi R}$, between the peaks analysed. Figure 6.10 shows the resulting distribution of toroidal filament separations at the separatrix. For separations $\gtrsim 8$ cm the distribution is well described by an exponential distribution with an e-folding length of $\lambda_{\Delta\phi R} = 24$ cm.

As we discuss in [168], this experimental measurement supports the assertion that filaments are randomly emitted from the core, such that the filaments are uncorrelated and uniformly distributed along the separatrix at the mid-plane, with no clear mode structure. This provides important verification of the assumptions of the analytic frameworks. This behaviour in the SOL is in contrast to instabilities in the core plasma which typically exhibit correlated modal properties.

For separations below 8 cm, indicated by the vertical dashed line, there is a roll over in frequency, indicating interactions between filaments over short distances. 2D and 3D non-linear simulations of seeded filaments with the STORM code [187] have shown that filaments separated by more than around 5 widths are effectively blind to each other's electrostatic potentials. However, filaments within a few filament widths of one another start to feel each other's presence and interact, leading to the merger of filaments and the reduction in filament frequency observed here. As shall be seen shortly, toroidal filament widths are measured to be be around 1-2 cm, consistent with

filaments interacting over length scales of less than 5-10 cm.

Given the toroidal circumference at the mid-plane is $1.6 \times 2\pi \approx 10$ m, the measured average toroidal separation of 24 cm corresponds to an average of ∼40 filaments around the toroidal extent of the machine, supporting the use of 40 filaments per frame in the synthetic data sets used in Chapter 5. These values are also in line with past measurements on MAST [51].

6.1.9 Comparison of 1D and 2D filament detections

The experimental camera measurements of filaments presented in the preceding sections have focused on identification of filaments in terms of peaks in 1D temporal or spatial intensity signals, similar to the analysis of Langmuir probe data. However, the full 2D camera inversions uniquely contain a wealth of simultaneous information about an extended region of the SOL, which we wish to utilise.

Direct analysis of the full camera inversions facilitates detailed measurements of filaments using the watershed filament detection algorithm described in §4.3. The 2D technique has the benefit of being able to simultaneously measure the radial and toroidal properties of the filaments and therefore will be the focus of the following sections.

In order to demonstrate that the 0D and 2D filament detection techniques examine the same filament structures, Figure 6.11 shows a sample of 0D time varying probe field line amplitude alongside the distance of the probe field line from the centres of the nearest filaments detected with the watershed algorithm. If the same filament is tracked in multiple frames, it is represented in the figure by the same colour.

Peaks in the probe field line intensity are observed to coincide with minima in distance to 2D filament detections. Similar behaviour is observed across longer time series. The probe field line amplitude is negatively correlated with filament distance with a Pearson correlation coefficient of $\rho = -0.29$. The relatively low correlation coefficient is a consequence of the filament detection amplitude threshold, leading to the omission of small 2D filament detections and the non-linear width profiles of the filaments. Large spikes in probe field line intensity typically occur when a 2D filament detection is within 3 cm of the probe field line, consistent with the passage of filaments with widths of a few centimetres. indicating the two approaches measure the same structures.

6.2 Application of experimental inversion data in turbulence modelling

We now discuss an application of a small number of filament measurements from the inversion technique, before moving on to discuss the analysis of large numbers of filaments.

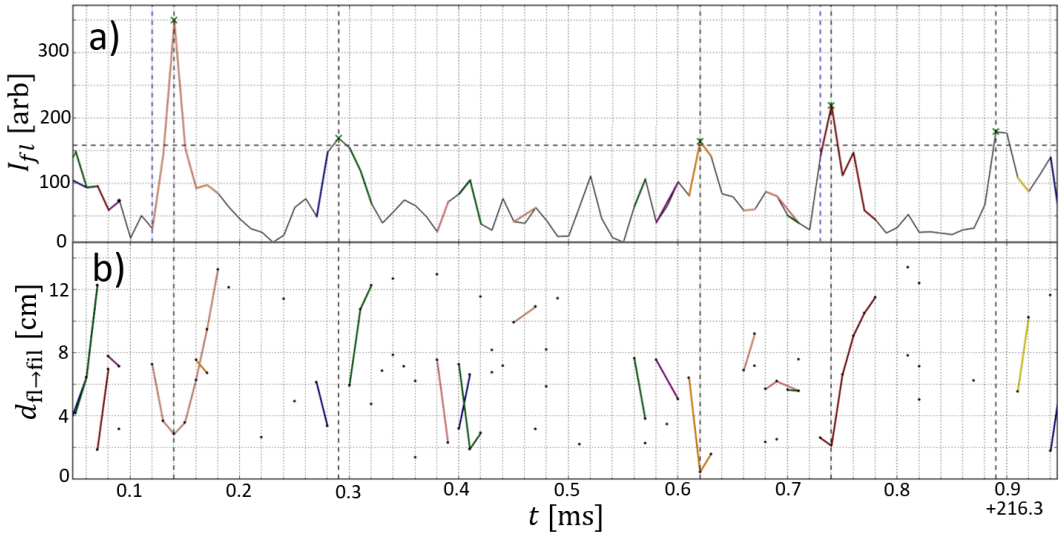

Figure 6.11: a) Time varying intensity of a probe field line as in Figure 6.4. The threshold amplitude $\mu+2.5\sigma$ is indicated with a horizontal dashed line. Vertical dashed lines and green crosses show peaks exceeding the threshold. b) Distance from the probe field line to the centre of the closest filaments detected with the 2D watershed based identification algorithm. Multiple detections of the same filament tracked between frames are highlighted with lines of a fixed colour. Peaks in the probe field line intensity are observed to coincide with minima in distance to 2D filament detections indicating the two approaches measure the same structures.

Experimental filament measurements performed with the Elzar technique have made it possible to conduct a benchmarking exercise [99] comparing the turbulence modelling codes STORM [169], GBS [180], HESEL [176] and TOKAM3X [181] with experiment. This work extended a similar code validation study in the TORPEX toroidal plasma device [185] to the more complex tokamak magnetic geometry in MAST. As each code implements slightly different equations to describe the SOL, comparing the outputs of the codes provides a means of testing the validity of different physical approximations, as well as validating the ability of the codes to reproduce experimental findings. In order to assess the performance of the codes their output needed to be compared to a reference set of experimental measurements.

The higher temperature of the tokamak SOL and more limited diagnostic access precludes probe array measurements of the type used in the previous TORPEX study. Therefore, the study was made possible by the unique capabilities of the fast camera measurements on MAST, which could provide reference filament dynamics.

Two particularly bright filaments in MAST discharge 29852 were selected as reference filaments based on their persistence over at least 4 consecutive frames, their brightness (which facilitated easier measurements), and their differing dynamics despite their similar widths and brightnesses. The filaments were first detected at times of 0.21642 s and 0.21702 s respectively. The first reference filament (referred to from now

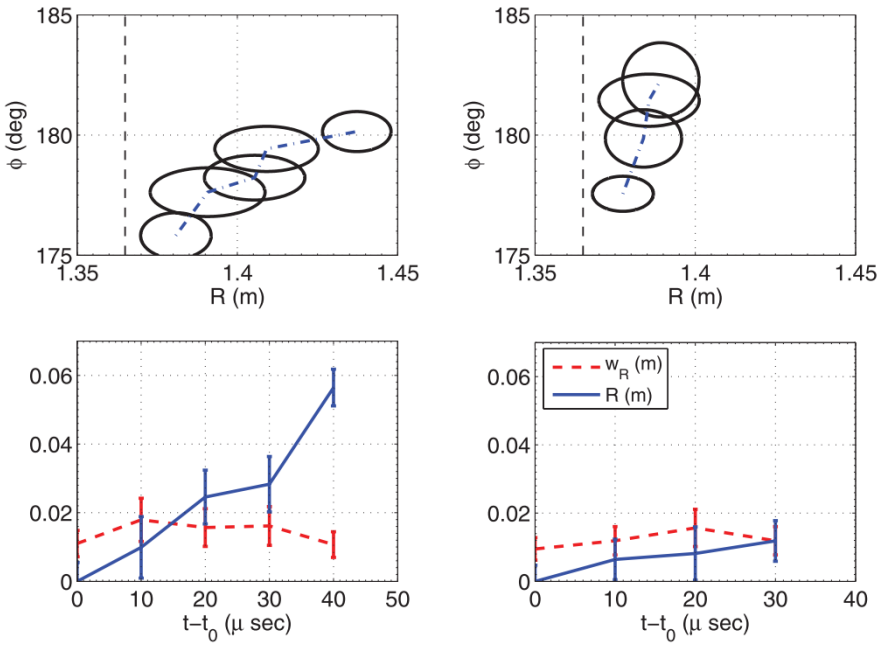

Figure 6.12: The top two panels show the positions and sizes of two filaments as measured with the Elzar inversion technique. The bottom two panels show the time evolution of the filaments' radial positions (blue solid line) and radial widths (red dashed lines) for the two filaments in the panels above. Reproduced from Ref. 99.

on as filament one) persisted for 5 frames (50 μs) and rapidly propagated across the SOL travelling radially 5 cm from the separatrix over its lifetime. Filament two by contrast had a much lower radial velocity, resulting in a total radial displacement of only 1 cm in its 40 μs lifetime. The paths of the two filaments along with the temporal variations of their radial positions and widths are shown in Figure 6.12.

In addition to the insight gained into the relative capabilities of the turbulence codes discussed in §2.4.4, the comparison of the experimental data with the modelling data provided valuable insight into the experimental measurements. As the intensity of light emitted by a filament depends on both its density and temperature, according to Equation 1.21, such that in the absence of a secondary measurement with a different dependence on density and temperature, the relative contributions of the density and temperature to the observed emission are unknown. A second approach to identifying the relative contributions of the density and temperature to a filament's pressure perturbation is to simulate the experimentally observed filament with different possible temperatures and densities and identify which parametrisation best reproduces the dynamics observed in experiment.

The magnitude of the pressure perturbation in each filament was estimated by comparing the intensity of the filaments in the background subtracted images to the intensity of the background emission, whose corresponding temperature and density was known from high resolution Thomson scattering (HRTS) measurements. With the rela-

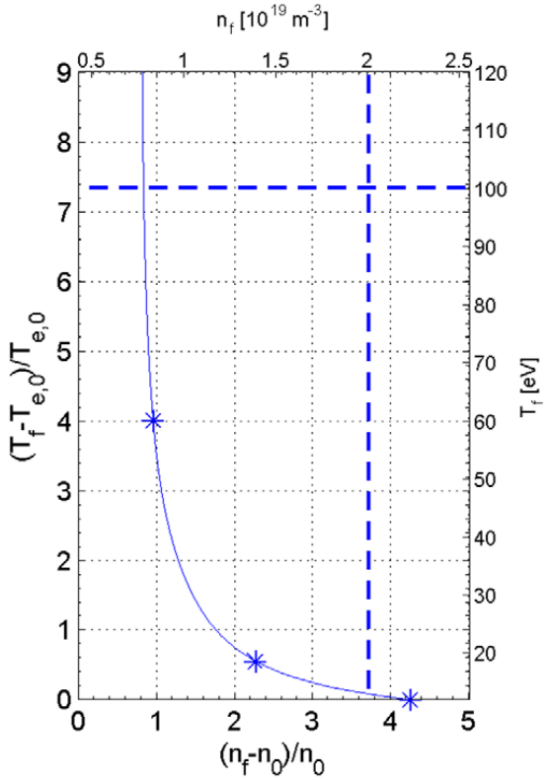

Figure 6.13: Emission curve for filament 1, representing the combinations of filament density, n_f, and temperature, T_f, required to reproduced the filament's measured light intensity. $T_{e,0}$ and n_0 represent the background temperature and density in the SOL respectively. Dashed lines mark the density and temperature values 7 cm inside the separatrix. The asterisks show the combination of n_f and T_f used in the simulations. Reproduced from Ref. 99.

tive pressure of each filament identified relative to the pressure of the background SOL plasma, the size of the density and temperature perturbation in each filament was constrained to a curve in $\delta n - \delta T$ space. The curve for filament 1 is shown in Figure 6.13. The simulations were then performed at three possible points in this phase space and the dynamics of the simulated filaments compared to those observed with the camera. In order to perform a like for like comparison, the synthetic camera diagnostics was used to forward model the output of the turbulence codes so that the motions of the filaments could be analysed consistently with the Elzar technique. The ADAS atomic physics library [241] was used to calculate the light emission for the output distributions of simulated temperature and density.

The STORM and GBS codes were able to produce behaviour most consistent with experiment. Both the radial and toroidal motions of the filaments were examined to find the best parametrisation of the pressure perturbation in each filament. The toroidal velocity of the filaments includes both their intrinsic toroidal motion and the bulk rotation of the plasma in which they are carried. By comparing the relative toroidal

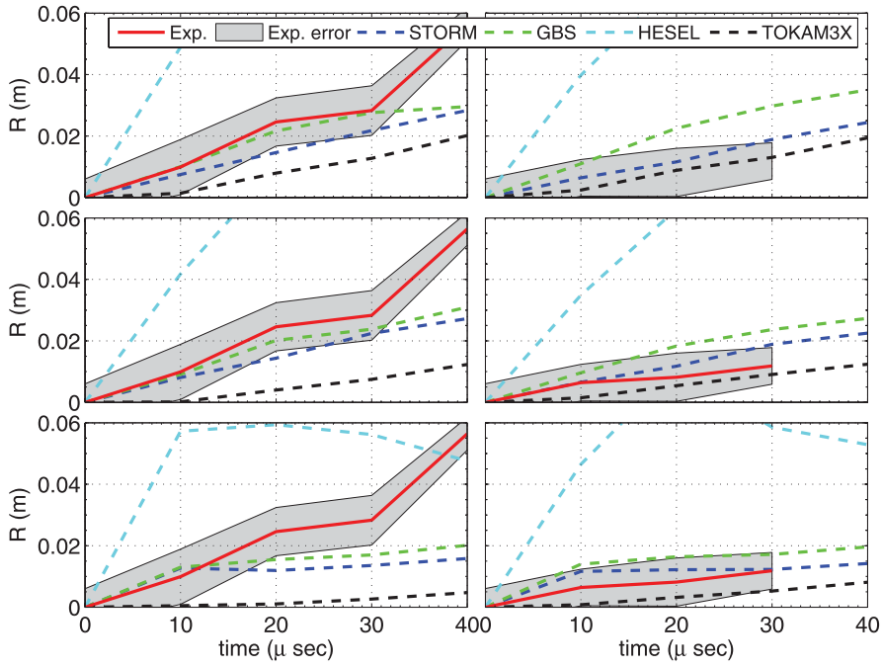

Figure 6.14: Comparison of the experimentally observed (red and grey) radial motion of filament one (left) and filament two (right) with that predicted by each of the four codes (dashed lines) for each of the three simulated parametrisations of the filaments' pressure perturbations. The top, middle and bottom rows of panels correspond to pure density filament perturbations, half the peak density perturbation and primarily temperature perturbations respectively. Filament one (left) is best described by the STORM and GBS codes for an intermediate density perturbation (middle), while filament two (right) is best reproduced for a primarily temperature perturbation (bottom). Reproduced from Ref. 99.

motion of the two filaments in the simulations the intrinsic filament motion could be separated from the bulk plasma rotation. In this way, the intrinsic toroidal rotation of the filaments were found to be 0.2âĂŞ1 km s^{-1}, while the mean SOL plasma rotation was found to be ~ 3 km s^{-1}. Figure 6.14 compares the experimentally observed radial motion of the two filaments to the simulated radial motion in each of the three simulated parametrisations of the filaments pressure perturbations.

The STORM and GBS codes were able to best reproduce the experimentally observed radial and toroidal motion for filament 1 with a pressure perturbation primarily due to density, corresponding to a density perturbation over two times the background density and a temperature perturbation around 50% greater than the background. Filament two's behaviour by contrast was best reproduced for a density perturbation just below 1 times the background level and a temperature perturbation around 3.5 times the background temperature.

This showed that the slower radial motion of the second filament could be attributed to a larger relative temperature contribution to its pressure perturbation than for the first filament. This is unsurprising as large relative temperature perturbations ($\delta T/T_0 \gg$

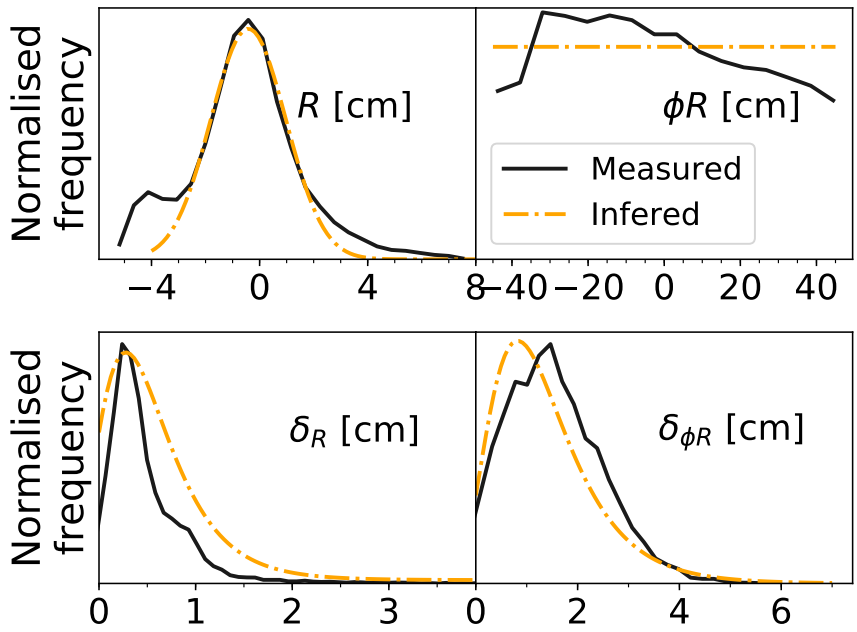

Figure 6.15: Distributions of filament parameters for experimental data from MAST shot 29852, t=[0.205,0.245] s, with $\epsilon_{thresh} = 1.5 \times 10^{-2}$. Solid black lines show the measured distributions with the systematic error offset correction applied. Dashed orange lines are sketches of the inferred most probable true distributions shapes, given the insight gained from the synthetic analysis.

$\delta n/n_0$) drive the formation of competing monopole potential structures, that suppress the radial diamagnetic drive of the filament by increasing its intrinsic bulk rotation and breaking up the filament's electric dipole. In this way, the experimental camera analysis, combined with the synthetic camera diagnostic, provided a unique means of benchmarking the turbulence codes, while the codes were able to provide additional insight into the experimental measurements.

6.3 Experimental measurements of filament statistics

6.3.1 Parameter distribution functions

In this section, experimental measurements of filament parameter PDFs equivalent to those performed on synthetic data in benchmarking exercise in Chapter 5 are presented. The full Elzar filament analysis technique is applied to MAST discharge 29852. The analysis was performed on 4000 frames, over the time window 0.205-0.245 s. The filament detection amplitude threshold was set to the optimal value identified in Chapter 5 of $\epsilon_{thresh} = 0.015$ and the systematic offset corrections applied as determined by the synthetic filament study.

Figure 6.15 shows the measured experimental distributions of filament positions and widths, obtained with the new inversion technique, in black. Given the findings from the

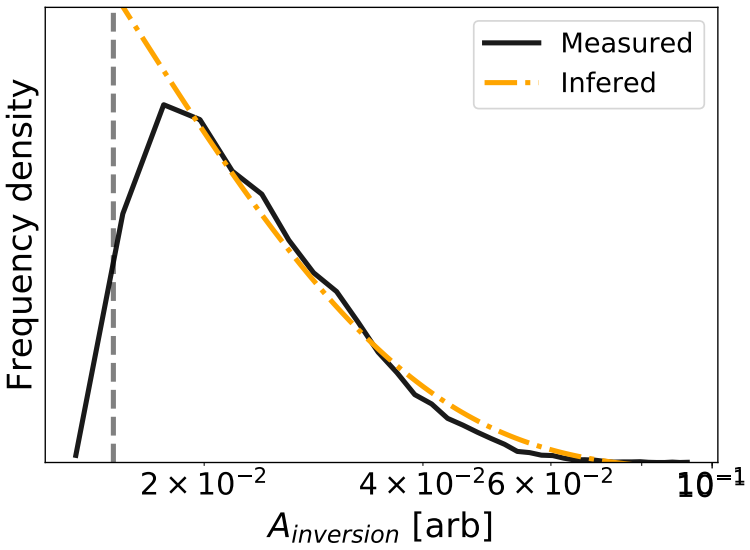

Figure 6.16: Distributions of filament amplitudes for experimental data from MAST shot 29852, t=[0.205,0.245] s, with $\epsilon_{thresh} = 1.5 \times 10^{-2}$. The solid black line shows the measured distribution. The dashed orange line is a sketch of the inferred most probable true distributions shapes given the insight gained from the synthetic analysis. The vertical dashed line indicates ϵ_{thresh}.

synthetic data analysis (see Figure 5.8), orange dashed lines have been added to sketch, without any pretence of rigour, possible inferred distributions for each parameter.

A bi-modal distribution is recovered for the R positions, with the peak at low major radius likely due to ghost filaments which have their highest relative frequently at low R. The primary peak is described well by a log-normal distribution as observed previously and assumed for the synthetic data sets. The peak in detections occurs around the separatrix where both the number of filaments and the neutral particle density is high. Inside the separatrix the detection rate falls off sharply due to the rapid fall off in neutral density, given the neutral ionisation mean free path at the separatrix is ~ 1 cm.

As far as the toroidal distribution is concerned, a reduced detection rate is observed at large ϕR which is significantly stronger than is seen in the synthetic data set, but is still believed to be a diagnostic artefact from line of sight effects. The region of increased detection density at low ϕR does not coincide with the highest region of false detections in the synthetic data, so it is likely the increased detections at low ϕR are a result of greater sensitivity rather than increased false positives.

Both radial and toroidal width distributions have similar shapes to the those measured from the reference synthetic data set, indicating the widths are well described by log-normal distributions.

Figure 6.16 shows the measured distributions of filament amplitudes for the experimental data set in black, again with an orange dashed line added to sketch out the possible distribution that can be inferred from the synthetic data analysis. An

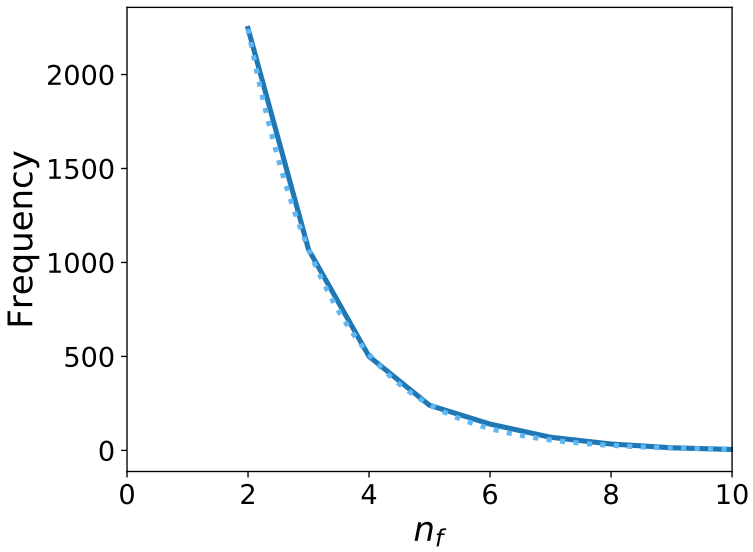

Figure 6.17: Distribution of filament lifetimes, measured in terms of the number of consecutive frames, n_f, over which each filament is observed. An exponential fit to the distribution is given by the dotted line.

exponential amplitude distribution is observed with a roll over around ϵ_{thresh}, similar to that observed for the reference synthetic dataset. This observation of an exponential amplitude distribution is in agreement with previous measurements with Langmuir probes [66, 90] and GPI [165].

6.4 Filament dynamics

We now present measurements of filament dynamics using the filament tracking technique described in §4.4.

6.4.1 Filament lifetimes

Of the 19654 individual instantaneous filament detections over the analysis interval, 0.205-0.245 s, in discharge 29852, 4324 were identified as belonging to chains of observations of filaments across two or more frames.

The distribution of the number of frames, n_f, over which a filament is observed is shown in Figure 6.17. The longest lived filament was observed to persist for 14 frames, although only one such example was observed in the data set. Lifetimes of 11, 12 and 13 frames were also very rare, with 5, 6 and 0 observations respectively. The distribution is described well by an exponential distribution with with an e-folding length of $\lambda_{\tau_{life}} = 1.35$ frames. Therefore, the average filament can be expected to have a lifetime corresponding to 1.35 frame intervals.

Given the frame rate of the camera was 100 kHz, the interval between frames is $\Delta t = 10$ μs. (For simplicity of the discussion, we shall approximate the frame acquisition

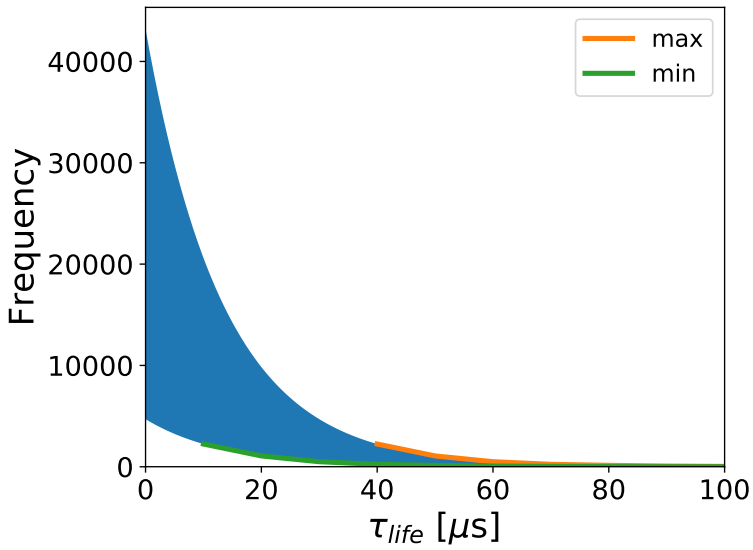

Figure 6.18: Envelope indicating the bounds placed on the distribution of filament lifetimes (blue) Orange and green curves represent the bounding curves produced by interpreting the curve in Figure 6.17 as if all detections were due to filaments with their bounding maximum and minimum lifetimes respectively.

as being instantaneous, while in reality the integration time was 3 μs followed by 7 μs of sensor dead-time.) For a filament to be observed across n frames, its minimum lifetime must be $(n-1)\Delta t$ or it would not have survived long enough to be seen in the bounding frames. Equally, its maximum lifetime is $(n+2)\Delta t$, or it would have been observed in one of the neighbouring frames. Therefore, lifetime of a filament observed in n consecutive frames can be constrained to the interval $\big[(n-1)\Delta t, (n+2)\Delta t\big)$.

Figure 6.18 shows the envelope of possible lifetime distributions that can be inferred from this data. The orange and green curves represent the bounding curves produced by interpreting the curve in Figure 6.17 as if all detections were due to filaments with their bounding maximum and minimum lifetimes respectively.

6.4.2 Filament velocity distributions

With the chains of consecutive detections assigned to filaments with differing lifetimes, we now calculate the average radial and toroidal velocities of each filament. Figure 6.19 shows an example of linear fits applied to chains of detected filament positions in order to calculate their average radial and toroidal velocities. The toroidal velocity of the filaments is quite constant over their lifetimes, so that linear fits form a good approximation of their overall motions. This is in large part expected, as a large component of the filaments' toroidal velocities is due to the bulk rotation of the plasma, which varies relatively slowly on the temporal and spatial scales of the filaments. The radial velocities show more complex structure but are none the less relatively constant over their lifetimes. As discussed in Chapter 5, the slower radial motion of the filaments relative

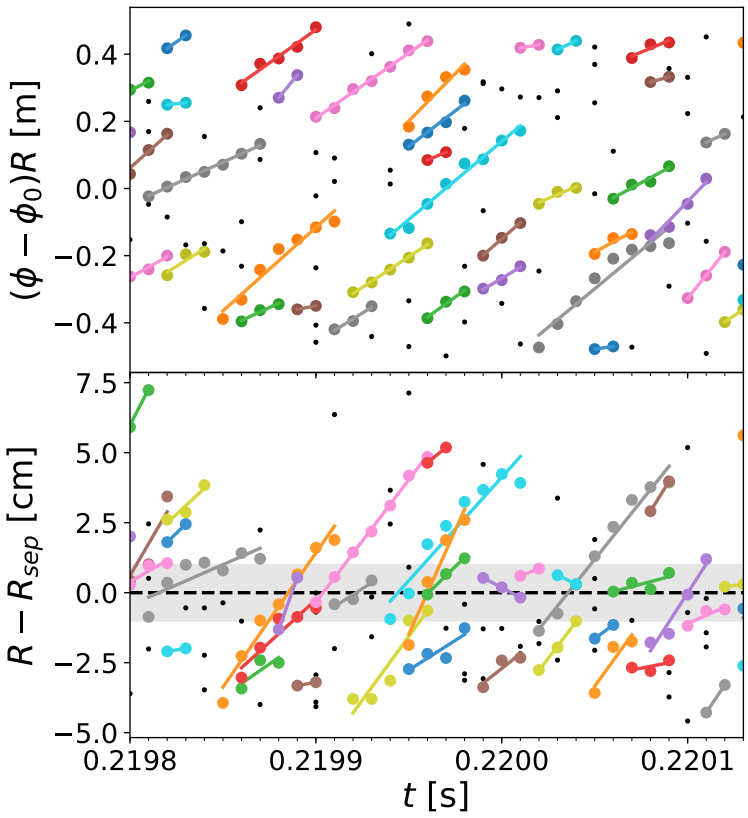

Figure 6.19: Toroidal (top) and radial (bottom) locations of filament detections as a function of time. Filament detections assigned to the same chain are given the same colour. Linear fits to each chain of filament detections, used to calculate the average toroidal and radial velocities for each filament, are shown as straight lines of the same colour. Isolated filaments that are only observed for one frame are shown in black. The position of the separatrix and it's uncertainty are shown by the black dashed line and shaded region.

to their toroidal motion leads to a greater fractional error in their radial velocities, despite the errors in the radial and toroidal positions both being around 2 mm. Therefore, fractional errors on the instantaneous (frame-to-frame) radial and toroidal velocities are estimated to be around 40% and 7% respectively. Thus, by treating the radial velocities as also being approximately constant and calculating their average velocities with linear fits, the errors on the radial velocities can be reduced by an approximate factor of $\sim\frac{1}{\sqrt{n_f}}$.

Following this rationale, Figure 6.20 shows the distributions of the filament averaged radial and toroidal velocities. The mean and mode of the radial filament velocities are 0.73 km/s and 0.55 km/s respectively, while the corresponding mean and mode of the toroidal filament velocities are 3.8 km/s and 3.3 km/s (vertical dashed and dotted lines). Both distributions are described well by log-normal distributions. The log-normal distribution describes positive quantities whose values arise from the accumulation of many

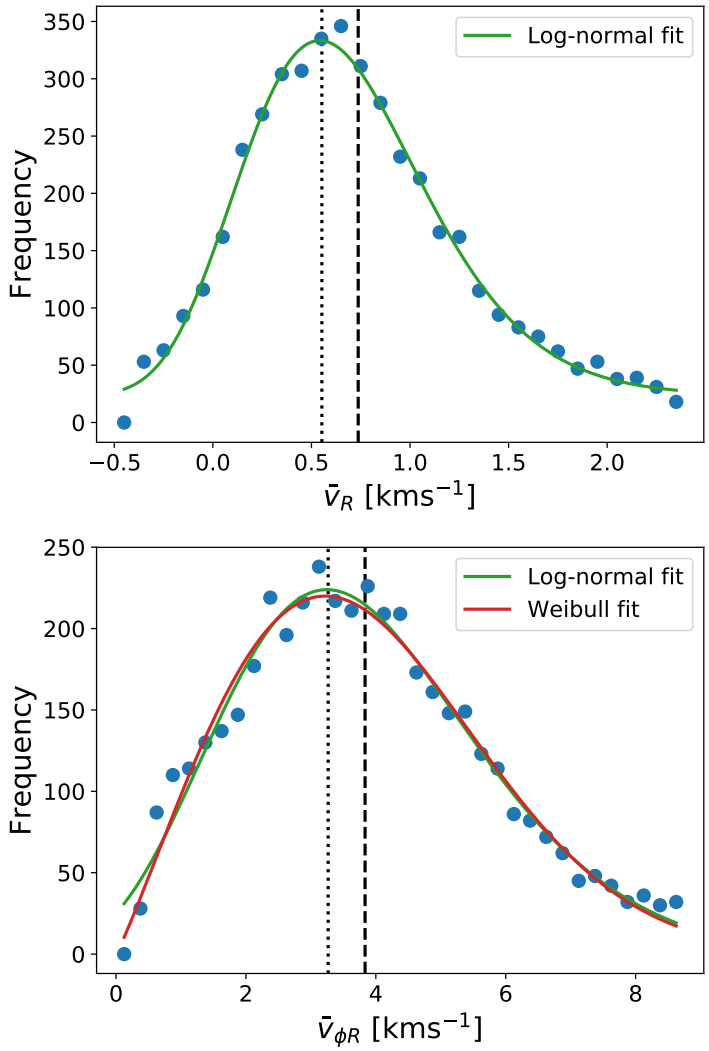

Figure 6.20: Distributions of average radial (top) and toroidal (bottom) filament velocities. Log-normal distribution fits are shown in green and the Weibull distribution fit to the toroidal velocities in red. The mean and mode of each distribution are indicated by dashed and dotted lines respectively.

small fraction changes, and is commonly observed in many physical systems where a quantity is positively defined but finite. The PDFs of radial and toroidal velocities agree well with past measurements from MAST presented in [51].

Interestingly, the toroidal velocity distribution is equally well described by the Weibull distribution. The Weibull distribution describes the magnitude of a vector quantity resulting from the combination of two independent, Gaussian distributed orthogonal components. For this reason, it is often used to describe distributions of wind velocity for which the northerly and easterly components are independent and normally distributed. However, the Weibull distribution only describes positive quantities and therefore does not describe the radial velocities as well, due a small proportion of

filaments with negative radial velocities that result in a tail at the lower end of the radial velocity distribution. The physical interpretation of this result, in the context of filaments, requires further investigation.

6.5 Conclusions

The Elzar filament analysis technique has been applied to experimental data from the MAST tokamak. Comparison of reduced dimensionality camera inversion data with reciprocating Langmuir probe data from a similar discharge has shown the camera data can be used to reproduce statistical findings from Langmuir probe data.

In particular, time series intensity fluctuations from camera and RCP diagnostics both show the same symmetrical conditionally averaged waveforms, exponential waiting time distributions and properties of the higher order moments of the intensity fluctuations. Uniquely, the fast camera measurements are able to examine the toroidal filament separation distribution, indicating that the filaments satisfy the independence criteria required for application of existing analytical filament frameworks and consistent with the distribution of the waiting times.

Experimental measurements combined with the synthetic camera diagnostic discussed in the previous chapter have enabled a multi-code validation exercise of four turbulence codes. In addition to validating the codes, this work provided additional insight into the pressure factorisation of the observed filaments, showing filaments with slower radial velocities are likely to have higher relative temperature components contributing to their pressure.

Finally, experimental PDFs of filament parameters have been produced, showing the distributions of filament positions, widths, amplitudes and velocities. A number of these measurements are interpreted with the insight gained from the benchmarking exercise in Chapter 5. The radial and toroidal velocity distributions are found to be described well by log-normal distributions, with the toroidal velocity distribution also well represented by the Weibull distribution.

Chapter 7

Conclusions and Further Work

7.1 Overview of research outcomes

The research undertaken during the course of this work can be split into four main areas:

1. The development of the Elzar camera image inversion and filament detection technique.

2. The development of the Elzar fast camera analysis code in preparation for the automated large scale analysis of fast camera data from MAST and MAST-Upgrade.

3. The benchmarking and error assessment of the Elzar filament analysis technique using a synthetic camera forward model.

4. Application of the Elzar filament analysis technique to experimental data from MAST, including comparison with Langmuir probe data.

7.1.1 Development of the Elzar technique

A new inversion technique has been developed to map field aligned emission in unfiltered fast camera images onto a field aligned basis set for the analysis of filamentary plasma structures, which are now understood to be the 'quanta' of particle transport in the SOL.

By inverting the emission in the camera images onto a field aligned basis set, the quantity of information that can be extracted from the images is greatly increased in comparison to previous to previous techniques. The inversion makes it possible to simultaneously measure the positions, widths, amplitudes and velocities of filaments over a large region of the SOL. By contrast, past fast camera analysis techniques have tended to only measure one filament property at a time due to relying on the particular viewing properties of different regions of the images. Furthermore, the automation of the inversion analysis technique gives it the potential to be applied readily to large quantities of camera data.

Fast camera diagnostics have range of capabilities that make them particularly favourable for filament measurements. The nature of fast camera imaging enables the passive analysis of a breadth of plasma regimes, without restriction from required, potentially perturbative probe interactions, or gas puffing needed by other diagnostics. The diagnostic also collects filament data from a large region of the SOL both facilitating collection of larger datasets and further investigation of spatial dependencies, particularly toroidal dependencies, of the filaments properties. Finally, the 3D nature of the measurements has unique potential for analysis of the parallel structure of filaments.

A new technique has also been developed for the identification and measurement of filament properties based on the watershed algorithm and the fitting of 2D Gaussian functions to identified regions of enhanced density. This technique has been found to be much more flexible and robust than other approaches that were explored. The positions of the identified filaments are input into a tracking algorithm that identifies chains of detections in consecutive frames corresponding to the same filament over its lifetime. This facilitates the study of the velocities of the filaments time dependent dynamics of the filaments.

7.1.2 Benchmarking of the Elzar technique with synthetic data

The newly developed technique has been extensively benchmarked through analysis of forward modelled camera data produced using a synthetic camera diagnostic. By analysing synthetic camera images composed of synthetic filaments with experimentally representative distributions of parameters, the amplitude threshold for filament identification was tuned to optimise the precision and sensitivity of the detections. For the optimal threshold, the benchmarking exercise indicated the Elzar technique is capable of detecting 74% of the target filaments (i.e. those with amplitudes above the threshold), corresponding to 36% of the total filament population, while maintaining a precision (proportion of true positive detections) of 98.8%.

Fits to the distributions of errors in the measured parameters revealed standard errors in the radial and toroidal position measurements of around 2 mm and standard errors of around 3 mm and 7 mm respectively in the radial and toroidal widths. However, the kurtosis of the error distributions was high indicating that large errors of the order 3σ are more common than would be expected for purely Gaussian errors, but still rare.

Systematic errors implicit in the inversion technique, with magnitudes typically around twice that of the standard errors, were also identified which can be offset from future measurements. These error estimations provide assessments of the accuracy to which the different filament parameters can be recovered, and insight into artefacts inherent to the inversion technique. They are best case estimates due to the fact the analysis of synthetic data does not account for inaccuracies in the magnetic equilibrium and the camera calibration (discussed shortly). However, these additional sources of

error are not expected to change the overall magnitude of the errors.

Finally, the ability of the technique to reproduce the PDFs of the input filament parameter distributions was assessed. The technique was found to successfully qualitatively reproduce the true input distributions and enabled the identification of deviations from the input distributions that can be used to better inform the interpretation of experimental measurements.

7.1.3 Application of the Elzar technique to experimental data

The Elzar technique has been applied to experimental data and a direct comparison with RCP data performed in order to further benchmark the camera technique against well established Langmuir probe measurement techniques. While the background subtraction and regularisation steps of the camera analysis result a differing ratio between the mean and standard deviation of the signals from the two diagnostics, this can be compensated for with a suitable adjustment of threshold parameter. In general, good agreement is found between the two diagnostics. Both diagnostics show the same symmetrical conditionally averaged waveforms and similar exponential waiting time distributions. Both datasets also exhibit the same statistical properties in the higher order moments of the signals.

These measurements independently verify a range of past findings with Langmuir probes and demonstrate that the Elzar technique is equally capable of investigating the statistical properties of filaments. However, the nature of the camera measurements make the Elzar technique far more flexible as it is capable of making simultaneous measurements of radial and toroidal profiles throughout the duration of a discharge, in contrast to the single time series of 0D measurements made by a Langmuir probe. This enables much larger sample statistics to be collected and enables data corruption from transient events such as sawtooth crashes to be avoided. Furthermore, the Elzar technique provides the unique capability to measure the toroidal distribution of the filaments. This has been found to be broadly exponential indicating that the filaments are generated uniformly around the toroidal extent of the machine with no common toroidal mode number and are thus generated independently. This, along with the observation of exponentially distributed waiting times is strong evidence that the generation of filaments is a Poisson process. This is an important property to verify, as it validates the fundamental assumptions of analytic frameworks [45, 131] which enable the interpretation of SOL profiles and their statistical properties in terms of the statistical properties of filaments.

Peaks in time series signals extracted from camera inversions, which are identified as filaments in the Langmuir probe style analysis, are observed to correspond to filaments identified with 2D filament detection technique, with peaks in the time varying signals coinciding with the closest approach of filaments to the probe field line. This

demonstrates that the 0D Langmuir probe style analysis and 2D inversion approaches measure the same objects.

Fast camera measurements of filaments in MAST have made it possible to perform a verification exercise of four turbulence codes in tokamak relevant simulations. In addition to providing insight into the the capabilities of the codes, the simulations were able to better inform the experimental measurements, indicating that slower filaments observed with the fast camera may have larger relative temperature perturbations than faster neighbouring filaments of equal brightness.

Analysis of experimental data with the full Elzar filament detection technique has provided valuable measurements of the PDFs of filament widths, amplitudes and velocities. These have been interpreted with greater insight as a result of the synthetic benchmarking exercise. These and future measurements will provide useful constraints on the distributions of free parameters used in analytical models of filamentary SOL transport.

7.2 Further work

Now that the fast camera inversion and filament detection technique has been developed and benchmarked, there are a great many avenues for future work and application of the technique.

7.2.1 Analysis of archived MAST data

The work in this thesis has focused on the detailed analysis of mid-plane view camera data for two MAST discharges. However, archive mid-plane fast camera data from MAST exists for 190 discharges from the M9 campaign in 2013, of which around 76 have useful data of varying quality. Large quantities of past camera data can now be analysed with the technique, in order to compose a large database of filaments and their properties. This should provide the breadth and quantity of data necessary in order to perform detailed statistical analyses of filament properties as a function of machine and physics parameters, in order to understand their roles in shaping SOL profiles.

7.2.2 Analysis of MAST-Upgrade Data

MAST-U is scheduled to start it's first experimental campaign at the start of 2020. It will be equipped with two Photron SA-X2 fast framing cameras which should enable wide angle imaging of the plasma at 200 kHz. MAST-U studies will provide an excellent opportunity to understand the influence of alternative divertor configurations on filamentary transport. While the cameras will be collecting data for all the discharges, the design of dedicated experiments will enable direct comparisons between the fast cameras and other diagnostics in the same discharge, as well as systematic analysis of

different operating regimes.

A number of challenges will be faced on MAST-U. While operations in conventional divertor configurations should be very similar to those on MAST, the baffle is predicted to significantly reduce the main chamber neutral density in the super-X magnetic configuration, leading to a concomitant decrease in D_α intensity. It is believed is will be possible to offset this effect through a combination of a larger collection lens, longer integration times, reduced field of view and the higher sensitivity of the SA-X2's sensor. If necessary, weak outboard mid-plane fuelling could be used, or the frame rate can also be dropped to increase the exposure time further, at the expense of filament tracking capabilities. However, this is unlikely to be necessary.

The technique could also be exported to other machines, provided that good equilibrium reconstruction and sufficiently bright filaments are present.

7.2.3 Further error quantification and benchmarking

Magnetic equilibrium and camera registration inaccuracies

A limitation of the synthetic error analysis in Chapter 5 is that the synthetic images are produced and analysed with the same magnetic equilibrium and camera calibration descriptions, which ensures a strong mapping between the inversions and synthetic images. The magnetic equilibrium and camera calibrations will, by contrast, not be exact mappings in the case of experimental data. Thus, the synthetic analysis indicates how well the technique can perform with optimal equilibria and camera calibrations. Any inaccuracies in these inputs will lead to further inaccuracies in the inversions and inferred filament properties. Errors in the magnetic equilibrium could be emulated by using magnetic gfiles from different portions of a shot, or by perturbing the pressure profiles used to generate the gfiles. Errors in the camera calibrations could be explored by generating different camera registrations with perturbed input point pairs.

Influence of impurity species

A further source of unquantified error, that could have a small effect on the width measurements of filaments, is emission from impurity ion species. Since the camera is operated without a filter, impurity emission, primarily from carbon, will be recorded that has differing spatial variation to the primary D_α emission. With the exception of the centre column (which has a low surface area) and the P4 and P5 poloidal field coils, there are no first wall surfaces close the mid-plane plasma in MAST. Therefore large carbon releases from the walls in the main chamber are not expected and the assumption of dominant D_α emission should be valid. However, in order to fully justify not using a filter more sophisticated spectral forward modelling could be performed using the CHERAB synthetic diagnostic tool [242]. CHERAB performs spectrally resolved ray tracing and should be able to quantify these effects.

Improvements to the synthetic camera diagnostic

While the combination of field aligned emission and noise used in the current synthetic camera diagnostic captures the most important features of experimental background subtracted camera data, further features could be added. Parallel variation along field lines, alternative basis functions and techniques to avoid potentially unphysical filament overlap could be introduced. Filament tracking techniques could also be benchmarked by evolving the synthetic filaments forward in time using the principles of the analytic filament framework. Furthermore, the fidelity of the forward modelling could be improved by using the more sophisticated CHERAB forward modelling tool, although this would be accompanied by a large increase in complexity and computational time.

Improvements to the benchmarking parametrisation

Some well established methods exist in the machine learning community for optimising the performance of detection algorithms by assessing the inter-relation between the sensitivity, precision and specificity of an algorithm.

In particular the Receiver Operating Characteristics (ROC) is accessed by plotting the proportion of true positives (sensitivity) against the proportion of false positives (1-specificity) and measuring the area under the curve (AUC) [237]. In an ideal case the AUC is zero indicating the model can perfectly separate true and false positives. Conversely, an AUC of 0.5 indicates the model has no capability to discriminate between true and false positives. In additional to providing a figure of merit for the model, the curve can be used in a cost-benefit analysis, with the top left most region of the curve indicating the most favourable operating space. This technique should be applied alongside the curve optimisation technique shown in Figure 5.5 to select the best possible filament detection threshold.

7.2.4 Improvements to the inversion algorithm

In this work, the optimal regularisation level was chosen by eye by minimising the level of distortion and banding in the inversions, while avoiding over smoothing of the filament structures. While this approach is sufficient, it would be valuable to establish a more quantitative method of refining the regularisation strength parameter. Therefore, future work should explore the application of techniques such as the L-curve and Generalized Cross-Validation methods described in [243, 244], to objectively select the regularisation strength.

Alternative regularisation schemes should also be explored in more detail. In particular, using a minimum Fisher information functional [218] in Equation 4.11 is promising, as it modifies the local level of smoothing depending on the local value of the solution and also naturally includes a positivity constraint. This enables the peaks in emission to be smoothed less, potentially reducing over-smoothing from regularisation. However,

the possible increase in accuracy of this technique is accompanied by the possibility of overestimating peak intensities [206]. Therefore it introduces a further element of uncertainty with respect to other regularisation techniques that are known to only ever over-smooth. For this reason, it is avoided in some applications [206]. Minimum Fisher information regularisation has been applied successfully in X-ray tomography [218] and so is a promising technique for investigation here.

More sophisticated non-negativity constrains should be explored. In particular the Accelerated Projected Conjugate Gradient (APCG) technique described in Ref. [224] is a promising technique that uses the non-negativity of the solution to better inform the least-squares gradient decent method, potentially producing faster convergence and smoother solutions.

Using a non-regular inversion grid with a sparse set of broadened field lines extending the inversion domain may also hold potential for improving the results, by better constraining the emission currently outside the inversion domain.

7.2.5 Filament identification with neural networks

For future development, the use of convolutional neural networks for identification of filaments in inverted camera data is an active area of interest. Neural networks have the potential to overcome geometric line of sight effects and even greater detection accuracy than that achieved with the watershed algorithm technique described in §4.3. The pursuit of further gains in identification accuracy has lead to a collaboration with the University of Cagliari has exploring the application of a Faster Region-based Convolutional Neural Network (Faster R-CNN) to the detection of filaments in Elzar inversion data. The inverted mappings of the 5000 frame reference synthetic data set used in the benchmarking exercise in Chapter 5 was used as a training set for the Faster R-CNN. This network has demonstrated a precision of 93% (see Equation 5.6 for definition) for a sensitivity of 73% on this synthetic data [236]. Therefore, the neural network is able to identify over twice as many filaments for only a 6% increase in false positive detection rate compared to the watershed detection algorithm. This is due to the neural network's ability to identify faint filaments that fall below the detection threshold applied in the watershed technique. However, the neural network is yet to be applied to experimental camera data, so it remains to be seen how well the neural network is able to learn to interpret the full complexity of experimental camera data from synthetic training data.

7.2.6 Parallel filament structure

While the identification technique assumes filaments are homogeneous along field lines, once the filaments have been found their parallel structure in the images can also be investigated, opening up many new possibilities to explore 3D effects. This could provide valuable insight into the physics governing the draining of filaments towards the material

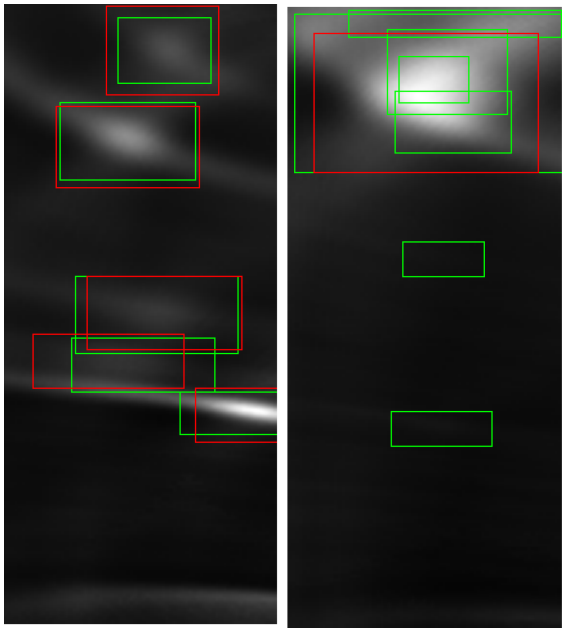

Figure 7.1: Two synthetic Elzar inversion maps with filaments identified using a Faster R-CNN. Synthetic input filaments are shown in green and output from the neural network in red. The neural network performs well on the left frame, while it performs poorly in the complex case on the right containing many faint and overlapping filaments. Reproduced from Ref. 236.

surfaces they are in contact with.

7.2.7 Application to alternative fields of view

Divertor views

In addition to mid-plane camera data, fast camera data exists that views the lower MAST divertor. The presence of the X-point in the divertor region introduces a broad range of additional phenomena such as shearing of mid-plane filaments and local generation of different types of divertor filaments, unconnected with those at the mid-plane.. Therefore, the application of the Elzar technique to the divertor region holds the potential to yield a wealth of valuable information, especially with the operation of the super-X divertor in MAST-U.

Alongside the investigation of filaments in the main-chamber of MAST which is the subject of this thesis, the inversion algorithm described in this chapter has also been successfully applied to the divertor camera view. Figure 7.2 shows a recent application of the Elzar inversion technique to divertor camera data from MAST shot 29720. The magnetic field in the divertor is nearly toroidal which makes the poloidal plane a good basis for these inversions since it is close to perpendicular to the magnetic field. In all other respects the method is identical to that used for the mid-plane view. The inner and outer divertor legs are treated separately.

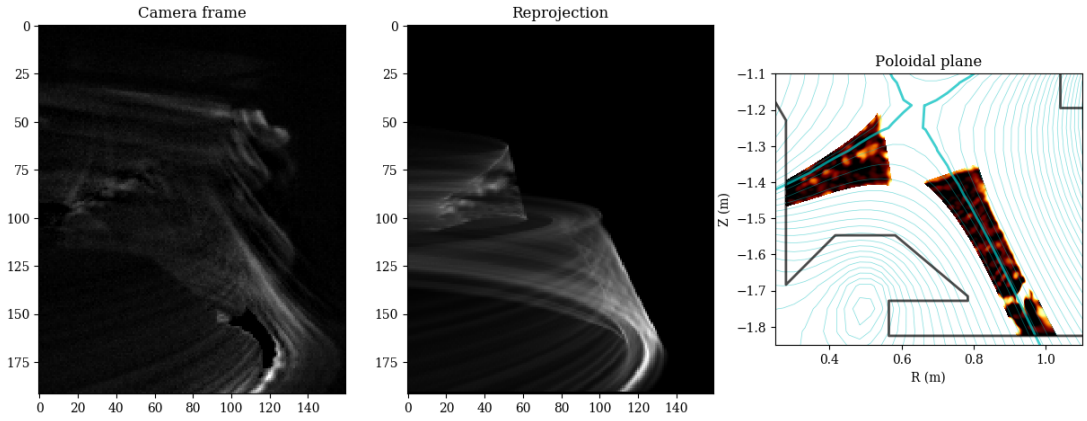

Figure 7.2: Application of the Elzar inversion technique to fast camera data viewing the MAST divertor. The left two panels show the background subtracted camera frame and the re-projection of the inverted data. The right hand panel shows the inverted emission along the inner and outer divertor legs in the poloidal plane. Filaments seen in the inversion are produced locally in the divertor close to the separatrix. Heavily sheared filaments generated upstream at the mid-plane can be seen at the top of the camera image, although this region is not included in the inversion domain due to the magnetic distortion. Figure courtesy of Dr N.Walkden.

Some artefacts appear in the inversions close to the divertor target, where the strike point is present in the camera images. The light emission from the strike-point does not align to the magnetic field, but rather aligns toroidally along the surface of the divertor. As such it cannot be properly represented by the field-line basis used for the inversion. Nevertheless the inversion process localizes filamentary structures in the poloidal plane on both the inner and outer divertor legs. This demonstrates that the inversion algorithm outlined here is not specific to a single camera view on MAST, but may be more widely applicable. Further work on these divertor inversions is being pursued by Dr N.Walkden at CCFE.

Stereoscopic imaging

A small set of stereoscopic camera data already exists from MAST and MAST-U's dual SA-X2 cameras will enable further collection of stereoscopic data. Application of the technique to stereoscopic fast camera data may help assess and better understand the effects of line of sight effects in the inversion results and provide alternative methods of quantifying the errors in the technique. Furthermore, stereoscopic imaging has the potential for inversion of the camera emission independently of a magnetic equilibrium description and can be used to better qualify the degree to which filaments are strongly field aligned. A collaboration with William and Mary University is underway to explore these possibilities.

7.2.8 Application to inter-ELM and ELM filaments

The technique has thus far been developed looking at L-mode filaments as a first case. This is because their physics is marginally better understood and they are brighter and more frequent than inter-ELM filaments [97], yet avoid the complications of higher energies and currents involved in ELMs. However, this should not prohibit applications to inter-ELM and ELM filaments. Application to inter-ELM filaments should be relatively straightforward, while in the case of ELMs, significant ballooning is expected to occur, distorting the filament's local magnetic field from that of the magnetic equilibrium and weakening the assertion that filaments are well aligned to the known toroidally symmetric magnetic equilibrium. This should not prevent the identification of ELM filaments but will likely introduce large uncertainties in their widths and velocities.

Bibliography

[1] World Bank Group. "Catalyzing Markets for Modern Energy Lighting". Technical report, The World Bank, Lagos, 2008. URL https://www.lightingafrica.org/wp-content/uploads/2017/12/Lighting-Africa-Brochure_October-2017.pdf.

[2] International Energy Agency. "2018 World Energy Outlook: Executive Summary". Technical report, International Energy Agency, 2018. URL https://webstore.iea.org/download/summary/190?fileName=English-WEO-2018-ES.pdf.

[3] World Energy Council. "World energy scenarios - Composing energy futures to 2050". Technical report, World Energy Council, 2013. URL https://www.worldenergy.org/assets/downloads/World-Energy-Scenarios_Composing-energy-futures-to-2050_Full-report1.pdf.

[4] United Nations Department of Economic and Social Affairs. "World Population Prospects 2017". Technical report, United Nations Department of Economic and Social Affairs, 2017. arXiv:1011.1669v3, URL http://dx.doi.org/10.1016/j.matchemphys.2004.09.004.

[5] International Energy Agency. "Global Energy and CO2 Status Report 2017". Technical report, International Energy Agency, 2018. URL https://www.iea.org/publications/freepublications/publication/GECO2017.pdf.

[6] L. Capuano. "International Energy Outlook 2018 (IEO2018)". Technical report, Center for Strategic and International Studies, Washington, DC, 2018. URL https://www.eia.gov/pressroom/presentations/capuano_07242018.pdf.

[7] G. Hartfield, J. Blunden and D. S. Arndt. "State of the Climate in 2017". Technical Report 8, American Meteorological Society, 2018. URL http://dx.doi.org/10.1175/2018BAMSStateoftheClimate.1.

[8] W. V. Sweet, R. E. Kopp, C. P. Weaver, J. Obeysekera, R. M. Horton, E. R. Thieler and C. Zervas. "Global and regional sea level rise scenarios for the United States". Technical Report January, National Oceanic and Atmospheric Ad-

ministration, 2017. URL https://tidesandcurrents.noaa.gov/publications/techrpt83_Global_and_Regional_SLR_Scenarios_for_the_US_final.pdf.

[9] World Health Organisation. "Ambient (outdoor) air quality and health", 2018. URL https://www.who.int/en/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health.

[10] World Health Organisation. "Household air pollution and health", 2018. URL https://www.who.int/en/news-room/fact-sheets/detail/household-air-pollution-and-health.

[11] E. Chivian and A. Bernstein. "How Our Health Depends on Biodiversity". Technical report, Center for Health and the Global Environment, Harvard Medical School, Boston, 2010. URL https://www.bu.edu/sph/files/2012/12/Chivian_and_Bernstein_2010_How_our_Health_Depends_on_Biodiversity.pdf.

[12] M. R. Allen, O. P. Dube and W. Solecki. "Framing and Context". *In* "Special Report: Global Warming of 1.5^{o}C", The Intergovernmental Panel on Climate Change, 2018. URL https://www.ipcc.ch/site/assets/uploads/sites/2/2019/05/SR15_Chapter1_Low_Res.pdf.

[13] B. C. V. Hirschhausen, C. Gerbaulet, C. Kemfert, F. Reitz and C. Ziehm. "German Nuclear Phase-Out Enters the Next Stage : Electricity Supply Remains Secure âĂŤ Major Challenges and High Costs for Dismantling and Final Waste Disposal". Technical report, German Institute for Economic Research, 2015. URL https://www.diw.de/documents/publikationen/73/diw_01.c.506840.de/diw_econ_bull_2015-22-1.pdf.

[14] K. Otsuki. "r-process Nucleosynthesis in Supernovae". *In* "Proceedings of 10th Symposium on Nuclei in the Cosmos", 47–53. Sissa Medialab, Trieste, Italy, 2009. ISBN 0031922804100. URL http://dx.doi.org/10.22323/1.053.0141.

[15] "Average binding energy per nucleon of the common isotopes". URL https://commons.wikimedia.org/w/index.php?curid=1540082.

[16] J. Wesson. *Tokamaks*. Oxford Univerity Press, Oxford, 3rd edition, 2004.

[17] J. Ongena and G. Van Oost. "Energy for future centuries Will fusion be an inexhaustible, safe and clean energy source?" *In* "Fusion Science and Technology", volume 41, 3–14. 2002. ISSN 15361055. URL http://dx.doi.org/10.13182/FST04-A464.

[18] M. Kovari, M. Coleman, I. Cristescu and R. Smith. "Tritium resources available for fusion reactors". *Nuclear Fusion*, **58**(2), 2018. ISSN 17414326. URL http://dx.doi.org/10.1088/1741-4326/aa9d25.

[19] T. B. Cochran, W. M. Arkin, R. S. Norris and M. M. Hoenig. *Nuclear Weapons Databook: Appendix C*, volume 2. Ballinger pulishing company, Cambridge, Massachesetts, 1987.

[20] G. Van Oost and J. Ongena. "Energy For Future Centuries - Will fusion be an inexhaustible, safe and clean energy source?" *Fusion Technology*, **29**(CONF-9509335), 1996. URL http://ippwww.ipp.kfa-juelich.de/papers/2009/docs/IN1_Ongena_paper.pdf.

[21] M. Romanelli, K. G. Mcclements, J. Cross, P. J. Knight, A. Thyagaraja and J. Callagham. "Full orbit simulations of ion collisional and turbulent transport in the MAST spherical tokamak". *Science*, 30, 2010.

[22] F. F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer, 1984. ISBN 9780306413322. URL https://books.google.co.uk/books?id=ToAtqnznr80C.

[23] A. Longo. "Magnetic coils and magnetic flux surfaces in the Wendelstein 7-X stellarator", 2018. URL https://www.colorado.edu/amath/sites/default/files/styles/medium/public/article-image/stellaratorwords.png?itok=DnAet9oN.

[24] A. H. Boozer. "What is a stellarator?" *Physics of Plasmas*, **5**(5), 1647–1655, 1998. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.872833.

[25] H. S. Bosch, R. C. Wolf, T. Andreeva, J. Baldzuhn, D. Birus, T. Bluhm, T. Bräuer, H. Braune, V. Bykov, A. Cardella *et al.* "Technical challenges in the construction of the steady-state stellarator Wendelstein 7-X". *Nuclear Fusion*, **53**(12), 2013. ISSN 00295515. URL http://dx.doi.org/10.1088/0029-5515/53/12/126001.

[26] L. A. El-Guebaly, P. Wilson, D. Henderson, M. Sawan, G. Sviatoslavsky, T. Tautges, R. Slaybaugh, B. Kiedrowski and A. Ibrahim. "Nuclear challenges and progress in designing stellarator fusion power plants". *Energy Conversion and Management*, **49**(7), 1859–1867, 2008. ISSN 01968904. URL http://dx.doi.org/10.1016/j.enconman.2007.09.032.

[27] EUROfusion. "The tokamak". URL https://www.euro-fusion.org/fileadmin/user_upload/Archive/wp-content/uploads/2011/09/jg05-537-1c.jpg.

[28] T. Noda, F. Abe, H. Araki and M. Okada. "Materials selection for reduced activation of fusion reactors". *Journal of Nuclear Materials*, **155-157**(PART 2), 581–584, 1988. ISSN 00223115. URL http://dx.doi.org/10.1016/0022-3115(88)90375-3.

[29] K. Ikeda. "ITER on the road to fusion energy". *Nuclear Fusion*, **50**(1), 014002, 2010. ISSN 00295515. URL http://dx.doi.org/10.1088/0029-5515/50/1/014002.

[30] M. Keilhacker. "Fusion physics progress on the Joint European Torus (JET)". *Plasma Physics and Controlled Fusion*, **41**(B), 1–23, 1999. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/41/12B/301.

[31] S. Ciattaglia, G. Federici, L. Barucca, A. Lampasi, S. Minucci and I. Moscato. "The European DEMO fusion reactor: Design status and challenges from balance of plant point of view". *In* "1st IEEE International Conference on Industrial and Commercial Power Systems", June, 1–7. 2017. ISBN 9781538639160. URL http://dx.doi.org/10.1109/EEEIC.2017.7977853.

[32] ITER Council. "ITER Research Plan within the Staged Approach". Technical Report ITR-18-003, ITER Council, 2018.

[33] R. Dux. "Impurity Transport in Tokamak Plasmas". Technical report, Max-Planck-Institut fur Plasmaphysik, 2004.

[34] M. Harrison. "Boundary Plasma". *In* "Plasmas", 395–439. Elsevier, 1984. URL http://dx.doi.org/10.1016/B978-0-12-478802-2.50020-8.

[35] M. R. Gilbert, S. L. Dudarev, S. Zheng, L. W. Packer and J. C. Sublet. "An integrated model for materials in a fusion power plant: Transmutation, gas production, and helium embrittlement under neutron irradiation". *Nuclear Fusion*, **52**(8), 2012. ISSN 00295515. URL http://dx.doi.org/10.1088/0029-5515/52/8/083019.

[36] F. Wagner, G. Becker, K. Behringer, D. Campbell, A. Eberhagen, W. Engelhardt, G. Fussmann, O. Gehre, J. Gernhardt, G. v. Gierke *et al.* "Regime of Improved Confinement and High Beta in Neutral-Beam-Heated Divertor Discharges of the ASDEX Tokamak". *Physical Review Letters*, **49**(19), 1408–1412, 1982. ISSN 00319007. URL http://dx.doi.org/10.1103/PhysRevLett.49.1408.

[37] M. Keilhacker. "H-mode confinement in tokamaks". *Plasma Physics and Controlled Fusion*, **29**(10A), 1401 –1413, 1987.

[38] P. Stangeby. *The Plasma Boundary of Magnetic Fusion Devices*. IOP Publishing, London, 2000. ISBN 0 7503 0559 2.

[39] P. C. Stangeby. "Modeling plasma contact with the main vessel walls of a divertor tokamak". *Physics of Plasmas*, **9**(8), 3489–3507, 2002. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.1493207.

[40] W. Fundamenski. *Power Exhaust in Fusion Plasmas*. Cambridge University Press, Cambridge, 1st edition, 2010. ISBN 9780521851718. URL https://books.google.co.uk/books?id=0Rgmo6625fkC.

[41] B. LaBombard, M. Umansky, R. Boivin, J. Goetz, J. Hughes, B. Lipschultz, D. Mossessian, C. Pitcher, J. Terry and A. Group. "Cross-field plasma transport and main-chamber recycling in diverted plasmas on Alcator C-Mod". *Nuclear Fusion*, **40**(12), 2041–2060, 2000. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/40/12/308.

[42] S. I. Krasheninnikov, D. a. D'Ippolito and J. R. Myra. "Recent theoretical progress in understanding coherent structures in edge and SOL turbulence". *Journal of Plasma Physics*, **74**(05), 679–717, 2008. ISSN 0022-3778. URL http://dx.doi.org/10.1017/S0022377807006940.

[43] B. Lipschultz, B. LaBombard, C. S. Pitcher and R. Boivin. "Investigation of the origin of neutrals in the main chamber of Alcator C-Mod". *Plasma Physics and Controlled Fusion*, **44**, 733–748, 2002. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/44/6/309.

[44] O. Garcia, R. Pitts, J. Horacek, A. Nielsen, W. Fundamenski, J. Graves, V. Naulin and J. J. Rasmussen. "Turbulent transport in the TCV SOL". *Journal of Nuclear Materials*, **363**, 575–580, 2016. ISSN 00223115. 1604.08790, URL http://dx.doi.org/10.1088/0029-5515/56/10/104004.

[45] F. Militello and J. T. Omotani. "Scrape off layer profiles interpreted with filament dynamics". *Nuclear Fusion*, **56**(10), 104004, 2016. 1604.08790v1, URL http://dx.doi.org/10.1088/0029-5515/56/10/104004.

[46] B. LaBombard, R. L. Boivin, M. Greenwald, J. Hughes, B. Lipschultz, D. Mossessian, C. S. Pitcher, J. L. Terry and S. J. Zweben. "Particle transport in the scrape-off layer and its relationship to discharge density limit in Alcator C-Mod". *Physics of Plasmas*, **8**(5 II), 2107–2117, 2001. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1352596.

[47] B. D. Dudson. *Edge turbulence in the Mega-Amp Spherical Tokamak*. Ph.D. thesis, University of Oxford, 2007.

[48] J. A. Boedo, D. L. Rudakov, R. A. Moyer, G. R. McKee, R. J. Colchin, M. J. Schaffer, P. G. Stangeby, W. P. West, S. L. Allen, T. E. Evans *et al.* "Transport by intermittency in the boundary of the DIII-D tokamak". *Physics of Plasmas*, **10**(5 II), 1670–1677, 2003. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1563259.

[49] D. a. D'Ippolito, J. R. Myra and S. J. Zweben. "Convective transport by intermittent blob-filaments: Comparison of theory and experiment". *Physics of Plasmas*, **18**(6), 060501, 2011. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.3594609.

[50] G. Y. Antar, S. I. Krasheninnikov, P. Devynck, R. P. Doerner, E. M. Hollmann, J. a. Boedo, S. C. Luckhardt and R. W. Conn. "Experimental evidence of intermittent convection in the edge of magnetic confinement devices." *Physical Review Letters*, **87**(6), 065001, 2001. ISSN 0031-9007. URL http://dx.doi.org/10.1103/PhysRevLett.87.065001.

[51] A. Kirk, A. J. Thornton, J. R. Harrison, F. Militello, N. R. Walkden, the MAST Team and the EUROfusion MST1 Team. "L-mode filament characteristics on MAST as a function of plasma current measured using visible imaging". *Plasma Physics and Controlled Fusion*, **58**(8), 085008, 2016. ISSN 0741-3335. 1602.03021, URL http://dx.doi.org/10.1088/0741-3335/58/8/085008.

[52] A. J. Thornton, G. Fishpool, A. Kirk, the MAST Team, the EUROfusion MST1 Team, t. E. M. M. Team and t. E. M. M. Team. "The effect of L mode filaments on divertor heat flux profiles as measured by infrared thermography on MAST". *Plasma Phys. Control. Fusion*, **57**(11), 115010, 2015. ISSN 13616587. URL http://dx.doi.org/10.1088/0741-3335/57/11/115010.

[53] W. Fundamenski. "Power and particle exhaust in tokamaks: Integration of plasma scenarios with plasma facing materials and components". *Journal of Nuclear Materials*, **390-391**(1), 10–19, 2009. ISSN 00223115. URL http://dx.doi.org/10.1016/j.jnucmat.2009.01.038.

[54] A. Loarte, B. Lipschultz, A. Kukushkin, G. Matthews, P. Stangeby, N. Asakura, G. Counsell, G. Federici, A. Kallenbach, K. Krieger *et al.* "Chapter 4: Power and particle control". *Nuclear Fusion*, **39**(12), S203–S263, 1999. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/47/6/S04.

[55] B. LaBombard, A. Kuang, D. Brunner, I. Faust, R. Mumgaard, M. Reinke, J. Terry, N. Howard, J. Hughes, M. Chilenski *et al.* "Impurity screening behavior of the high-field side scrape-off layer in near-double-null configurations: prospect for mitigating plasma-material interactions on RF actuators and first-wall components". *Nuclear Fusion*, **57**(7), 076021, 2017. ISSN 0029-5515. URL http://dx.doi.org/10.1088/1741-4326/aa6dd2.

[56] J. R. Myra and D. A. D'Ippolito. "Slow wave propagation and sheath interaction for ICRF waves in the tokamak SOL". *AIP Conference Proceedings*, **1187**, 153–156, 2009. ISSN 0094243X. URL http://dx.doi.org/10.1063/1.3273717.

[57] S. Krasheninnikov, A. Pigarov and D. Sigmar. "Plasma recombination and divertor detachment". *Physics Letters A*, **214**(5-6), 285–291, 1996. ISSN 03759601. URL http://dx.doi.org/10.1016/0375-9601(96)00185-5.

[58] B. Lipschultz, X. Bonnin, G. Counsell, A. Kallenbach, A. Kukushkin, K. Krieger, A. Leonard, A. Loarte, R. Neu, R. Pitts *et al.* "PlasmaâĂŞsurface interaction, scrape-off layer and divertor physics: implications for ITER". *Nuclear Fusion*, **47**(9), 1189–1205, 2007. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/47/9/016.

[59] J. Roth, E. Tsitrone, A. Loarte, T. Loarer, G. Counsell, R. Neu, V. Philipps, S. Brezinsek, M. Lehnen, P. Coad *et al.* "Recent analysis of key plasma wall interactions issues for ITER". *Journal of Nuclear Materials*, **390-391**(1), 1–9, 2009. ISSN 00223115. URL http://dx.doi.org/10.1016/j.jnucmat.2009.01.037.

[60] B. Kurzan, L. D. Horton, H. Murmann, J. Neuhauser, W. Suttrop and t. A. U. Team. "Thomson scattering analysis of large scale fluctuations in the ASDEX Upgrade edge". *Plasma Physics and Controlled Fusion*, **49**(6), 825–844, 2007. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/49/6/010.

[61] O. Garcia, J. Horacek and R. Pitts. "Intermittent fluctuations in the TCV scrape-off layer". *Nuclear Fusion*, **55**(6), 062002, 2015. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/55/6/062002.

[62] A. Fasoli, B. Labit, M. McGrath, S. H. Müller, G. Plyushchev, M. Podestà and F. M. Poli. "Electrostatic turbulence and transport in a simple magnetized plasma". *Physics of Plasmas*, **13**(5), 2006. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.2178773.

[63] S. H. Müller, A. Diallo, A. Fasoli, I. Furno, B. Labit, M. PodestÌĂ and M. Podestà. "Plasma blobs in a basic toroidal experiment: Origin, dynamics, and induced transport". *Physics of Plasmas*, **14**(11), 110704, 2007. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.2813193.

[64] J. A. Boedo, D. Rudakov, R. Moyer, S. Krasheninnikov, D. Whyte, G. McKee, G. Tynan, M. Schaffer, P. Stangeby, P. West *et al.* "Transport by intermittent convection in the boundary of the DIII-D tokamak". *Physics of Plasmas*, **8**(11), 4826–4833, 2001. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1406940.

[65] Y. H. Xu, S. Jachmich and R. R. Weynants. "On the properties of turbulence intermittency in the boundary of the TEXTOR tokamak". *Plasma Physics and Controlled Fusion*, **47**(10), 1841–1855, 2005. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/47/10/014.

[66] J. P. Graves, J. Horacek, R. a. Pitts and K. I. Hopcraft. "Self-similar density turbulence in the TCV tokamak scrape-off layer". *Plasma Physics and Controlled Fusion*, **47**(3), L1–L9, 2005. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/47/3/L01.

[67] F. Militello, P. Tamain, W. Fundamenski, A. Kirk, V. Naulin and A. H. Nielsen. "Experimental and numerical characterization of the turbulence in the scrape-off layer of MAST". *Plasma Physics and Controlled Fusion*, **55**(2), 025005, 2013. ISSN 0741-3335. arXiv:1305.5064v1, URL http://dx.doi.org/10.1088/0741-3335/55/2/025005.

[68] N. R. Walkden, F. Militello, J. Harrison, T. Farley, S. Silburn and J. Young. "Identification of intermittent transport in the scrape-off layer of MAST through high speed imaging". *Nuclear Materials and Energy*, **12**, 175–180, 2017. ISSN 23521791. URL http://dx.doi.org/10.1016/j.nme.2016.10.024.

[69] G. Y. Antar, G. Counsell, Y. Yu, B. Labombard and P. Devynck. "Universality of intermittent convective transport in the scrape-off layer of magnetically confined devices". *Physics of Plasmas*, **10**(2), 419–428, 2003. ISSN 08631042. URL http://dx.doi.org/10.1002/ctpp.200410031.

[70] S. J. Zweben, J. L. Terry, D. P. Stotler and R. J. Maqueda. "Invited Review Article: Gas puff imaging diagnostics of edge plasma turbulence in magnetic fusion devices". *Review of Scientific Instruments*, **88**(4), 2017. ISSN 10897623. URL http://dx.doi.org/10.1063/1.4981873.

[71] S. J. Zweben and S. S. Medley. "Visible imaging of edge fluctuations in the TFTR tokamak". *Phys. Fluids B*, **1**(10), 2058–2065, 1989. ISSN 08998221. URL http://dx.doi.org/10.1063/1.859070.

[72] A. Huber, U. Samm, B. Schweer and P. Mertens. "Results from a double Li-beam technique for measurement of both radial and poloidal components of electron density fluctuations using two thermal beams". *Plasma Phys. Control. Fusion*, **47**(3), 409, 2005. URL http://dx.doi.org/10.1088/0741-3335/47/3/002.

[73] E. Wolfrum, F. Aumayr, D. Wutte, H. P. Winter, E. Hintz, D. Rusbüldt and R. P. Schorn. "Fast lithium-beam spectroscopy of tokamak edge plasmas". *Review of Scientific Instruments*, **64**(8), 2285–2292, 1993. ISSN 00346748. URL http://dx.doi.org/10.1063/1.1144460.

[74] S. Zoletnik, G. H. Hu, B. Tál, D. Dunai, G. Anda, O. Asztalos, G. I. Pokol, S. Kálvin, J. Németh and T. Krizsanóczi. "Ultrafast two-dimensional lithium beam emission spectroscopy diagnostic on the EAST tokamak". *Review of Scientific*

Instruments, **89**(6), 2018. ISSN 10897623. URL http://dx.doi.org/10.1063/1.5017224.

[75] D. R. Smith, R. J. Fonck, G. R. McKee, N. Schoenbeck, I. Uzun-Kaymak, G. Winz, F. Feder, R. Feder, G. Labik, B. C. Stratton *et al.* "Overview of the beam emission spectroscopy diagnostic system on the National Spherical Torus Experiment." *Review of Scientific Instruments*, **81**(10), 10D717, 2010. ISSN 1089-7623. URL http://dx.doi.org/10.1063/1.3478660.

[76] A. R. Field, D. Dunai, N. J. Conway, S. Zoletnik and J. Sárközi. "Beam emission spectroscopy for density turbulence measurements on the MAST spherical tokamak". *Review of Scientific Instruments*, **80**(7), 073503, 2009. ISSN 00346748. URL http://dx.doi.org/10.1063/1.3170034.

[77] C. R. Negus, C. Giroud, A. G. Meigs, K. D. Zastrow and D. L. Hillis. "Enhanced core charge exchange recombination spectroscopy system on Joint European Torus". *Review of Scientific Instruments*, **77**(10), 10F102, 2006. ISSN 00346748. URL http://dx.doi.org/10.1063/1.2222170.

[78] J. Mlynar. "Diagnostics Neutral Beam Injector at the TCV Tokamak". Technical Report October, IAEA, 2001. URL http://infoscience.epfl.ch/record/121250/files/lrp_710_01_hq.pdf?version=1.

[79] G. Anda, A. Bencze, M. Berta, D. Dunai, P. Hacek, J. Krbec, D. Réfy, T. Krizsanóczi, S. Bató, T. Ilkei *et al.* "Lithium beam diagnostic system on the COMPASS tokamak". *Fusion Engineering and Design*, **108**, 1–6, 2016. ISSN 09203796. URL http://dx.doi.org/10.1016/j.fusengdes.2016.04.022.

[80] S. Zweben, R. Maqueda, D. Stotler, A. Keesee, J. Boedo, C. Bush, S. Kaye, B. LeBlanc, J. Lowrance, V. Mastrocola *et al.* "High-speed imaging of edge turbulence in NSTX". *Nuclear Fusion*, **44**(1), 134–153, 2004. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/44/1/016.

[81] Y. Sechrest, D. Smith, D. P. Stotler, T. Munsat and S. J. Zweben. "Comparison of beam emission spectroscopy and gas puff imaging edge fluctuation measurements in National Spherical Torus Experiment". *Physics of Plasmas*, **22**(5), 2015. ISSN 10897674. URL http://dx.doi.org/10.1063/1.4921215.

[82] G. R. McKee, R. Ashley, R. D. Durst, R. J. Fonck, M. Jakubowski, K. Tritz, K. H. Burrell, C. Greenfield and J. Robinson. "The beam emission spectroscopy diagnostic on the DIII-D tokamak". *Review of Scientific Instruments*, **70**(1), 913–916, 1999. ISSN 0034-6748. URL http://dx.doi.org/10.1063/1.1149416.

[83] G. R. McKee, R. J. Fonck, M. Jakubowski, K. H. Burrell, K. Hallatschek, R. A. Moyer, D. L. Rudakov, W. Nevins, G. D. Porter, P. Schoch *et al.* "Experimental characterization of coherent, radially-sheared zonal flows in the DIII-D tokamak". *Physics of Plasmas*, **10**(5), 1712–1719, 2003. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.1559974.

[84] S. J. Zweben, J. R. Myra, W. M. Davis, D. a. D'Ippolito, T. K. Gray, S. M. Kaye, B. P. LeBlanc, R. J. Maqueda, D. a. Russell and D. P. Stotler. "Blob structure and motion in the edge and SOL of NSTX". *Plasma Physics and Controlled Fusion*, **58**(4), 044007, 2016. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/58/4/044007.

[85] R. J. Maqueda, G. A. Wurden, D. P. Stotler, S. J. Zweben, B. LaBombard, J. L. Terry, J. L. Lowrance, V. J. Mastrocola, G. F. Renda, D. A. D'lppolito *et al.* "Gas puff imaging of edge turbulence (invited)". *Review of Scientific Instruments*, **74**(3 II), 2020–2026, 2003. ISSN 00346748. URL http://dx.doi.org/10.1063/1.1535249.

[86] J. R. Myra, D. A. D'Ippolito, D. P. Stotler, S. J. Zweben, B. P. LeBlanc, J. E. Menard, R. J. Maqueda, J. Boedo, D. A. D'Ippolito, D. P. Stotler *et al.* "Blob birth and transport in the tokamak edge plasma: Analysis of imaging data". *Physics of Plasmas*, **13**(9), 092509, 2006. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.2355668.

[87] S. Zweben, W. Davis, S. Kaye, J. Myra, R. Bell, B. LeBlanc, R. Maqueda, T. Munsat, S. Sabbagh, Y. Sechrest *et al.* "Edge and SOL turbulence and blob variations over a large database in NSTX". *Nuclear Fusion*, **55**(9), 093035, 2015. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/55/9/093035.

[88] S. J. Zweben, D. P. Stotler, J. L. Terry, B. Labombard, M. Greenwald, M. Muterspaugh, C. S. Pitcher, K. Hallatschek, R. J. Maqueda, B. Rogers *et al.* "Edge turbulence imaging in the Alcator C-Mod tokamak". *Physics of Plasmas*, **9**(5), 1981–1989, 2002. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1445179.

[89] R. Kube, O. E. Garcia, B. LaBombard, J. L. Terry and S. J. Zweben. "Blob sizes and velocities in the Alcator C-Mod scrape-off layer". *Journal of Nuclear Materials*, **438**(SUPPL), S505–S508, 2013. ISSN 00223115. URL http://dx.doi.org/10.1016/j.jnucmat.2013.01.104.

[90] O. E. Garcia, S. M. Fritzner, R. Kube, I. Cziegler, B. Labombard and J. L. Terry. "Intermittent fluctuations in the Alcator C-Mod scrape-off layer". *Physics*

of Plasmas, **20**(5), 055901, 2013. ISSN 1070664X. 1801.00612, URL http://dx.doi.org/10.1063/1.4802942.

[91] G. Fuchert, G. Birkenmeier, D. Carralero, T. Lunt, P. Manz, H. W. Müller, B. Nold, M. Ramisch, V. Rohde and U. Stroth. "Blob properties in L- and H-mode from gas-puff imaging in ASDEX upgrade". *Plasma Physics and Controlled Fusion*, **56**(12), 125001, 2014. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/56/12/125001.

[92] O. Grulke, J. L. Terry, I. Cziegler, B. LaBombard and O. E. Garcia. "Experimental investigation of the parallel structure of fluctuations in the scrape-off layer of Alcator C-Mod". *Nuclear Fusion*, **54**(4), 43012, 2014. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/54/4/043012.

[93] A. Kirk, H. R. Wilson, R. Akers, N. J. Conway, G. F. Counsell, S. C. Cowley, J. Dowling, B. Dudson, A. Field, F. Lott *et al.* "Structure of ELMs in MAST and the implications for energy deposition". *Plasma Physics and Controlled Fusion*, **47**(2), 315–333, 2005. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/47/2/008.

[94] A. Kirk, N. Ben Ayed, G. Counsell, B. Dudson, T. Eich, A. Herrmann, B. Koch, R. Martin, A. Meakins, S. Saarelma *et al.* "Filament structures at the plasma edge on MAST". *Plasma Physics and Controlled Fusion*, **48**(12B), B433–B441, 2006. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/48/12B/S41.

[95] A. Kirk, B. Koch, R. Scannell, H. R. Wilson, G. Counsell, J. Dowling, A. Herrmann, R. Martin, M. Walsh and M. Team. "Evolution of filament structures during edge-localized modes in the MAST tokamak". *Physical Review Letters*, **96**(18), 185001, 2006. ISSN 0031-9007. URL http://dx.doi.org/10.1103/PhysRevLett.96.185001.

[96] B. D. Dudson, N. B. Ayed, A. Kirk, H. R. Wilson, G. Counsell, X. Xu, M. Umansky, P. B. Snyder, B. Lloyd, N. Ben Ayed *et al.* "Experiments and simulation of edge turbulence and filaments in MAST". *Plasma Physics and Controlled Fusion*, **50**(12), 124012, 2008. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/50/12/124012.

[97] N. Ben Ayed, A. Kirk, B. Dudson, S. Tallents, R. G. L. Vann and H. R. Wilson. "Inter-ELM filaments and turbulent transport in the Mega-Amp Spherical Tokamak". *Plasma Physics and Controlled Fusion*, **51**(3), 035016, 2009. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/51/3/035016.

[98] J. R. Harrison, A. Kirk, I. T. Chapman, P. Cahyna, Y. Liu, E. Nardon and A. J. Thornton. "Characteristics of X-point lobe structures in single-null discharges

on MAST". *Nuclear Fusion*, **54**(6), 064015, 2014. ISSN 17414326. URL http://dx.doi.org/10.1088/0029-5515/54/6/064015.

[99] F. Militello, N. R. Walkden, T. Farley, W. A. Gracias, J. Olsen, F. Riva, L. Easy, N. Fedorczak, I. Lupelli, J. Madsen *et al.* "Multi-code analysis of scrape-off layer filament dynamics in MAST". *Plasma Physics and Controlled Fusion*, **58**(10), 105002, 2016. ISSN 13616587. URL http://dx.doi.org/10.1088/0741-3335/58/10/105002.

[100] N. R. Walkden, J. Harrison, F. Militello, P. Innocente, B. Labit, H. Reimerdes and T. Farley. "Fluctuation characteristics of the TCV snowflake divertor measured with high speed visible imaging". *Plasma Physics and Controlled Fusion*, **60**, 115008, 2018. ISSN 0741-3335. URL http://dx.doi.org/10.1088/1361-6587/aae005.

[101] J. Cavalier, N. Lemoine, F. Brochard, V. Weinzettl, J. Seidl, S. Silburn, P. Tamain, R. Dejarnac, J. Adamek and R. Panek. "Tomographic reconstruction of tokamak edge turbulence from single visible camera data and automatic turbulence structure tracking". *Nuclear Fusion*, **59**(5), 056025, 2019. ISSN 0029-5515. URL http://dx.doi.org/10.1088/1741-4326/ab0d4c.

[102] G. S. Yun, W. Lee, M. J. Choi, J. Lee, H. K. Park, B. Tobias, C. W. Domier, N. C. Luhmann, A. J. H. Donne and J. H. Lee. "Two-dimensional visualization of growth and burst of the edge-localized filaments in KSTAR H-mode plasmas". *Physical Review Letters*, **107**(4), 1–4, 2011. ISSN 00319007. URL http://dx.doi.org/10.1103/PhysRevLett.107.045004.

[103] A. Kirk, T. Eich, A. Herrmann, H. W. Muller, L. D. Horton, G. F. Counsell, M. Price, V. Rohde, V. Bobkov, B. Kurzan *et al.* "The spatial structure of type-I ELMs at the mid-plane in ASDEX Upgrade and a comparison with data from MAST". *Plasma Physics and Controlled Fusion*, **47**(7), 995–1013, 2005. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/47/7/003.

[104] D. Iraji, A. Diallo, A. Fasoli, I. Furno and S. Shibaev. "Fast visible imaging of turbulent plasma in TORPEX". *Review of Scientific Instruments*, **79**(10), 10F508, 2008. ISSN 00346748. URL http://dx.doi.org/10.1063/1.2953677.

[105] D. Goodall. "High speed cine film studies of plasma behaviour and plasma surface interactions in tokamaks". *Journal of Nuclear Materials*, **111-112**, 11–22, 1982. ISSN 00223115. URL http://dx.doi.org/10.1016/0022-3115(82)90174-X.

[106] J. L. Terry, S. Ballinger, D. Brunner, B. LaBombard, A. E. White and S. J. Zweben. "Fast imaging of filaments in the X-point region of Alcator C-Mod".

Nuclear Materials and Energy, **12**, 989–993, 2017. ISSN 23521791. URL http://dx.doi.org/10.1016/j.nme.2016.11.020.

[107] R. Nguyen Van Yen, N. Fedorczak, F. Brochard, G. Bonhomme, K. Schneider, M. Farge and P. Monier-Garbet. "Tomographic reconstruction of tokamak plasma light emission from single image using wavelet-vaguelette decomposition". *Nuclear Fusion*, **52**(1), 013005, 2012. ISSN 00295515. URL http://dx.doi.org/10.1088/0029-5515/52/1/013005.

[108] R. Baude and M. Desecures. "TRACK software", 2018. URL https://www.aprex-solutions.com/.

[109] A. Theodorsen, O. E. Garcia, J. Horacek, R. Kube and R. A. Pitts. "Scrape-off layer turbulence in TCV: Evidence in support of stochastic modelling". *Plasma Physics and Controlled Fusion*, **58**(4), 044006, 2016. ISSN 13616587. 1604.08163, URL http://dx.doi.org/10.1088/0741-3335/58/4/044006.

[110] S. J. Zweben. "Search for coherent structure within tokamak plasma turbulence". *Physics of Fluids*, **28**(3), 974, 1985. ISSN 00319171. URL http://dx.doi.org/10.1063/1.865069.

[111] M. Endler, H. Niedermeyer, L. Giannone, E. Kolzhauer, A. Rudyj, G. Theimer and N. Tsois. "Measurements and modelling of electrostatic fluctuations in the scrape-off layer of ASDEX". *Nuclear Fusion*, **35**(11), 1307–1339, 1995. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/35/11/I01.

[112] M. Endler. "Turbulent SOL transport in stellarators and tokamaks". *Journal of Nuclear Materials*, **266**, 84–90, 1999. ISSN 00223115. URL http://dx.doi.org/10.1016/S0022-3115(98)00659-X.

[113] M. V. Umansky, S. I. Krasheninnikov, B. LaBombard and J. L. Terry. "Comments on particle and energy balance in the edge plasma of Alcator C-Mod". *Physics of Plasmas*, **5**(9), 3373–3376, 1998. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.873051.

[114] H. S. Bosch, J. Neuhauser, R. Schneider, A. Field, A. Herrmann, G. Lieder, W. Junker, C. S. Pitcher, D. Reiter, B. Braams *et al.* "2D modelling of the ASDEX-Upgrade scrape-off layer and divertor plasma". *Journal of Nuclear Materials*, **220-222**, 558–562, 1995. ISSN 00223115. URL http://dx.doi.org/10.1016/0022-3115(94)00539-7.

[115] S. I. Krasheninnikov. "On scrape off layer plasma transport". *Physics Letters, Section A: General, Atomic and Solid State Physics*, **283**(5-6), 368–370, 2001. ISSN 03759601. URL http://dx.doi.org/10.1016/S0375-9601(01)00252-3.

[116] H. Q. Liu, K. Hanada, N. Nishino, R. Ogata, M. Ishiguro, H. Zushi, K. Nakamura, M. Sakamoto, H. Idei, A. Fujisawa *et al.* "Study of blob-like structures in QUEST". *Journal of Nuclear Materials*, **415**(1 SUPPL), S620–S623, 2011. ISSN 00223115. URL http://dx.doi.org/10.1016/j.jnucmat.2010.12.225.

[117] D. L. Rudakov, J. A. Boedo, R. A. Moyer, S. Krasheninnikov, P. C. Stangeby, J. G. Watkins, W. P. West, D. G. Whyte and G. Antar. "Fluctuation-driven transport in the DIII-D boundary". *Plasma Phys. Control. Fusion*, **44**, 717–731, 2002.

[118] G. S. Xu, V. Naulin, W. Fundamenski, C. Hidalgo, J. A. Alonso, C. Silva, B. Gonçalves, A. H. Nielsen, J. Juul Rasmussen, S. I. Krasheninnikov *et al.* "Blob/hole formation and zonal-flow generation in the edge plasma of the JET tokamak". *Nuclear Fusion*, **49**(9), 092002, 2009. ISSN 17414326. URL http://dx.doi.org/10.1088/0029-5515/49/9/092002.

[119] B. Nold, G. D. Conway, T. Happel, H. W. Müller, M. Ramisch, V. Rohde and U. Stroth. "Generation of blobs and holes in the edge of the ASDEX Upgrade tokamak". *Plasma Physics and Controlled Fusion*, **52**(6), 065005, 2010. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/52/6/065005.

[120] S. I. Krasheninnikov, A. Y. Pigarov, S. A. Galkin, G. Q. Yu, D. A. D'Ippolito, J. R. Myra, D. R. Mccarthy, W. M. Nevins, T. D. Rognlien, X. Q. Xu *et al.* "Blobs and Cross-Field Transport in the Tokamak Edge Plasmas". *19th Fusion Energy Conference 14 - 19 October 2002 Lyon, France*, IAEA–CN–94/TH/4–1, 2003.

[121] N. R. Walkden. *Properties of Intermittent Transport in the Mega Ampere Spherical Tokamak*. Ph.D. thesis, University of York, 2014.

[122] L. Easy. *Three Dimensional Simulations of Scrape-Off Layer Filaments*. Ph.D. thesis, University of York, 2016.

[123] D. A. D'Ippolito, J. R. Myra and S. I. Krasheninnikov. "Cross-field blob transport in tokamak scrape-off-layer plasmas". *Physics of Plasmas*, **9**(1), 222, 2002. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1426394.

[124] W. Fundamenski, O. E. Garcia, V. Naulin, R. A. Pitts, A. H. Nielsen, J. Juul Rasmussen, J. Horacek and J. P. Graves. "Dissipative processes in interchange driven scrape-off layer turbulence". *Nuclear Fusion*, **47**(5), 417–433, 2007. ISSN 17414326. URL http://dx.doi.org/10.1088/0029-5515/47/5/006.

[125] O. Grulke, J. L. Terry, B. LaBombard and S. J. Zweben. "Radially propagating fluctuation structures in the scrape-off layer of Alcator C-Mod". *Physics of Plas-*

mas, **13**(1), 012306, 2006. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.2164991.

[126] R. Moyer, J. Cuthbertson, T. Evans, G. Porter and J. Watkins. "The role of turbulent transport in DIII-D edge and divertor plasmas". *Journal of Nuclear Materials*, **241-243**, 633–638, 1997. ISSN 0022-3115. URL http://dx.doi.org/10.1016/S0022-3115(97)80113-4.

[127] C. Silva, B. Gonçalves, C. Hidalgo, M. A. Pedrosa, K. Erents, G. Matthews and R. A. Pitts. "Fluctuation measurements using a five-pin triple probe in the Joint European Torus boundary plasma". *Review of Scientific Instruments*, **75**(10 II), 4314–4316, 2004. ISSN 00346748. URL http://dx.doi.org/10.1063/1.1787578.

[128] D. Carralero, P. Manz, L. Aho-Mantila, G. Birkenmeier, M. Brix, M. Groth, H. W. Müller, U. Stroth, N. Vianello and E. Wolfrum. "Experimental Validation of a Filament Transport Model in Turbulent Magnetized Plasmas". *Physical Review Letters*, **115**(21), 1–5, 2015. ISSN 10797114. arXiv:1505.07741v1, URL http://dx.doi.org/10.1103/PhysRevLett.115.215002.

[129] S. H. Müller, C. Theiler, A. Fasoli, I. Furno, B. Labit, G. R. Tynan, M. Xu, Z. Yan and J. H. Yu. "Studies of blob formation, propagation and transport mechanisms in basic experimental plasmas (TORPEX and CSDX)". *Plasma Physics and Controlled Fusion*, **51**(5), 2009. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/51/5/055020.

[130] C. Theiler, I. Furno, A. Fasoli, P. Ricci, B. Labit and D. Iraji. "Blob motion and control in simple magnetized plasmas". *Physics of Plasmas*, **18**(5), 2011. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.3562944.

[131] O. E. Garcia. "Stochastic modeling of intermittent scrape-off layer plasma fluctuations". *Physical Review Letters*, **108**(265001), 1–5, 2012. ISSN 00319007. URL http://dx.doi.org/10.1103/PhysRevLett.108.265001.

[132] D. A. Russell, J. R. Myra and D. A. D'Ippolito. "Saturation mechanisms for edge turbulence". *Physics of Plasmas*, **16**(12), 122304, 2009. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.3270051.

[133] J. R. Myra, D. A. Russell, D. A. D'Ippolito, J.-W. Ahn, R. Maingi, R. J. Maqueda, D. P. Lundberg, D. P. Stotler, S. J. Zweben, J. Boedo *et al.* "Reduced model simulations of the scrape-off-layer heat-flux width and comparison with experiment". *Physics of Plasmas*, **18**(1), 012305, 2011. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.3526676.

[134] J. L. Terry, N. P. Basse, I. Cziegler, M. Greenwald, O. Grulke, B. Labombard, S. J. Zweben, E. M. Edlund, J. W. Hughes, L. Lin *et al.* "Transport phenomena in the edge of Alcator C-Mod plasmas". *Nuclear Fusion*, **45**(11), 1321–1327, 2005. ISSN 17414326. URL http://dx.doi.org/10.1088/0029-5515/45/11/013.

[135] I. Furno, B. Labit, M. Podestà, A. Fasoli, S. H. Müller, F. M. Poli, P. Ricci, C. Theiler, S. Brunner, A. Diallo *et al.* "Experimental observation of the blob-generation mechanism from interchange waves in a plasma". *Physical Review Letters*, **100**(5), 8–11, 2008. ISSN 00319007. URL http://dx.doi.org/10.1103/PhysRevLett.100.055004.

[136] J. Cheng, L. W. Yan, W. Y. Hong, K. J. Zhao, T. Lan, J. Qian, A. D. Liu, H. L. Zhao, Y. Liu, Q. W. Yang *et al.* "Statistical characterization of blob turbulence across the separatrix in HL-2A tokamak". *Plasma Physics and Controlled Fusion*, **52**(5), 055003, 2010. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/52/5/055003.

[137] D. A. D'Ippolito and J. R. Myra. "Thermal transport catastrophe and the tokamak edge density limit". *Physics of Plasmas*, **13**(6), 062503, 2006. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.2206168.

[138] J. R. Myra, D. A. Russell and D. A. D'Ippolito. "Transport of perpendicular edge momentum by drift-interchange turbulence and blobs". *Physics of Plasmas*, **15**(3), 032304, 2008. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.2889419.

[139] J. Myra, J. Boedo, B. Coppi, D'Ippolito, S. Krasheninnikov, B. P. LeBlanc, M. Lontano, R. Maqueda, D. A. Russell, D.P.Stotler *et al.* "Blob Transport Models, Experiments, and the Accretion Theory of Spontaneous Rotation". *In* "21st IAEA Fusion Energy Conference 2006", October. Chengdu, 2006. URL http://www-naweb.iaea.org/napc/physics/FEC/FEC2006/papers/th_p6-21.pdf.

[140] J. R. Myra. "Current carrying blob filaments and edge-localized-mode dynamics". *Physics of Plasmas*, **14**(10), 102314, 2007. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.2776900.

[141] G. S. Xu, V. Naulin, W. Fundamenski, J. J. Rasmussen, A. H. Nielsen and B. N. Wan. "Intermittent convective transport carried by propagating electromagnetic filamentary structures in nonuniformly magnetized plasma". *Physics of Plasmas*, **17**(2), 022501, 2010. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.3302535.

[142] S. I. Krasheninnikov, A. I. Smolyakov and T. K. Soboleva. "On anomalous cross-field edge plasma convection in fusion devices". *Physics of Plasmas*, **12**(7), 1–5, 2005. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1940061.

[143] S. I. Krasheninnikov and A. I. Smolyakov. "On neutral wind and blob motion in linear devices". *Physics of Plasmas*, **10**(7), 3020–3021, 2003. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.1579692.

[144] F. L. Hinton and C. W. Horton. "Amplitude limitation of a collisional drift wave instability". *Physics of Fluids*, **14**(1), 116–123, 1971. ISSN 10706631. URL http://dx.doi.org/10.1063/1.1693260.

[145] R. Fitzpatrick. *Plasma Physics: An Introduction*. CRC Press, 2014. ISBN 978-1466594265. URL http://farside.ph.utexas.edu/teaching/plasma/Plasma/Plasmahtml.html.

[146] O. E. Garcia, R. A. Pitts, J. Horacek, J. Madsen, V. Naulin, A. H. Nielsen and J. J. Rasmussen. "Collisionality dependent transport in TCV SOL plasmas". *Plasma Physics and Controlled Fusion*, **49**(12B), B47, 2007. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/49/12B/S03.

[147] G. Q. Yu and S. I. Krasheninnikov. "Dynamics of blobs in scrape-off-layer/shadow regions of tokamaks and linear devices". *Physics of Plasmas*, **10**(11), 4413–4418, 2003. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.1616937.

[148] S. Krasheninnikov, D. Ryutov and G. Yu. "Large plasma pressure perturbations and radial convective transport in a tokamak". *Journal of Plasma and Fusion Research*, **6**, 139, 2004.

[149] D. D. Ryutov and R. H. Cohen. "Instability driven by sheath boundary conditions and limited to divertor legs". *Contributions to Plasma Physics*, **44**(1-3), 168–175, 2004. ISSN 08631042. URL http://dx.doi.org/10.1002/ctpp.200410023.

[150] C. K. Tsui, J. A. Boedo, J. R. Myra, B. Duval, B. Labit, C. Theiler, N. Vianello, W. A. J. Vijvers, H. Reimerdes, S. Coda *et al.* "Filamentary velocity scaling validation in the TCV tokamak". *Physics of Plasmas*, **25**(7), 072506, 2018. ISSN 10897674. URL http://dx.doi.org/10.1063/1.5038019.

[151] D. A. Russell, J. R. Myra, D. A. D'Ippolito, D. A. Russell and D. A. D'Ippolito. "Collisionality and magnetic geometry effects on tokamak edge turbulent transport. II. Many-blob turbulence in the two-region model". *Physics of Plasmas*, **13**(11), 112502, 2006. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.2364858.

[152] D. A. D'Ippolito, J. R. Myra, D. A. Russell and G. Q. Yu. "Rotational stability of plasma blobs". *Physics of Plasmas*, **11**(10), 4603–4609, 2004. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1785791.

[153] D. A. D'Ippolito and J. R. Myra. "Blob stability and transport in the scrape-off-layer". *Physics of Plasmas*, **10**(10), 4029–4039, 2003. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1606447.

[154] R. Kube and O. E. Garcia. "Velocity scaling for filament motion in scrape-off layer plasmas". *Physics of Plasmas*, **18**(10), 102314, 2011. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.3647553.

[155] J. R. Angus and S. I. Krasheninnikov. "Inviscid evolution of large amplitude filaments in a uniform gravity field". *Physics of Plasmas*, **21**(11), 112504, 2014. ISSN 10897674. URL http://dx.doi.org/10.1063/1.4901237.

[156] C. Theiler, I. Furno, P. Ricci, A. Fasoli, B. Labit, S. H. Müller and G. Plyushchev. "Cross-field motion of plasma blobs in an open magnetic field line configuration". *Physical Review Letters*, **103**(6), 2–5, 2009. ISSN 00319007. URL http://dx.doi.org/10.1103/PhysRevLett.103.065001.

[157] J. T. Omotani, F. Militello, L. Easy and N. R. Walkden. "The effects of shape and amplitude on the velocity of scrape-off layer filaments". *Plasma Physics and Controlled Fusion*, **58**(1), 014030, 2015. ISSN 13616587. 1506.02814, URL http://dx.doi.org/10.1088/0741-3335/58/1/014030.

[158] F. Militello, L. Garzotti, J. Harrison, J. T. Omotani, R. Scannell, S. Allan, A. Kirk, I. Lupelli and A. J. Thornton. "Characterisation of the L-mode scrape off layer in MAST: Decay lengths". *Nuclear Fusion*, **56**(1), 16006, 2016. ISSN 17414326. URL http://dx.doi.org/10.1088/0029-5515/56/1/016006.

[159] D. G. Whyte, B. L. Lipschultz, P. C. Stangeby, J. Boedo, D. L. Rudakov, J. G. Watkins and W. P. West. "The magnitude of plasma flux to the main-wall in the DIII-D tokamak". *Plasma Physics and Controlled Fusion*, **47**(10), 1579–1607, 2005. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/47/10/002.

[160] N. Walkden, A. Wynn, F. Militello, B. Lipschultz, G. Matthews, C. Guillemaut, J. Harrison and D. Moulton. "Statistical analysis of the ion flux to the JET outer wall". *Nuclear Fusion*, **57**(3), 036016, 2017. ISSN 0029-5515. URL http://dx.doi.org/10.1088/1741-4326/aa515a.

[161] D. Carralero, J. Madsen, S. A. Artene, M. Bernert, G. Birkenmeier, T. Eich, G. Fuchert, F. Laggner, V. Naulin, P. Manz *et al.* "A study on the density shoulder formation in the SOL of H-mode plasmas". *Nuclear Materials and Energy*, **12**(August), 1189–1193, 2017. ISSN 23521791. URL http://dx.doi.org/10.1016/j.nme.2016.11.016.

[162] A. Q. Kuang, B. LaBombard, D. Brunner, O. E. Garcia, R. Kube and A. Theodorsen. "Plasma fluctuations in the scrape-off layer and at the divertor target in Alcator C-Mod and their relationship to divertor collisionality and density shoulder formation". *Nuclear Materials and Energy*, **19**(March), 295–299, 2019. ISSN 23521791. URL http://dx.doi.org/10.1016/j.nme.2019.02.038.

[163] A. Wynn, B. Lipschultz, G. Matthews, B. Tal, F. Militello, N. Walkden, C. Guillemaut, J. Harrison, A. Huber, U. Kruezi *et al.* "Scrape-Off layer density shoulder formation and evolution in JET". *In* "43rd European Physical Society Conference on Plasma Physics, EPS 2016", O4.114. 2016.

[164] N. Vianello, C. Tsui, C. Theiler, S. Allan, J. Boedo, B. Labit, H. Reimerdes, K. Verhaegh, W. A. Vijvers, N. Walkden *et al.* "Modification of SOL profiles and fluctuations with line-average density and divertor flux expansion in TCV". *Nuclear Fusion*, **57**(11), 116014, 2017. ISSN 17414326. URL http://dx.doi.org/10.1088/1741-4326/aa7db3.

[165] O. E. Garcia, I. Cziegler, R. Kube, B. LaBombard and J. L. Terry. "Burst statistics in Alcator C-Mod SOL turbulence". *Journal of Nuclear Materials*, **438**(SUPPL), S180–S183, 2013. ISSN 00223115. arXiv:1205.6996v1, URL http://dx.doi.org/10.1016/j.jnucmat.2013.01.054.

[166] F. Militello. "Predicting Scrape-Off Layer profiles and filamentary transport for reactor relevant devices". *In* "27th IAEA Fusion Energy Conference", Gandhinagar, 2018. URL https://conferences.iaea.org/indico/event/151/contributions/6215/.

[167] F. Militello and J. T. Omotani. "On the relation between non-exponential scrape off layer profiles and the dynamics of filaments". *Plasma Physics and Controlled Fusion*, **58**(12), 125004, 2016. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/58/12/125004.

[168] F. Militello, T. Farley, K. Mukhi, N. Walkden and J. T. Omotani. "A two-dimensional statistical framework connecting thermodynamic profiles with filaments in the scrape off layer and application to experiments". *Physics of Plasmas*, **25**(5), 056112, 2018. ISSN 1070-664X. URL http://dx.doi.org/10.1063/1.5017919.

[169] L. Easy, F. Militello, J. Omotani, N. R. Walkden and B. Dudson. "Investigation of the effect of resistivity on scrape off layer filaments using three-dimensional simulations". *Physics of Plasmas*, **23**(1), 012512, 2016. ISSN 10897674. 1508.04085, URL http://dx.doi.org/10.1063/1.4940330.

[170] N. R. Walkden, A. Wynn, F. Militello, B. Lipschultz, G. Matthews, C. Guillemaut, J. Harrison and D. Moulton. "Interpretation of scrape-off layer profile evolution and first-wall ion flux statistics on JET using a stochastic framework based on fillamentary motion". *Plasma Physics and Controlled Fusion*, **59**(8), 085009, 2017. ISSN 0741-3335. URL http://dx.doi.org/10.1088/1361-6587/aa7365.

[171] M. Kočan, J. P. Gunn, S. Carpentier-Chouchana, A. Herrmann, A. Kirk, M. Komm, H. W. Müller, J. Y. Pascal, R. A. Pitts, V. Rohde *et al.* "Measurements of ion energies in the tokamak plasma boundary". *Journal of Nuclear Materials*, **415**, 1133–1138, 2011. ISSN 00223115. URL http://dx.doi.org/10.1016/j.jnucmat.2010.08.036.

[172] S. Elmore, S. Y. Allan, A. Kirk, G. Fishpool, J. Harrison, P. Tamain, M. Kočan, R. Gaffka, R. Stephen and J. W. Bradley. "Upstream and divertor ion temperature measurements on MAST by retarding field energy analyser". *Plasma Physics and Controlled Fusion*, **54**(6), 065001, 2012. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/54/6/065001.

[173] M. Kočan, F. P. Gennrich, A. Kendl and H. W. Müller. "Ion temperature fluctuations in the ASDEX Upgrade scrape-off layer". *Plasma Physics and Controlled Fusion*, **54**(8), 085009, 2012. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/54/8/085009.

[174] S. Y. Allan, S. Elmore, G. Fishpool and B. Dudson. "Ion temperature measurements of L-mode filaments in MAST by retarding field energy analyser". *Plasma Physics and Controlled Fusion*, **58**(4), 045014, 2016. ISSN 13616587. URL http://dx.doi.org/10.1088/0741-3335/58/4/045014.

[175] F. Militello, W. Fundamenski, V. Naulin and A. H. Nielsen. "Simulations of edge and scrape off layer turbulence in mega ampere spherical tokamak plasmas". *Plasma Physics and Controlled Fusion*, **54**(9), 095011, 2012. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/54/9/095011.

[176] A. H. Nielsen, J. Madsen, G. S. Xu, V. Naulin, J. J. Rasmussen and N. Yan. "2D fluid simulations of interchange turbulence with ion dynamics". *40th European Physical Society Conference on Plasma Physics*, 3–6, 2013.

[177] N. Fedorczak, A. Gallo, P. Tamain, H. Bufferand, G. Ciraolo and P. Ghendrih. "On the dynamics of blobs in scrape-off layer plasma: Model validation from two-dimensional simulations and experiments in Tore Supra". *Contributions to Plasma Physics*, **58**(6-8), 471–477, 2018. ISSN 15213986. URL http://dx.doi.org/10.1002/ctpp.201700169.

[178] O. E. Garcia, V. Naulin, a. H. Nielsen and J. J. Rasmussen. "Turbulence and intermittent transport at the boundary of magnetized plasmas". *Physics of Plasmas*, **12**(2005), 1–14, 2005. ISSN 1070664X. URL http://dx.doi.org/10.1063/1.1925617.

[179] N. Katz, J. Egedal, W. Fox, A. Le and M. Porkolab. "Experiments on the propagation of plasma filaments". *Physical Review Letters*, **101**(1), 1–4, 2008. ISSN 00319007. URL http://dx.doi.org/10.1103/PhysRevLett.101.015003.

[180] P. Ricci, F. D. Halpern, S. Jolliet, J. Loizu, A. Mosetto, A. Fasoli, I. Furno and C. Theiler. "Simulation of plasma turbulence in scrape-off layer conditions: the GBS code, simulation results and code validation". *Plasma Physics and Controlled Fusion*, **54**(12), 124047, 2012. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/54/12/124047.

[181] P. Tamain, H. Bufferand, G. Ciraolo, C. Colin, P. Ghendrih, F. Schwander and E. Serre. "3D Properties of Edge Turbulent Transport in Full-Torus Simulations and their Impact on Poloidal Asymmetries". *Contributions to Plasma Physics*, **54**(4-6), 555–559, 2014. ISSN 15213986. URL http://dx.doi.org/10.1002/ctpp.201410017.

[182] B. D. Dudson and J. Leddy. "Hermes: Global plasma edge fluid turbulence simulations". *Plasma Physics and Controlled Fusion*, **59**(5), 054010, 2017. ISSN 13616587. URL http://dx.doi.org/10.1088/1361-6587/aa63d2.

[183] A. Stegmeir, D. Coster, A. Ross, O. Maj, K. Lackner and E. Poli. "GRILLIX: A 3D turbulence code based on the flux-coordinate independent approach". *Plasma Physics and Controlled Fusion*, **60**(3), 035005, 2018. ISSN 13616587. URL http://dx.doi.org/10.1088/1361-6587/aaa373.

[184] F. Nespoli, P. Tamain, N. Fedorczak, G. Ciraolo, D. Galassi, R. Tatali, E. Serre, Y. Marandet, H. Bufferand and P. Ghendrih. "3D structure and dynamics of filaments in turbulence simulations of WEST diverted plasmas". *Nuclear Fusion*, **59**(9), 1–25, 2019. URL http://dx.doi.org/10.1088/1741-4326/ab2813.

[185] F. Riva, C. Colin, J. Denis, L. Easy, I. Furno, J. Madsen, F. Militello, V. Naulin, A. H. Nielsen, J. M. Olsen *et al.* "Blob dynamics in the TORPEX experiment: A multi-code validation". *Plasma Physics and Controlled Fusion*, **58**(4), 044005, 2016. ISSN 13616587. URL http://dx.doi.org/10.1088/0741-3335/58/4/044005.

[186] W. Morris. "MAST The Spherical Tokamak". *Europhysics News*, **29**(6), 233–234, 1998. ISSN 0531-7479. URL http://dx.doi.org/10.1007/s00770-998-0233-7.

[187] F. Militello, B. Dudson, L. Easy, A. Kirk and P. Naylor. "On the interaction of scrape off layer filaments". *Plasma Physics and Controlled Fusion*, **59**(12), 125013, 2017. ISSN 13616587. URL http://dx.doi.org/10.1088/1361-6587/aa9252.

[188] J. R. Harrison. *Characterisation of detached plasmas on the MAST tokamak*. Ph.D. thesis, University of York, 2010. URL http://dx.doi.org/10.1016/j.jnucmat.2010.12.226.

[189] E. Havlíčková, J. Harrison, B. Lipschultz, G. Fishpool, A. Kirk, A. Thornton, M. Wischmeier, S. Elmore and S. Allan. "SOLPS analysis of the MAST-U divertor with the effect of heating power and pumping on the access to detachment in the Super-x configuration". *Plasma Physics and Controlled Fusion*, **57**(11), 115001, 2015. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/57/11/115001.

[190] M. Cox. "Mega Amp Spherical Tokamak". *Fusion Engineering and Design*, **46**(2-4), 397–404, 1999. ISSN 09203796. URL http://dx.doi.org/10.1016/S0920-3796(99)00031-9.

[191] P. REBUT, R. BICKERTON and B. KEEN. "The Joint European Torus: installation, first results and prospects". *Nuclear Fusion*, **25**, 1011, 1985. URL http://iopscience.iop.org/0029-5515/25/9/003/pdf/0029-5515_25_9_003.pdf.

[192] R. Aymar, P. Barabaschi and Y. Shimomura. "The ITER design". *Plasma Physics and Controlled Fusion*, **44**, 519–565, 2002. ISSN 1098-6596. arXiv:1011.1669v3, URL http://dx.doi.org/10.1017/CBO9781107415324.004.

[193] R. S. Hemsworth, D. Boilson, G. Serianni, P. Zaccaria, D. Shah, J. Graceffa, M. Singh, H. P. L. de Esch, F. Geli, M. D. Palma *et al.* "Overview of the design of the ITER heating neutral beam injectors". *New Journal of Physics*, **19**(2), 025005, 2017. URL http://dx.doi.org/10.1088/1367-2630/19/2/025005.

[194] S. Silburn, J. Harrison, M. Smithies, A. Wynn, T. Farley and J. Cavalier. "Calcam (Version 1.9.4)", 2018. URL http://dx.doi.org/10.5281/zenodo.1478555.

[195] G. Bradski. "The OpenCV Library". *Dr. Dobb's Journal of Software Tools*, 2000.

[196] J. Heikkila and O. Silven. "A Four-step Camera Calibration Procedure with Implicit Image Correction". *In* "Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition", San Juan, 1997. ISSN 1063-6919. URL http://dx.doi.org/10.1109/CVPR.1997.609468.

[197] J. Mlynar, S. Coda, A. Degeling, B. P. Duval, F. Hofmann, T. Goodman, J. B. Lister, X. Llobet and H. Weisen. "Investigation of the consistency of magnetic and

soft x-ray plasma position measurements on TCV by means of a rapid tomographic inversion algorithm". *Plasma Physics and Controlled Fusion*, **45**(2), 169–180, 2003. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/45/2/308.

[198] D. Mazon, D. Vezinet, D. Pacella, D. Moreau, L. Gabelieri, A. Romano, P. Malard, J. Mlynar, R. Masset and P. Lotte. "Soft x-ray tomography for real-time applications: Present status at Tore Supra and possible future developments". *Review of Scientific Instruments*, **83**(6), 063505, 2012. ISSN 00346748. URL http://dx.doi.org/10.1063/1.4730044.

[199] A. Wingen, M. W. Shafer, E. A. Unterberg, J. C. Hill and D. L. Hillis. "Regularization of soft-X-ray imaging in the DIII-D tokamak". *Journal of Computational Physics*, **289**, 83–95, 2015. ISSN 10902716. URL http://dx.doi.org/10.1016/j.jcp.2015.02.040.

[200] A. Huber, K. McCormick, P. Andrew, M. R. de Baar, P. Beaumont, S. Dalley, J. Fink, J. C. Fuchs, K. Fullard, W. Fundamenski *et al.* "Improved radiation measurements on JET - First results from an upgraded bolometer system". *Journal of Nuclear Materials*, **363-365**(1-3), 365–370, 2007. ISSN 00223115. URL http://dx.doi.org/10.1016/j.jnucmat.2007.01.124.

[201] J. Jang, B. J. Peterson, S. Oh, K. Mukai, S. H. Hong and W. Choe. "Reconstruction of radiation profiles near the plasma boundary using an infrared imaging video bolometer in KSTAR". *Review of Scientific Instruments*, **89**(10), 10E111, 2018. ISSN 10897623. URL http://dx.doi.org/10.1063/1.5038904.

[202] M. Carr, A. Meakins, M. Bernert, P. David, C. Giroud, J. Harrison, S. Henderson, B. Lipschultz and F. Reimold. "Description of complex viewing geometries of fusion tomography diagnostics by ray-tracing". *Review of Scientific Instruments*, **89**(8), 083506, 2018. ISSN 10897623. URL http://dx.doi.org/10.1063/1.5031087.

[203] W. H. Meyer, S. L. Allen, C. M. Samuell and M. E. Fenstermacher. "Tomographic analysis of tangential viewing cameras (invited)". *Review of Scientific Instruments*, **89**(10), 10K110, 2018. ISSN 10897623. URL http://dx.doi.org/10.1063/1.5038586.

[204] M. Carr, A. Meakins, S. A. Silburn, J. Karhunen, M. Bernert, C. Bowman, A. Callarelli, P. Carvalho, C. Giroud, J. R. Harrison *et al.* "Physically principled reflection models applied to filtered camera imaging inversions in metal walled fusion machines". *Review of Scientific Instruments*, **90**(4), 043504, 2019. ISSN 0034-6748. URL http://dx.doi.org/10.1063/1.5092781.

[205] D. Iraji, D. Ricci, G. Granucci, S. Garavaglia, I. Furno, A. Cremona and D. Minelli. "Imaging of turbulent structures and tomographic reconstruction of TORPEX plasma emissivity". *Fusion Science and Technology*, **62**(3), 428–435, 2012. ISSN 15361055. URL http://dx.doi.org/10.13182/FST12-A15342.

[206] L. C. Ingesson, B. Alper, B. J. Peterson and J. C. Vallet. "Chapter 7: Tomography diagnostics: Bolometry and soft-X-ray detection". *Fusion Science and Technology*, **53**(2), 528–576, 2008. ISSN 15361055. URL http://dx.doi.org/10.13182/FST53-528.

[207] L. Montgomery Smith, D. R. Keefer and S. I. Sudharsanan. "Abel inversion using transform techniques". *Journal of Quantitative Spectroscopy and Radiative Transfer*, **39**(5), 367–373, 1988. ISSN 00224073. URL http://dx.doi.org/10.1016/0022-4073(88)90101-X.

[208] C. A. Michael. *Doppler spectroscopy of argon plasmas in H-1NF using a coherence imaging camera*. Ph.D. thesis, ANU, 2003. URL https://openresearch-repository.anu.edu.au/handle/1885/151238.

[209] Y. C. Ghim, A. R. Field, D. Dunai, S. Zoletnik, L. Bardóczi and A. A. Schekochihin. "Measurement and physical interpretation of the mean motion of turbulent density patterns detected by the beam emission spectroscopy system on the mega amp spherical tokamak". *Plasma Physics and Controlled Fusion*, **54**(9), 2012. ISSN 07413335. URL http://dx.doi.org/10.1088/0741-3335/54/9/095012.

[210] T. Munsat and S. J. Zweben. "Derivation of time-dependent two-dimensional velocity field maps for plasma turbulence studies". *Review of Scientific Instruments*, **77**(10), 2006. ISSN 00346748. URL http://dx.doi.org/10.1063/1.2356851.

[211] C. Tomasi and R. Manduchi. "Bilateral filtering for gray and color images". *In* "Sixth International Conference on Computer Vision", 839–846. Bombay, 1998. ISBN 81-7319-221-9. ISSN 1873622X. arXiv:1011.1669v3, URL http://dx.doi.org/10.1109/ICCV.1998.710815.

[212] B. M. ter Haar Romeny, L. M. J. Florack, M. de Swart, J. Wilting and M. A. Viergever. "Deblurring Gaussian Blur". *In* "Mathematical Methods in Medical Imaging III", 139–148. 1994. URL http://dx.doi.org/10.1117/12.179245.

[213] L. C. Appel and I. Lupelli. "Equilibrium reconstruction in an iron core tokamak using a deterministic magnetisation model". *Computer Physics Communications*, **223**, 1–17, 2018. ISSN 00104655. URL http://dx.doi.org/10.1016/j.cpc.2017.09.016.

[214] L. C. Appel, G. T. A. Huysmans, L. L. Lao, P. J. McCarthy, D. G. Muir, E. R. Solano, J. Storrs, D. Taylor and W. Zwingmann. "A Unified Approach to Equilibrium Reconstruction". *In* "33rd EPS Conference on Plasma Physics", Rome, 2006. ISBN 9781622763337. URL http://www.iop.org/Jet/fulltext/EFDC060307.pdf.

[215] R. C. Thurber, B. Aster and C. H. Borchers. *Parameter Estimation and Inverse Problems*. Elsevier, 3rd edition, 2013. ISBN 9780128046517. URL http://dx.doi.org/10.1016/b978-0-12-804651-7.00003-1.

[216] R. Rosipal and L. J. Trejo. "Kernel Partial Least Squares Regression in Reproducing Kernel Hilbert Space". *Journal of Machine Learning Research*, **1**, 97–123, 2001. ISSN 0003-6951. URL http://dx.doi.org/10.1162/15324430260185556.

[217] P. C. Hansen. "Numerical tools for analysis and solution of Fredholm integral equations of the first kind". *Inverse Problems*, **8**(6), 849–872, 1992. ISSN 02665611. URL http://dx.doi.org/10.1088/0266-5611/8/6/005.

[218] M. Anton, H. Weisen, M. J. Dutch, W. V. D. Linden, F. Buhlmann, P. Marmillod, P. Paris, R. Chavan and B. Marletaz. "X-ray tomography on the TCV tokamak". *Plasma Physics and Controlled Fusion*, **38**(11), 1849–1878, 1996. URL http://iopscience.iop.org/0741-3335/38/11/001.

[219] T. Hobiger, T. Kondo and Y. Koyama. "Constrained simultaneous algebraic reconstruction technique (C-SART)–a new and simple algorithm applied to ionospheric tomography". *Earth Planets and Space (EPS)*, **60**(2), 727–735, 2008. URL http://www.springerlink.com/index/T56625R7G486M711.pdf%5Cnhttp://www.terrapub.co.jp/journals/EPS/pdf/2008/6007/60070727.pdf.

[220] J. Demmel and W. Kahan. "Accurate Singular Values of Bidiagonal Matrices". *SIAM Journal on Scientific and Statistical Computing*, **11**(5), 873–912, 1990. ISSN 0196-5204. URL http://dx.doi.org/10.1137/0911052.

[221] M. G. H. Bell. "The estimation of origin-destination matrices by constrained generalised least squares". *Transportation Research Part B*, **25**(1), 13–22, 1991. ISSN 01912615. URL http://dx.doi.org/10.1016/0191-2615(91)90010-G.

[222] G. Golub and W. Kahan. "Calculating the Singular Values and Pseudo-Inverse of a Matrix". *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis*, **2**(2), 205–224, 2005. ISSN 0887-459X. URL http://dx.doi.org/10.1137/0702016.

[223] G. Strang. "Linear Algebra and Its Applications". *In* "Linear Algebra and Its Applications, Fourth Edition Gilbert Strang", volume 1. Thomson Brooks, Belmont, 4th edition, 2006. ISBN 0-03-010567-6.

[224] L. Ding, X. Luís Deán-Ben, C. Lutzweiler, D. Razansky and V. Ntziachristos. "Efficient non-negative constrained model-based inversion in optoacoustic tomography". *Physics in Medicine and Biology*, **60**(17), 6733–6750, 2015. ISSN 13616560. URL http://dx.doi.org/10.1088/0031-9155/60/17/6733.

[225] N. Beucher and S. Bucher. "Hierarchical Queues: general description and implementation in MAMBA image library", 2011. URL http://cmm.ensmp.fr/$\sim$beucher/publi/HQ_algo_desc.pdf.

[226] D. Carralero, G. Birkenmeier, H. Müller, P. Manz, P. DeMarne, S. Müller, F. Reimold, U. Stroth, M. Wischmeier and E. Wolfrum. "An experimental investigation of the high density transition of the scrape-off layer transport in ASDEX Upgrade". *Nuclear Fusion*, **54**(12), 123005, 2014. ISSN 0029-5515. arXiv:1407.3618v1, URL http://dx.doi.org/10.1088/0029-5515/54/12/123005.

[227] J. Ball, F. I. Parra, M. Landreman and M. Barnes. "Optimized up-down asymmetry to drive fast intrinsic rotation in tokamaks". *Nuclear Fusion*, **58**(2), 026003, 2018. ISSN 17414326. URL http://dx.doi.org/10.1088/1741-4326/aa9a50.

[228] D. Comaniciu, V. Ramesh and P. Meer. "Kernel-based object tracking". *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **25**(5), 564–577, 2003. ISSN 01628828. URL http://dx.doi.org/10.1109/TPAMI.2003.1195991.

[229] G. M. Quénot, J. Pakleza and T. A. Kowalewski. "Particle image velocimetry with optical flow". *Experiments in Fluids*, **25**(3), 177–189, 1998. ISSN 07234864. URL http://dx.doi.org/10.1007/s003480050222.

[230] M. J. Black and A. D. Jepson. "EigenTracking: Robust matching and tracking of articulated objects using a view-based representation". *International Journal of Computer Vision*, **26**(1), 63–84, 1998. ISSN 09205691. URL http://dx.doi.org/10.1007/BFb0015548.

[231] F. Nespoli, B. Labit, I. Furno, F. Musil, F. D. Halpern, P. Ricci and F. Riva. "Blob properties in nonlinear simulations of the TCV Scrape-Off Layer". *Plasma Physics and Controlled Fusion*, **59**(5), 055009, 2017.

[232] R. Rossi, A. Malizia, L. A. Poggi, J. F. Ciparisse, E. Peluso and P. Gaudio. "Flow Motion and Dust Tracking Software for PIV and Dust PTV". *Journal of Failure Analysis and Prevention*, **16**(6), 951–962, 2016. ISSN 15477029. URL http://dx.doi.org/10.1007/s11668-016-0204-0.

[233] A. Bruhn, J. Weickert and C. Schnörr. "Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods". *International Journal of Computer Vision*, **61**(3), 1–21, 2005. ISSN 09205691. 1035417, URL http://dx.doi.org/10.1023/B:VISI.0000045324.43199.43.

[234] B. K. Horn and B. G. Schunck. "Determining Optical Flow". *Artificial Intelligence*, **17**(1981), 185–203, 1981.

[235] E. Limpert, W. A. Stahel and M. Abbt. "Log-normal Distributions across the Sciences: Keys and Clues". *BioScience*, **51**(5), 341–352, 2001. ISSN 0006-3568. URL http://dx.doi.org/10.1641/0006-3568(2001)051[0341:LNDATS]2.0.CO;2.

[236] B. Cannas, S. Carcangiu, A. Fanni, T. Farley, F. Militello, A. Montisci, F. Pisano, G. Sias and N. Walkden. "Towards an automatic filament detector with a Faster R-CNN on MAST-U". *Fusion Engineering and Design*, **146**(A), 374–377, 2019. ISSN 09203796. URL http://dx.doi.org/10.1016/j.fusengdes.2018.12.071.

[237] M. Sokolova, N. Japkowicz and S. Szpakowicz. "Beyond accuracy, F-score and ROC: A family of discriminant measures for performance evaluation". *Advances in Artificial Intelligence*, **4304**, 24–29, 2006. URL http://dx.doi.org/https://doi-org.liverpool.idm.oclc.org/10.1007/11941439_114.

[238] A. Kirk, N. Ben Ayed, G. Counsell, B. Dudson, T. Eich, A. Herrmann, B. Koch, R. Martin, A. Meakins, S. Saarelma *et al.* "Filament structures at the plasma edge on MAST". *Plasma Physics and Controlled Fusion*, **48**(12B), B433–B441, 2006. ISSN 0741-3335. URL http://dx.doi.org/10.1088/0741-3335/48/12b/S41.

[239] O. Garcia, J. Horacek, R. Pitts, a.H Nielsen, W. Fundamenski, V. Naulin and J. J. Rasmussen. "Fluctuations and transport in the TCV scrape-off layer". *Nuclear Fusion*, **47**(7), 667–676, 2007. ISSN 0029-5515. URL http://dx.doi.org/10.1088/0029-5515/47/7/017.

[240] R. Kube and O. E. Garcia. "Convergence of statistical moments of particle density time series in scrape-off layer plasmas". *Physics of Plasmas*, **22**(1), 2015. ISSN 10897674. 1409.8319, URL http://dx.doi.org/10.1063/1.4905513.

[241] H. P. Summers. "ADAS User Manual, v2.6", 2004. URL http://www.adas.ac.uk/manual.php.

[242] M. Carr, A. Meakins, A. Baciero, M. Bernert, A. Callarelli, A. Field, C. Giroud, J. Harrison, N. Hawkes, S. Henderson *et al.* "Towards integrated data analysis of divertor diagnostics with ray-tracing". *In* "44th EPS Conference on Plasma Physics", June 2017, 2017. 2017.

[243] G. H. Golub, M. Heath, G. Wahba and G. H. Golub. "Generalized Cross-Validation as a Method for Choosing a Good Ridge Parameter". *Technometrics*, **21**(2), 215–223, 1979.

[244] D. Calvetti, S. Morigi, L. Reichel and F. Sgallari. "Tikhonov regularization and the L-curve for large discrete ill-posed problems". *Journal of Computational and Applied Mathematics*, **123**(1-2), 423–446, 2000. ISSN 03770427. URL http://dx.doi.org/10.1016/S0377-0427(00)00414-3.