

UNIVERSITY OF STRATHCLYDE  
DEPARTMENT OF PHYSICS

**Advancing Laser-Driven Ion  
Acceleration: Optimising with  
Machine Learning and  
Investigating Sources of  
Instability**



by

**Ewan James Dolier**

In partial fulfilment of the requirements for the degree of  
Doctor of Philosophy in Physics

**2024**

# Copyright Declaration

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed:

Date:

# Abstract

This thesis reports on numerical and experimental investigations of proton acceleration driven by intense laser pulse ( $\sim 10^{21}$  Wcm $^{-2}$ ) interactions with foil targets. The resultant beam of protons has unique properties compared to those produced in conventional accelerators. As a result, these novel accelerator sources are expected to have an important impact on both research and societal applications. For this to be realised, key properties such as the maximum proton energy and laser-to-proton energy conversion efficiency must be improved, in addition to the beam reproducibility and stability from one laser shot to another.

Progress towards this goal is presented in two main investigations, the first of which involves the development of methods to automatically optimise properties of laser-driven proton beams in numerical simulations, advancing beyond conventional grid-search optimisation. Optimal values for laser energy, pulse duration, target foil thickness, and pre-plasma density scale length are identified with  $\sim 200$  fewer data-points, corresponding to a reduction of  $\sim 48$  days in simulation time, by employing a newly developed code called BISHOP with an integrated machine learning (ML) model. This four parameter optimisation is made feasible because of this technique, and is found to double the maximum energy of protons produced in the target normal sheath acceleration (TNSA) regime, compared to optimising for only the laser energy and pulse duration. The ML model also uncovered novel optimal pre-plasma conditions that increase laser energy coupling to fast electrons in the pre-plasma, whilst mitigating their overall divergence upon propagation through the target foil, thus increasing the sheath field strength on the target normal axis, and, as a result, the maximum energy of TNSA protons.

A second, numerical and experimental investigation, demonstrates that the

onset of relativistic self-induced transparency (RSIT) enhances proton energies beyond those that are achieved solely via TNSA, but with less stability when RSIT is induced at an optimal interaction time that maximises proton energies. This sensitivity is significant when factoring in shot-to-shot fluctuations in laser energy and pulse duration, demonstrated to occur in experiments at high-power laser facilities. This exacerbates the known sensitivity of optimised RSIT-enhanced proton acceleration to target foil thickness and laser temporal-intensity contrast. Early onset of RSIT deoptimises this regime in terms of proton energy, but makes it less susceptible to fluctuations in laser pulse parameters, whilst still enabling proton energy enhancement compared to the TNSA regime.

Together, these investigations contribute to the development of laser-driven proton sources towards applications, by identifying new pathways to improve proton beam properties towards required specifications, whilst increasing understanding of how to produce these properties consistently.

# Acknowledgements

In the final weeks of writing this thesis I've had many moments to reflect on the journey, from beginning in late 2019 until now. To appreciate the opportunities I've had, those who've granted them, and those who have supported me along the way. It's difficult to fit 4 years, in some cases 26 years, of gratitude into a few pages. I'm not sure I'll do everyone justice. But I'll give it a go.

Firstly, I'm grateful to EPSRC for funding my PhD studentship, and to the British taxpayer who funds them. Without this funding I would never have embarked on a PhD. I'd like to thank Prof. Paul McKenna, who primarily secured this funding, and other funding which allowed me to take part in experiments not discussed in this thesis, and to present my research at national and international conferences. These experiences have developed my skills as a physicist, researcher, communicator, and as a person, and have enabled me to visit interesting places in the U.K, Germany, Portugal, and the Czech Republic. All of this was possible due to Paul's constant support of my work, and genuine interest in my development as a physicist and person. Clearly, this thesis would not be possible without Paul's support, but less obviously it would have been a poorer standard without Paul reviewing every section. Every comment from Paul is insightful and to the point. He is equally attuned to the 'big picture' of research, the 'zoomed in picture' of specific investigations, and the minor details of wording and grammar. I find this ability to 'see the forest from the trees' to be quite unique, and something to strive towards. Similarly, I admire Paul's dedication, professionalism, patience, empathy, knowledge and experience, and strive towards embodying these traits in my own research career. That is a long way to say, I am extremely lucky to have had Prof. Paul McKenna as a supervisor and mentor.

So too am I grateful for the mentoring and support of Dr. Ross Gray, who portrays many of these same traits. Ross has supported me in the ‘day-to-day’ of my PhD, and also reviewed my entire thesis. For this I am extremely grateful, and if Paul is reading, he will be too, as amendments suggested by Ross vastly improved Paul’s reading material. Ross’ day-to-day support is, in my experience, quite unique in academia, and something I am very grateful to have had. Working so closely with Ross, I am often impressed by his knowledge on a wide range of subjects. From laser-plasma interactions to virtually every area of physics, from numerical simulation and experimental techniques to computing systems and machine learning. Beyond this I am immensely grateful for Ross’ empathy, patience, and level-headedness, which is well needed in the sometimes stressful world of academia. I feel very fortunate to call Dr. Ross Gray a mentor, and strive towards the standard he sets.

I am grateful too, for the postdoctoral researchers who have mentored me throughout my PhD. Firstly, I am grateful to Dr. Martin King. For someone who never considered themselves to be ‘good with computers’ to submit a thesis with so many numerical simulations performed on national supercomputers is a testament to Martin’s knowledge on these matters, and his patience and empathy in guiding students through the process. Beyond this, Martin’s knowledge of physics is to be admired, and it vastly improved the insights presented in this thesis. Next, I am grateful to those who have taught me everything I know about experimental work. First and foremost, I am extremely grateful to Dr. Robbie Wilson. Since supervising my undergraduate project back in 2019, Robbie has taken me under his wing, and showed me the ropes of experimental laser-plasma physics. It’s said that it takes 10,000 hours to become a master in a given field. Robbie has certainly put in these hours, and has spent several thousand mentoring me on experiments in Scotland, England, and Germany, often as the deputy target area operator to his TAO. Robbie’s deep passion for physics, and generosity in sharing his knowledge and experience, is something to admire, and to strive for in my own research career. I am also grateful to Dr. Tim Frazer, who has taught me more about experimental work and laser-plasma physics than I think

he would ever give himself credit for, a testament to his humble nature. I am grateful for this, and for the very funny situations we got ourselves into on long, sleep deprived days on experiment. Finally, I am grateful to staff at the central laser facility, where I gained most of my experimental experience. There are too many to name, but I will express my gratitude for Dr. Tom Dzelzainis. Link scientists are crucial to any successful experiment, but in Tom's case this was more true than most. Tom always went above and beyond his responsibilities, working with us from 9am to well beyond 5pm, often late into the night, and even occasionally on weekends. Tom was therefore truly one of the team, and I am immensely grateful for his passion, dedication, knowledge, and experience. I credit Tom alongside Robbie as my most important mentor in experimental laser-plasma physics, even though this was well beyond his job description. Finally, I am grateful for Tom's incredible sense of humour and fun, which lightened the most arduous and tiring situations on experiment.

On this note, I owe a great deal of thanks to my fellow students whom I've been fortunate enough to work alongside. To Alana Horne, Antonio Lofrese, Ben Torrance, Christopher McQueen, Ewan Bacon, Jack Goodman, Jesel Patel, Maia Peat, Matthew Alderton and Radhika Nayli, I am extremely grateful for the work you have put in alongside me in experiments. More importantly, I am very glad to call you all friends. Our conversations and laughter have carried me throughout the PhD.

This leads me nicely to my friends and family, who have supported me throughout the PhD and, more importantly, throughout my life. Firstly, I feel incredibly lucky to have a large group of friends away from work. To Aaron S, Alana H, Cameron H, Craig H, Craig P, Daniel M, Dean M, Euan H, Justin Y, Lewis F, Matthew G, Megan H, Nathan S, Ross M, Sam M and Sarah T, thank you for the days out, and the nights out, for the golf and 5-a-side football games. Thank you for the redzone Sundays, and for the holidays abroad. Your friendship, and that of my fellow PhD students, has given me important respite from my thesis, and you can all credit yourself with its completion. In return, I'll do you all a favour by never banging on about it again.

And finally, to my family. To my Grans Margaret and Vera, to my Grandad Jimmy and my Papa Richard, to my Aunts Anne, Eleanor, Heather, Irene and Linda, and my Uncles Mark, Richard, Stuart and Willie. To my Cousins Andrea, Craig, David, Jonathon, Lisa, Rebekah, and Stuart. Thank you for everything you have done for me. Without you I would not be who I am today. In particular, I am grateful to my Aunt Irene and Gran Vera, whose time spent reading with me, and taking me to museums as a kid, made learning fun. I am also particularly grateful to my Cousin Lisa, who taught me that even the most difficult adversities can be overcome, and who has always inspired me.

Finally, I am most grateful of all to my close family. To my dad Stewart, mum Cynthia, and sister Rachel, without your love and support I would never have come this far. To Rachel, thank you for being a wonderful sister. Your sense of humour and fun has kept me going throughout the PhD. When things have been difficult or weren't going my way, you always remind me that really, I have it pretty good. To my mum and dad, you have always supported me without hesitation, in every aspect of my life. You worked hard all of your life to allow me the opportunities I've had throughout mine, without which I would never be writing the acknowledgements of a PhD thesis. You gave me all of this support but never, ever, pushed me. From exams in high school until now, you always believed I'd 'get it done', that I'd 'get there in the end'. The confidence you instilled in me, the work ethic to do things properly, and not to cut corners, has been fundamental throughout my life and PhD. So too have been the constant cups of tea, and biscuits, sometimes snuck into my room during a zoom meeting by way of crawling on the floor. In the final weeks of long, difficult experiments, where energy and motivation was flagging, what got me through was the thought of returning through the gates at Glasgow airport to my dad, waiting with a smile on his face, and with a big hug. Then returning home, to a cup of tea, a slice of toast, and a long hug from my mum, fighting to stay awake at 1am because she insisted on welcoming me home. Mum, dad, Rachel, I am extremely lucky to be part of this family, and I love you very much.

# Role of Author

The author had a primary role in planning, executing, and analysing data from numerical simulations and experiments, the results of which are presented in this thesis. This was achieved as part of a team of researchers led by Prof. P. McKenna, as is now briefly summarised:

- **Chapter 4:** The author assisted Dr. R.J. Gray in developing and testing the BISHOP code, and led testing of different machine learning algorithms, eventually integrating the open source Bayesian optimisation algorithm from scikit-learn into BISHOP. The author initiated all BISHOP driven simulation campaigns, in consultation with Dr. M. King. and Dr. R.J. Gray. The author analysed all simulation data, and collaborated closely with Dr. R. Wilson, Dr. M. King, Dr. R.J. Gray, and Prof. P. McKenna to interpret and publish the findings.
- **Chapter 5:** The author reviewed results from a 2016 experiment, primarily setup by Dr. R.J. Gray, Dr. R.J. Dance, and Dr. N. M.H. Butler. Dr. A. Higginson also had a significant role in the setup, and led analysis of the results, interpreting and publishing the findings in collaboration with the mentioned authors amongst a larger team led by Prof. P. McKenna. The author's review of these findings motivated objectives on a 2021 experiment, setup by the author and Dr. R. Wilson as deputy target area operator (DTAO), and TAO, respectively. Central laser facility staff supported this experiment, and Ms. A. Horne and Mr. J. Goodman set up diagnostics. As did Mr. E. F.J. Bacon, who assembled and analysed the primary dosimetry film stack diagnostic, designed by the author and Dr. R. Wilson. The author had a primary role in analysing experimental results, performing and

analysing complementary simulation data, and interpreting overall findings, in close collaboration with Dr. R. Wilson, Dr. M. King, Dr. R.J. Gray, and Prof. P. McKenna.

# Publications

**Publications directly resulting from the investigations presented in this thesis:**

1. **E.J Dolier**, M King, R Wilson, R.J Gray, and P McKenna. Multi-parameter Bayesian optimisation of laser-driven ion acceleration in particle-in-cell simulations. *New Journal of Physics*, 24(7):073025, 2022.
2. J Goodman, M King, **E.J Dolier**, R Wilson, R.J Gray and P McKenna. Optimization and control of synchrotron emission in ultraintense laser–solid interactions using machine learning. *High Power Laser Science and Engineering*. 11:34, 2023.

**Additional publications resulting from the author’s PhD:**

1. **E.J Dolier**, R Wilson, T.P Frazer, J Goodman, A Lofrese, A Horne, M King, R.J Gray, P McKenna, T Dzelzainis and D Neely. Accounting for spectral bandwidth broadening due to self-phase modulation in measurements of laser pulse transmission through a dense plasma. *Central Laser Facility Annual Report*, 2020.
2. M King, R Wilson, E.F Bacon, **E.J Dolier**, T.P Frazer, J Goodman, R.J Gray, and P McKenna., Perspectives on laser-plasma physics in the relativistic transparency regime. *The European Physical Journal A*. 59(6):132, 2023.
3. J.K Patel, C.D Armstrong, R Wilson, M Alderton, **E.J Dolier**, T.P Frazer, A Horne, A Lofrese, M Peat, M Woodward, B Zielbauer, R.J Clarke, R

Deas, P.P Rajeev, R.J Gray, and P McKenna., A scintillating fiber imaging spectrometer for active characterisation of laser-driven proton beams. Submitted to *High Power Laser Science and Engineering*, 2024.

# Contents

<b>Abstract</b>	<b>i</b>
<b>Acknowledgements</b>	<b>iii</b>
<b>Role of Author</b>	<b>vii</b>
<b>Publications</b>	<b>ix</b>
<b>List of Figures</b>	<b>xvii</b>
<b>List of Tables</b>	<b>xviii</b>
<b>List of Commonly Used Variables</b>	<b>xix</b>
<b>List of Commonly Used Acronyms</b>	<b>xx</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Brief history of ion acceleration . . . . .	1
1.2 Machine learning to advance laser-driven ion acceleration . . . . .	6
1.3 Applications of radiation produced in laser-solid interactions . . . . .	8
1.3.1 Ultra-high dose rate proton therapy . . . . .	8
1.4 Thesis outline . . . . .	9
<b>2 Fundamentals of laser-solid interactions</b>	<b>11</b>
2.1 Properties of high-power lasers . . . . .	11
2.2 Laser-induced ionisation processes . . . . .	14
2.2.1 Multi-photon ionisation . . . . .	16
2.2.2 Barrier suppression ionisation . . . . .	16

2.3	Laser-induced plasma formation . . . . .	18
2.4	Response of electrons to intense laser light . . . . .	22
2.4.1	Electron motion in response to an infinite planar electro- magnetic wave . . . . .	22
2.4.2	Electron motion in response to the laser ponderomotive force	25
2.4.3	Laser-to-electron energy absorption . . . . .	30
2.5	Electron transport in overdense plasma . . . . .	34
2.6	Laser-driven ion acceleration . . . . .	39
2.6.1	Target normal sheath acceleration . . . . .	40
2.6.2	Radiation pressure acceleration . . . . .	45
2.6.3	Acceleration enhanced by relativistic self-induced transparency	51
2.7	Summary . . . . .	56
<b>3</b>	<b>Methodology</b>	<b>58</b>
3.1	Creation and amplification of short duration, high-power laser pulses at the Vulcan-PW facility . . . . .	59
3.2	Plasma mirrors . . . . .	66
3.3	Laser focal spot improvement and characterisation . . . . .	68
3.4	Experimental diagnostics . . . . .	71
3.4.1	Optical diagnostics . . . . .	72
3.4.2	Particle diagnostics . . . . .	75
3.5	Particle-in-cell simulations . . . . .	81
3.6	Summary . . . . .	86
<b>4</b>	<b>Multi-parameter Bayesian optimisation of laser-driven ion accel- eration in particle-in-cell simulations</b>	<b>87</b>
4.1	Introduction . . . . .	87
4.2	Development and operation of the BISHOP Code . . . . .	89
4.3	Selecting a machine learning-based optimisation technique . . . . .	93
4.4	Bayes theory and Gaussian process regression . . . . .	97
4.5	Bayesian optimisation . . . . .	101
4.6	Results . . . . .	103

4.6.1	Automated grid-search of laser parameters to optimise TNSA	103
4.6.2	Efficiency of Bayesian optimisation versus grid-search . . .	108
4.6.3	Multi-dimensional optimisation of proton acceleration in the TNSA regime . . . . .	113
4.6.4	Investigating optimal laser and target conditions for TNSA of protons in a multi-dimensional regime . . . . .	118
4.6.5	Novel optimisation of maximum proton energy with pre- plasma density scale length . . . . .	120
4.7	Conclusion . . . . .	126
<b>5</b>	<b>Stability of enhanced ion acceleration in the relativistic self- induced transparency regime</b>	<b>128</b>
5.1	Introduction . . . . .	128
5.2	Methodology . . . . .	131
5.2.1	Experimental set-up . . . . .	131
5.3	Results . . . . .	133
5.3.1	Fluctuations in Vulcan-PW laser properties . . . . .	133
5.3.2	Investigating the stability of proton acceleration in the TNSA and RSIT-enhanced regimes . . . . .	137
5.3.3	Influence of laser energy and pulse duration fluctuations on maximum proton energies . . . . .	140
5.3.4	Comparing the maximum proton energy sensitivity between experiment and simulations . . . . .	144
5.3.5	Diagnosing the sensitivity of RSIT-enhanced proton accel- eration to fluctuating laser conditions . . . . .	152
5.4	Conclusion . . . . .	156
<b>6</b>	<b>Conclusions and outlooks for future research</b>	<b>159</b>
6.1	Investigation into multi-parameter Bayesian optimisation of laser- driven ion acceleration in particle-in-cell simulations . . . . .	159
6.2	Investigation into the stability of enhanced ion acceleration in the relativistic self-induced transparency regime . . . . .	162

6.3 Final remarks and outlook for future research . . . . . 164

**Bibliography** **198**

# List of Figures

1.1	History of laser-driven ion acceleration . . . . .	3
2.1	Laser pulse Gaussian spatio-temporal intensity profile . . . . .	12
2.2	Realistic laser temporal-intensity profile . . . . .	13
2.3	Laser-driven ionisation mechanisms . . . . .	15
2.4	Illustration of charge shielding in an expanding plasma . . . . .	19
2.5	Influence of the laser ponderomotive force on plasma electrons . .	25
2.6	Laser pulse interacting with expanding plasma in a 2D particle-in-cell (PIC) simulation . . . . .	27
2.7	Illustration of laser-to-plasma electron energy absorption mechanisms	31
2.8	Fast electron energy spectrum from a 2D PIC simulation . . . . .	35
2.9	Illustration of fast electron transport in a solid target foil . . . . .	37
2.10	Illustration of target normal sheath acceleration (TNSA) of ions .	40
2.11	TNSA of ions in a 2D PIC simulation . . . . .	43
2.12	Illustration of radiation pressure acceleration (RPA) of ions . . . .	45
2.13	RPA of ions (hole-boring) in a 2D PIC simulation . . . . .	48
2.14	RPA of ions (light-sail) in a 2D PIC simulation . . . . .	50
2.15	RPA-TNSA hybrid acceleration of ions enhanced by relativistic self-induced transparency (RSIT) in a 2D PIC simulation . . . . .	53
2.16	Maximum energy of ions accelerated via different mechanisms in 2D PIC simulations . . . . .	55
3.1	Illustration of pulse amplification and temporal compression of the Vulcan-PW laser . . . . .	60
3.2	Photograph and operating principle of a plasma mirror . . . . .	67

3.3	Measured spatial-intensity profile of the Vulcan-PW laser focal spot	69
3.4	Illustration of diagnostics used to characterise laser-solid interactions	72
3.5	Illustration of a Thomson parabola ion spectrometer . . . . .	76
3.6	Illustration of the radiochromic film stack diagnostic technique . .	79
3.7	Diagram showing steps of the particle-in-cell (PIC) simulation method	83
3.8	Comparison of results from well-resolved and poorly-resolved PIC simulations . . . . .	85
4.1	Diagram of files required to operate the BISHOP code . . . . .	91
4.2	Diagram showing the operating steps of the BISHOP code . . . .	93
4.3	Illustration of Gaussian process regression for modelling data . . .	98
4.4	Effect of varying model parameters in Gaussian process regression	99
4.5	Illustration of Bayesian optimisation . . . . .	102
4.6	Effect of varying laser energy and pulse duration on fast electron temperature and TNSA maximum proton energy, identified via grid-search of 2D PIC simulations . . . . .	106
4.7	Fast electron energy spectrum in a 2D PIC simulation . . . . .	107
4.8	Performance of different Bayesian optimisation models in identify- ing optimal laser energy and pulse duration for TNSA maximum proton energies . . . . .	110
4.9	Bayesian optimisation vs grid-search for identifying optimal laser energy and pulse duration for TNSA maximum proton energies . .	113
4.10	Effect of varying laser energy, pulse duration, target foil thickness, and pre-plasma density scale length on TNSA maximum proton energies, identified via Bayesian optimisation . . . . .	115
4.11	2D PIC simulation results confirming TNSA is primarily responsi- ble for the highest energy protons being optimised . . . . .	117
4.12	Convergence of the Bayesian optimisation algorithm to the opti- mum maximum proton energy . . . . .	119
4.13	Effect of varying pre-plasma density scale length on the temporal evolution of sheath field strength and maximum proton energy . .	122

4.14	Effect of varying pre-plasma density scale length on overall laser-to-electron energy absorption, strength and transverse size of the sheath field, and maximum proton energy . . . . .	123
4.15	Location of the relativistically corrected critical density surface in pre-plasma of varying scale length . . . . .	125
5.1	Illustration of the experimental set up at the Vulcan-PW laser facility	131
5.2	Effect of varying ultra-thin foil target thickness on maximum proton energy, recorded during experiment A . . . . .	134
5.3	Fluctuations in Vulcan-PW laser properties during experiment A .	135
5.4	Effect of varying ultra-thin foil target thickness on maximum proton energy, recorded during experiment B . . . . .	137
5.5	Fluctuations in Vulcan-PW laser properties during experiment B .	139
5.6	Fluctuation and level of uncertainty in the Vulcan-PW pulse duration	140
5.7	Effect of varying Vulcan-PW laser intensity on maximum proton energies in the TNSA and RSIT-enhanced acceleration regimes . .	141
5.8	Effect of varying ultra-thin foil target thickness on maximum proton energy in 2D PIC simulations . . . . .	146
5.9	Stability of RSIT-enhanced proton acceleration to 8% experimental and simulated fluctuation of laser energy and pulse duration .	147
5.10	Stability of RSIT-enhanced proton acceleration to 35% experimental and simulated fluctuation of laser energy and pulse duration .	149
5.11	Stability of TNSA, RSIT-enhanced, and early onset RSIT proton acceleration to 8% experimental and simulated fluctuation of laser energy and pulse duration . . . . .	150
5.12	Effect of fluctuating simulated laser pulse duration on the temporal evolution and stability of RSIT-enhanced proton acceleration . . .	153
5.13	Effect of random simulation fluctuations on the temporal evolution and stability of RSIT-enhanced proton acceleration . . . . .	154
5.14	Effect of random simulation fluctuations on the RSIT-enhanced accelerating field structure experienced by protons . . . . .	155

# List of Tables

4.1	Operating modes of BISHOP . . . . .	92
4.2	Performance of different machine learning algorithms in optimising a 2D function . . . . .	95
4.3	Performance of Bayesian optimisation in optimising a 2D function under different model configurations . . . . .	111
5.1	Vulcan-PW laser parameter fluctuations during experiment A . .	136
5.2	Vulcan-PW laser parameter fluctuations during experiment B . .	138
5.3	Stability of RSIT-enhanced proton acceleration under fluctuation of Vulcan-PW laser parameters . . . . .	142

# List of Commonly Used Variables

Variable	Symbol	Units
Electric field	$E$	$\text{TVm}^{-1}$
Magnetic field	$B$	T
Laser pulse energy	$E_L$	J
Laser pulse duration	$\tau_L$	fs, ps, ns
Laser focal spot size	$\phi_L$	$\mu\text{m}$
Laser pulse intensity	$I_L$	$\text{Wcm}^{-2}$
Laser central wavelength	$\lambda_L$	nm
Laser frequency	$\omega_L$	$\text{rads}^{-1}$
Laser incident angle	$\theta_L$	$^\circ$
Normalised light amplitude	$a_0$	dimensionless
Plasma electron frequency	$\omega_p$	$\text{rads}^{-1}$
Plasma electron density	$n_e$	$\text{m}^{-3}$
Plasma critical density	$n_{\text{crit}}$	$\text{m}^{-3}$
Plasma electron temperature	$k_B T_e$	MeV
Plasma Debye length	$\lambda_D$	nm
Plasma density scale length	$L$	$\mu\text{m}$
Plasma collisional frequency	$\omega_c$	$\text{rads}^{-1}$
Fast electron current density	$\mathbf{j}_f$	$\text{Am}^{-2}$
Cold electron return current density	$\mathbf{j}_r$	$\text{Am}^{-2}$
Relativistic Lorentz factor	$\gamma$	dimensionless
Proton maximum energy	$\epsilon_{p\text{max}}$	MeV
Target foil thickness	$l$	nm, $\mu\text{m}$
Time	$t$	fs
Area	$S$	$\text{m}^2$

# List of Acronyms

<b>Acronym</b>	<b>Definition</b>
AI	Artificial intelligence
ASE	Amplified spontaneous emission
CPA	Chirped pulse amplification
CLF	Central laser facility
CSD3	Cambridge service for data-driven discovery
EI	Expected improvement
EM	Electromagnetic (wave)
FWHM	Full width half maximum
GP[R]	Gaussian process [regression]
HPC	High-performance computing
ML	Machine learning
OAP	Off-axis parabolic (mirror)
OP[CP]A	Optical parametric [chirped pulse] amplification
PI	Probability of improvement
PIC	Particle-in-cell (simulation)
QED	Quantum electrodynamics
RF	Radio frequency
RPA	Radiation pressure acceleration
RSIT	Relativistic self-induced transparency
TNSA	Target normal sheath acceleration
UCB	Upper confidence bound

# Chapter 1

## Introduction

This thesis investigates the stability and optimisation of beams of protons accelerated when short duration, high-power, laser pulses irradiate solid-density material. The motivation for doing so is presented in this introduction, alongside developments in technology and physical understanding that have enabled this research. This introduction provides context for the results presented in chapters 4 and 5, which further the development of laser-driven proton beams for potential applications in quantum electrodynamics (QED) research, nuclear fusion, radiobiology, and other important challenges.

### 1.1 Brief history of ion acceleration

Since the early 1900's, there has been interest in accelerating particles for fundamental physics research and societal applications. Machines constructed in the following decades accelerated ions to energies of tens to hundreds of keV [1], culminating in the first cyclotron, which accelerated protons to energies beyond 1 MeV, in 1932 [2]. Designs for a larger synchrocyclotron followed, and were implemented in 1950 [3]. This machine accelerated protons to energies exceeding 100 MeV, laying foundations for cyclotrons and synchrotrons to be constructed around the world [4]. This led to the discovery of new particles and elements [5–7], and improved understanding of particle physics and the creation of the universe [8]. As of writing in 2024, there are an estimated 30,000 particle accel-

erators in operation across the world, 1% of which are used for research, with the other 99% being utilised in societal and industrial applications such as cancer therapy, lithography, and medical and security imaging [8, 9]. Amongst the 1% used for research is the large hadron collider (LHC), a synchrotron that accelerates protons to TeV energies in a structure of 27 km circumference, constructed at a cost of several billion pounds, making it one of the largest and most expensive scientific instruments in history. This emphasises the importance of particle accelerators, but also their limitations. The size and cost of synchrotrons, cyclotrons, and linear accelerators mean they are not as widely available, and thus impactful, as they ideally would be. For example, there are currently only two conventional radio frequency (RF) based accelerators dedicated to proton therapy in the United Kingdom, treating a thousand or so cancer patients per year [10–12]. Ideally, these facilities would be available in hospitals across the country to meet patient demand. Making accelerators smaller and less expensive requires alternative approaches, the most promising of which utilise another of the 20th centuries most impactful technologies, the laser.

After their first demonstration in 1960 [13], lasers of intensity,  $I_L \sim 10^{10} \text{ Wcm}^{-2}$ , were immediately used to ionise solid-density material, creating plasma [14–19]. Thermal expulsion of this laser-induced plasma resulted in ion acceleration to 100 keV energies, comparable to those measured in early RF-based accelerators [20–23], but without the beam-like quality. Mode-locking improved laser pulse compression in 1964 [24], sparking a gradual increase in laser intensities over the next 20 years, as illustrated in figure 1.1. Ion energies increased accordingly, to a maximum of  $\sim 1 \text{ MeV}$  by 1985 [25]. Increases to laser intensities saturated around this time, before chirped pulse amplification (CPA) was demonstrated [26] to reduce pulse fluences below the damage threshold of solid-state optics, meaning intense pulses of shorter,  $\sim 1 \text{ ps}$ , duration could be produced, in contrast to  $\sim (30\text{--}30,000) \text{ ps}$  previously [25]. This facilitated laser pulses of intensity  $I_L \geq 10^{18} \text{ Wcm}^{-2}$ , that were demonstrated to accelerate ions in a beam-like profile with up to  $\sim 58 \text{ MeV}$  energy, in 2000 [27–29].

Various mechanisms were proposed to explain this ion acceleration in the

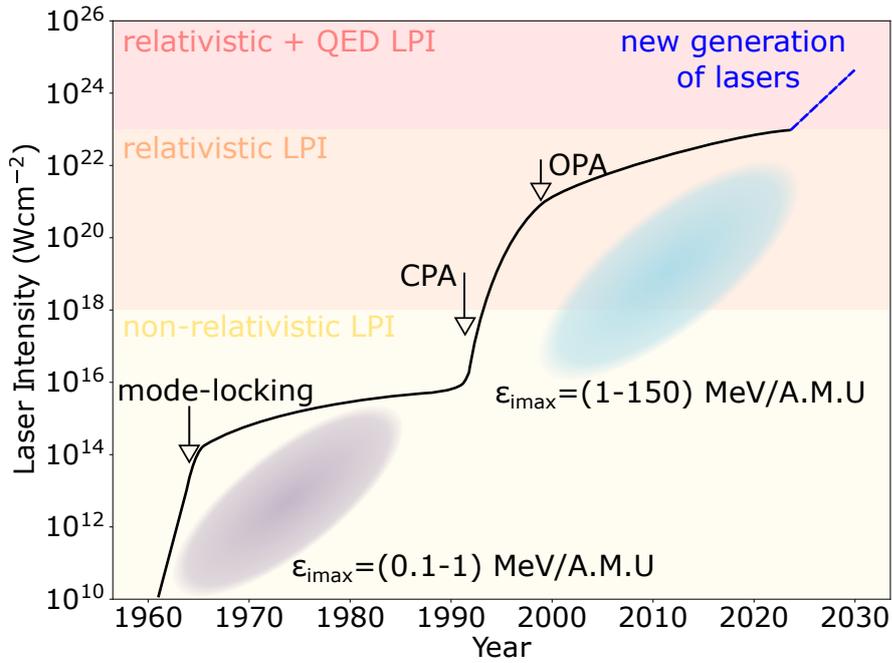


Figure 1.1: Illustration of increasing laser intensity with increased understanding and technological innovation over time. Initially non-relativistic laser-plasma interactions (LPI) and ion acceleration to 1 MeV/A.M.U was possible [25]. Now relativistic LPI and ion acceleration to 150 MeV/A.M.U is possible [30, 31].

following years, though they disagreed on whether the highest energy ions had originated from the target foil front or rear surface. Analysis of the results from these and further experiments [32] determined the highest energy protons had originated from hydrogen in a contaminant layer on the target rear, and that they were accelerated via the target normal sheath acceleration (TNSA) mechanism [33]. This now well-established mechanism begins with the laser inducing plasma on the front surface of a target foil, and accelerating plasma electrons to sufficient energy so that a large fraction propagate through the foil. A fraction of this population then escapes the foil, with the remainder building up near the rear surface. This induces charge separation with the positively charged ions, creating an electric sheath field which accelerates ions. The sheath field spans microns, resulting in an accelerating field strength of  $\sim \text{TVm}^{-1}$ , significantly higher than the  $\sim \text{MVm}^{-1}$  generated in RF-based accelerators [34, 35]. In principle, ions could therefore be accelerated to comparable energies by a smaller laser-driven system. Furthermore, this compact beam of ions has very low emittance, and is generated in a pulse of ultra-short duration, meaning the beam brightness is

orders of magnitude higher than that produced in an RF-based accelerator [36]. This provides access to the potentially beneficial FLASH regime of proton therapy [37], discussed in section 1.3 amongst other potential applications of laser-driven radiation sources which similarly make use of their unique properties.

Realising these applications requires laser-driven ion sources to be improved, so that these beneficial properties can be delivered stably, in a beam of improved energy and quality, with monoenergetic spectra, as delivered by RF-based accelerators. Such high quality beams are the result of over 100 years of RF-based accelerator research, and there is hope that continued development of laser-driven accelerators will enable them to stably deliver beams of ions with improved properties, building on the demonstration and development of this acceleration regime over the last 60 years. Various studies over these years demonstrated the energy, and other important properties, of TNSA ions to depend on the foil target thickness [38, 39], the laser energy, pulse duration [40, 41], focal spot size [42], and temporal-intensity contrast [43–45], the properties of beams of fast electrons generated and transported within plasma [46–50], and the accelerating field they induce [40, 51]. Laser properties were also improved by implementing adaptive optics [52] and plasma mirrors [53–55] whilst upgrading facilities [56–61], with these technological advances and new physical understandings combining to increase TNSA proton energies to a maximum of  $\sim 90$  MeV [62, 63].

Furthermore, this research contributed to experimental demonstrations [64–71] of alternative ion acceleration mechanisms driven by laser radiation pressure [72–74] and relativistic self-induced transparency (RSIT) of a target foil [75–78], the latter producing near  $\sim 100$  MeV proton energies [71]. New understanding of acceleration in this regime [45, 79] then contributed to even higher laser-driven proton energies of  $\sim 150$  MeV, reported in 2023 [31]. This was a significant step towards achieving the energy required for applications discussed in section 1.3, though to date RSIT enhancement of ion energies has not been demonstrated stably as laser and target conditions vary between interactions, as has been demonstrated for TNSA [37]. Instead, the energy of protons accelerated in this regime depends strongly on the time at which RSIT occurs relative to the peak laser in-

tensity irradiating the target, itself dependent on the foil target thickness [71, 79] and the laser temporal-intensity contrast [45]. Further instability in this regime is investigated in chapter 5, demonstrating that proton energies also vary due to relatively minor fluctuations in laser energy and pulse duration between interactions. These results also demonstrate that inducing RSIT earlier, by using thinner targets, results in proton acceleration to lower energies compared to inducing RSIT at the optimum time, but with less susceptibility to fluctuating laser parameters, and still higher proton energies than are possible from TNSA. Further investigation of RSIT-enhanced acceleration, only experimentally demonstrated in the last decade or so, is one promising route towards stably producing protons of sufficient energy so that other beneficial properties of these sources, such as ultra-short duration, compact size, and low emittance, can be utilised in fundamental physics research and societal applications.

Another pathway which has historically increased ion energies is to utilise lasers of increased intensity, as illustrated in figure 1.1. This is expected to continue as a new generation of laser facilities become available for experiments in the very near future [80–84], simultaneously unlocking a previously inaccessible regime of physics for investigation. The most intense lasers at these facilities are operated at  $\geq$ Hz repetition rates, much faster than past systems which delivered 1 pulse every 30 seconds, or even 30 minutes. This is enabling tens of thousands, rather than tens, of measurements to be made on experiments, enabling detailed statistical analysis of the most intense laser-driven ion acceleration for the first time. This presents huge opportunities but also challenges, in that methods currently employed to operate and diagnose experiments must be adapted to this increased operating rate. Work to address these challenges is underway in the form of new target designs [85–91], diagnostics [92–97], optics [98, 99], and software [100–102], each of which are being implemented on current  $\sim$ 1 Hz laser systems, to produce large data-sets that will soon be commonplace. Analysing these data-sets requires data-driven and machine learning (ML) based approaches [103, 104], that have already advanced research in various scientific fields.

## 1.2 Machine learning to advance laser-driven ion acceleration

ML is a subset of artificial intelligence (AI) models, first proposed in 1950 [105], and soon developed to play games such as checkers and chess [106], which have well defined rules and objectives for a machine to learn and achieve. Improvements to the machine models meant they soon became proficient at these games, eventually resulting in the Deep Blue supercomputer [107] defeating then chess world champion Garry Kasparov, in 1997. This was a landmark achievement for AI machines, though despite Deep Blue’s victory, “it wasn’t intelligent... at least not [as] founders of computer science had hoped.” [108]. Developments in AI somewhat stagnated in the next decade, before several major breakthroughs in the 2010’s, some of which came from Google DeepMind in their development of AlphaGo, a program which defeated the world champion of Go, a board game more complex than chess. This was another landmark moment for AI, during which AlphaGo played the now famous ‘move 37’, provoking laughter from commentators who “unanimously stated [that] not a single human player would have chosen move 37, [and that it was] a mistake” [109]. AlphaGo agreed, its model attributing a 1 in 10,000 probability that move 37 would have been chosen by a human. But move 37 it chose, in a pivotal point of a match it went on to win. Here, AlphaGo had gone “beyond its human guide, and come up with something new, something different” [109].

The promise of this achievement was that ML could contribute novel insight to more important challenges, as demonstrated when the Google DeepMind team applied similar models in AlphaFold, to solve “one of the most important unsolved issues of modern science”, in protein folding [110]. Before AlphaFold, researchers had spent decades identifying  $\sim 17\%$  of the 20,000 protein structures in the human body. AlphaFold correctly predicted 99% of these structures after only a few years of development, and has now predicted over 200 million proteins across various biological systems [111]. This tool is now open source, and

has accelerated research into drugs for treating cancer [112], Malaria [113], and Covid-19 [114], as well as research into reversing climate pollution [115]. Google DeepMind has also contributed to magnetic confinement fusion (MCF) research, accelerating progress towards clean nuclear energy as a replacement to fossil fuels [116]. Furthermore, ML was crucial in a major breakthrough towards this goal in 2022, when ignition, nuclear reactions releasing more energy than were used to compress the fuel and initiate the reactions, was achieved for the first time, in an inertial confinement fusion (ICF) reactor [117–119]. Here, ML was used to improve simulations of experiment designs, to propose novel designs [120–124], to enable stable high energy operation by automatically detecting optical damage [125], and to analyse results [126, 127], and improve diagnostics [128].

This work was made possible by open source algorithms which have democratised ML in the last decade, so that it is now being applied in various fields, including laser-plasma interactions [129]. ML-based image detection has identified anomalies in the laser beam profile during commissioning of the next-generation ELI-NP system [130], and detected shock waves, plasma waves, and plasma density distributions in laser-plasma interactions [131]. Genetic algorithms have proposed novel laser pulse profiles to improve electron beam properties in laser wake-field acceleration (LWFA) experiments [132], and novel target designs to optimise ion beam properties in simulated laser-solid interactions [133]. Bayesian optimisation in particular is being widely adopted, as it is more adept than other ML techniques in optimising sparse and noisy data-sets, which have historically been produced in laser-plasma studies due to the limited repetition rates of lasers, diagnostics, data capture, and target replenishment in experiments, and the limited number of computationally expensive simulations that can be performed. Bayesian optimisation has therefore been used to optimise electron and x-ray beam properties in a LWFA experiment [134], and TNSA ion beam properties in an experiment [135], and in simulations [100], the latter work being presented in detail in chapter 4. These are promising steps towards improving properties of radiation produced in laser-driven particle accelerators towards that required for fundamental physics and societal applications which will now be discussed.

## 1.3 Applications of radiation produced in laser-solid interactions

Laser-solid interactions produce a variety of radiation sources with properties of interest for applications [136, 137]. This includes THz radiation, x-rays, and  $\gamma$ -rays [138, 139], which can be used to analyse stress and strain in materials [140], to achieve nuclear fusion [118, 119, 141], and to induce strong-field QED effects when laser intensities reach  $I_L \geq 10^{23} \text{ Wcm}^{-2}$  at the newest generation of laser facilities. These interactions also support plasma optics to manipulate light, electron, and proton beams into high order modes [68–70], whilst still in a TEM<sub>00</sub> mode these electrons and protons are useful for a range of applications. In particular, this thesis investigates routes to stably optimise laser-driven proton beam properties of interest for nuclear fusion [142], radioisotope production for medical imaging [143, 144], radiography to detect electromagnetic field and plasma density perturbations in compressed matter [145], non-destructive diagnosis of cultural heritage artefacts [146], and stress testing of materials [147]. To understand the beneficial properties of laser-driven protons, and why they must be improved and stabilised, a further application is now described which would have huge societal impact if this was achieved.

### 1.3.1 Ultra-high dose rate proton therapy

Protons accelerated from laser-solid interactions form ultra-high dose pulsed beams of particular benefit to cancer therapy. Proton therapy was first proposed in 1946 [148] as an alternative to x-ray radiotherapy, since protons deposit energy in a very small region within a material, in contrast to much wider energy deposition from x-rays. The technique was first demonstrated in 1954 [149] but was not widely adopted, particularly due to poor cost effectiveness given large, expensive linear accelerators or cyclotrons were required to accelerate protons to sufficient energies for treatment. As discussed in section 1.1, this remains a problem to this day.

In principle, this problem could be solved using potentially more compact laser-driven proton accelerators, which have a further benefit in that they deliver a dose of  $\sim 10^9 \text{ Gys}^{-1}$ , much higher than the  $\sim \text{Gys}^{-1}$  delivered in conventional radiotherapy [37]. This ultra-high dose can trigger FLASH irradiation ( $\gtrsim 40 \text{ Gys}^{-1}$ ), reducing damage to healthy tissue surrounding a tumour [150–152] to enhance a patient’s quality of life after treatment [153, 154]. To this end, a stable,  $\sim 60 \text{ MeV}$  energy proton beam was recently deployed for in vivo studies of tumor irradiation in mice [37]. This is an extremely important development in the application of laser-driven proton beams for radiotherapy, though energies must be increased to  $\sim 250 \text{ MeV}$  to treat the deepest set tumours in human patients [155]. Several approaches have been proposed to reach this milestone, one of which makes use of laser-driven protons to achieve an ultra-high dose source at tens of MeV energy, which is then transported into an RF-based accelerator to achieve the hundreds of MeV required, rather than requiring a much larger linear accelerator to do so from scratch [156]. Another approach is to improve the maximum energy of laser-driven protons so that they can be utilised directly, which is the focus of this thesis.

## 1.4 Thesis outline

As discussed in the preceding sections, laser-driven particle accelerators are a relatively young technology compared to RF-based accelerators. Continued improvement of laser-driven radiation sources will enable their useful properties to be utilised in impactful fundamental physics and societal applications discussed in this chapter. Of particular interest to this thesis is the energy of laser-driven ion beams, which must be improved, and delivered stably over many laser shots, for applications of these beams to be realised. Contributions towards this goal are presented in this thesis, structured as follows:

- **Chapter 2** details the fundamental physics of laser-solid interactions, providing a theoretical base for the results to follow.

- **Chapter 3** introduces the methods used to achieve the results reported in later chapters. This includes methods used to amplify a laser pulse to intensities relevant for laser-solid interactions, to diagnose the radiation produced, and to improve and characterise laser spatio-temporal intensity properties. Methods used to simulate laser-solid interactions are also discussed.
- **Chapter 4** reports on an investigation into automating optimisation of laser-driven proton beams in simulations, using a new code called BISHOP. Optimisation is achieved more quickly, with less data-points, and for more interaction parameters, by employing Bayesian optimisation, rather than conventional techniques. This results in novel optimal interaction conditions, which are investigated.
- **Chapter 5** reports on an investigation into the stability of proton acceleration in the RSIT-enhanced, and TNSA regime, as a function of laser energy and pulse duration. Increased instability in the former regime is investigated, and is mitigated by inducing RSIT earlier, whilst still accelerating protons to energies beyond those achieved in TNSA.
- **Chapter 6** concludes the thesis by summarising key results, their significance to the field, and their promise in motivating future research.

# Chapter 2

## Fundamentals of laser-solid interactions

The work presented in this thesis focuses on improving properties of protons accelerated during the interaction between an intense laser pulse and a solid target foil. Expressed in one sentence this sounds like a relatively simple process. In reality it is quite complex, with many physical processes combining and competing throughout an interaction to determine how proton acceleration, and other radiation production, unfolds. This chapter will detail the most fundamental of these processes, starting with a brief overview of the properties of high-power laser pulses, and the plasma formed by these pulses, moving on to fast electron generation and transport, and ending with ion acceleration mechanisms.

### 2.1 Properties of high-power lasers

High-intensity laser pulses are fundamental in the physical processes described in this chapter. It is therefore useful to begin by characterising their properties. Typically, a laser pulse is considered as oscillating fields contained in an envelope with a temporally Gaussian distribution, defined by the peak amplitude and full width half maximum (FWHM). This is illustrated in figure 2.1(a) for an idealised temporal profile, with FWHM=900 fs and a peak electric field of  $|\mathbf{E}_0|=10^{14}$  Vm<sup>-1</sup>. Defining the FWHM as the pulse duration,  $\tau_L$ , and converting electric field to

energy ( $E_L \sim 207$  J), the peak power is  $P_L = E_L/\tau_L \sim 0.2$  PW for these parameters. This is typical of the laser energy, pulse duration, and peak power delivered by the Vulcan-PW laser used to accelerate protons to multi-MeV energies in chapter 5.

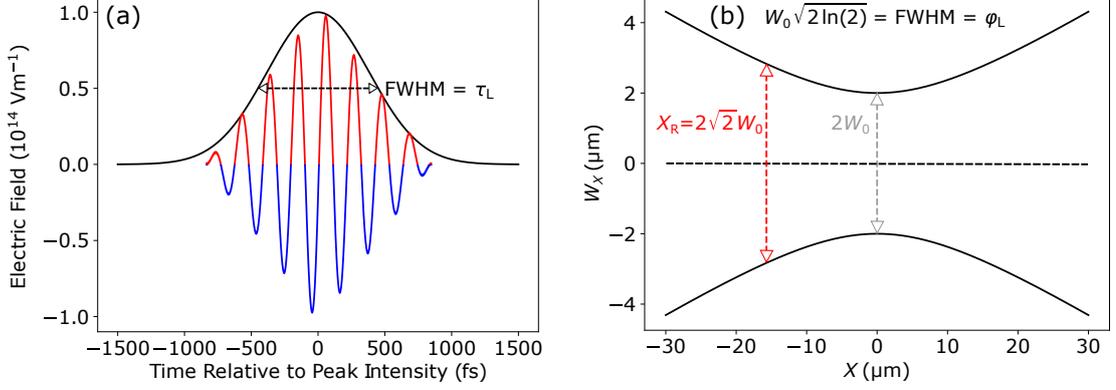


Figure 2.1: (a) Example of an idealised laser pulse consisting of an oscillating electric field (red/blue) enveloped by a Gaussian temporal profile (black). (b) Spatial evolution of a pulse focusing to a minimum spot size according to Gaussian optics.

The next important consideration is the spatial profile of the pulse. To reach the extreme intensities required to drive the physics investigated in this thesis, the initially collimated laser beam is focused to a spot with a Gaussian like spatial profile according to the following equation,

$$W_X = W_0 \sqrt{1 + \left(\frac{X}{X_R}\right)^2} \quad (2.1)$$

Here,  $W_0$  is the minimum radius to which the beam can be focused, typically characterised in terms of the beam waist,  $2W_0$  [157]. The beam expands along the propagation axis from this minimum area due to diffraction, increasing to a radius,  $W_X$ , over a distance,  $X$ , as characterised by the Rayleigh length,  $X_R = \pi W_0^2/\lambda_L$ , the distance over which the beam radius increases by a factor of  $\sqrt{2}$ , related to the laser central wavelength,  $\lambda_L$ . This process is illustrated in figure 2.1(b), where a beam focuses to a minimum waist of  $2W_0 = 4 \mu\text{m}$ . Converting to Gaussian FWHM according to the equation shown in figure 2.1(b), the focal spot size,  $\phi_L$ , is defined similarly to the pulse duration, and is  $\phi_L = 4.7 \mu\text{m}$  for this example. For a focal spot with an idealised Gaussian spatial profile, 50% of the laser energy is encircled within this FWHM. Multiplying the laser energy by this factor,  $E_{E\%}$ ,

and combining with the pulse duration and focal spot size, the peak laser intensity is defined as,

$$I_L = \frac{E_L \cdot E_{E\%}}{\tau_L} \frac{1}{\pi(\phi_L/2)^2} \quad (2.2)$$

In reality, the encircled energy contained within the FWHM is not 50% but typically 30%-40%, as explained in section 3.3. Using this approximate encircled energy fraction, the peak intensity of the example laser pulse considered thus far is calculated as  $I_L \sim 5 \times 10^{20} \text{ Wcm}^{-2}$ , which is typical of currently operational laser systems [82]. The peak intensity of the main pulse driving a laser-solid interaction is one of the most important parameters in defining the properties of radiation produced, as will be discussed in the results chapters of this thesis, though the main pulse is only part of an overall laser temporal-intensity profile which must be considered.

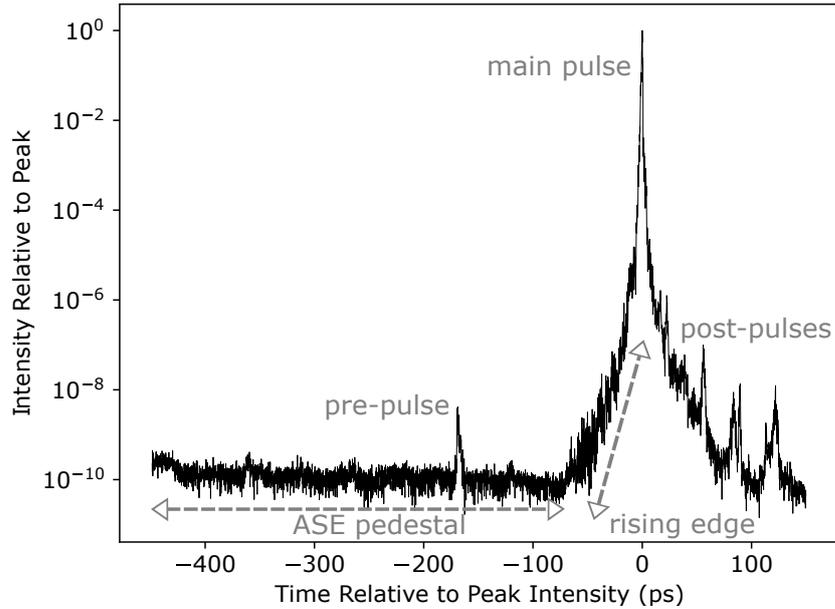


Figure 2.2: Temporal-intensity evolution of the Vulcan-PW laser pulse profile. The main pulse is preceded and succeeded by light from amplified spontaneous emission (ASE), and from other sources, described in chapter 3.

In practice, it is not only a main, high-intensity, laser pulse with an idealised Gaussian spatio-temporal profile that drives an interaction, but a main pulse with approximately these properties, that is temporally preceded, and succeeded, by lower intensity light. This is illustrated in figure 2.2, which shows an example

measurement of the temporal-intensity profile of the Vulcan-PW laser, made during the experiment investigated in chapter 5. The origins of these additional light sources differ, and will be discussed in more detail in section 3.1. For now, it is important to discuss the most influential preceding light, which typically comes from amplified spontaneous emission (ASE).

As the name suggests, ASE is a pedestal of spontaneous laser emission extending for hundreds of picoseconds, upon which the main pulse sits [58]. The ratio between the main pulse intensity and ASE intensity is defined as the temporal-intensity contrast and, though it is quite large for this example, corresponds to light of  $I_L \sim 10^{10} \text{ Wcm}^{-2}$ , for the Vulcan-PW laser pulse of  $I_L \sim 5 \times 10^{20} \text{ Wcm}^{-2}$  peak intensity, considered thus far. ASE light of this intensity is sufficient to pre-ionise the front surface of a target foil before the main pulse arrives, as is the more intense light contained within the shorter rising edge shown in figure 2.2. This pre-ionisation is therefore the first step in a laser-solid interaction, and has a significant influence on the main interaction which takes place when the peak intensity arrives [43, 44]. As such, the ionisation mechanisms induced by preceding light sources of various intensities are discussed in the following section.

## 2.2 Laser-induced ionisation processes

When discussing ionisation induced by high-power laser pulses, it is useful to begin by considering the most simple case, that is removal of an electron from the ground state of a hydrogen atom by an individual photon. For this to occur, the electric field strength which binds the electron to the hydrogen nucleus,  $\mathbf{E}_a$ , must be overcome. This is defined as,

$$\begin{aligned} \mathbf{E}_a &= \frac{e}{4\pi\epsilon_0 a_B^2} \\ &\sim 0.51 \times 10^{12} \text{ Vm}^{-1} \end{aligned} \tag{2.3}$$

where  $e$ ,  $\epsilon_0$  and  $a_B$  are the elementary charge, free-space permittivity, and Bohr

radius, respectively. The optical intensity required to overcome this electric field is calculated as,

$$I_a = \frac{c\epsilon_0\eta\mathbf{E}_a^2}{2} \quad (2.4)$$

$$\sim 3.5 \times 10^{16} \text{ Wcm}^{-2}$$

where  $c$  is the speed of light in a vacuum of refractive index  $\eta=1$ . Clearly, this is orders of magnitude higher than the intensity typically associated with an ASE pedestal,  $I_L \sim 10^{10} \text{ Wcm}^{-2}$ , despite this light being discussed as sufficient to ionise the front surface of a target foil. That is because ionisation can be induced either by a single photon of sufficient energy, as illustrated in figure 2.3(a), or by many lower energy photons.

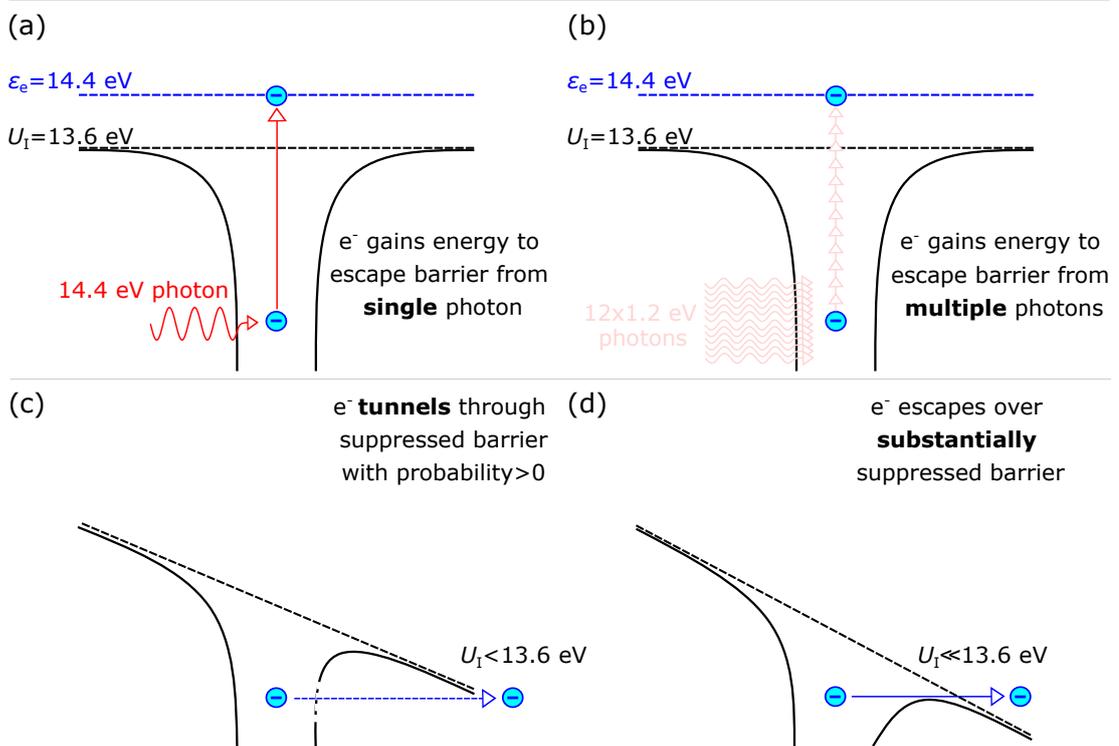


Figure 2.3: Illustration of (a) single photon, (b) multi-photon, (c) tunnelling, and (d) over-the-barrier ionisation.

### 2.2.1 Multi-photon ionisation

This process of so-called multi-photon ionisation (MPI) is illustrated in figure 2.3(b), where an electron gains sufficient energy to be freed from its parent atom if the incident light is intense enough so that there are many photons, making an  $n_{\text{th}}$  order ionisation event probable [158, 159]. Furthermore, an electron can absorb more photons than is necessary for multi-photon ionisation, in a process termed above-threshold ionisation [160, 161], resulting in an expelled electron gaining a final kinetic energy of,

$$E_e = (n + s)\hbar\omega_L - U_I \quad (2.5)$$

where  $n$  is the number of photons required to overcome the ionisation potential,  $U_I$ , via multi-photon ionisation,  $s$  is the excess number of photons absorbed, and  $\omega_L$  and  $\hbar$  are the photon angular frequency and the reduced Planck's constant, respectively. Light in the example ASE pedestal shown in figure 2.2 has a central wavelength of  $\lambda_L=1054$  nm, and so individual photons resulting from ASE have an energy of  $\sim 1.2$  eV, according to  $E=\hbar\omega_L/2\pi$ . The ionisation potential of hydrogen is  $U_I=13.6$  eV, meaning 12 of these photons will ionise a hydrogen atom through multi-photon ionisation if absorbed in  $\sim 0.5$  fs.

### 2.2.2 Barrier suppression ionisation

Also shown on figure 2.3(c-d) are two further ionisation mechanisms that occur not due to electrons gaining energy to overcome the ionisation potential, but rather due to the ionisation potential decreasing so that even relatively low energy electrons can escape their binding atoms. This requires a strong electric field, as associated with light of increased intensity,  $I_L \geq 10^{14}$  Wcm<sup>-2</sup>, in the rising edge of the pulse temporal-intensity profile, illustrated in figure 2.2. These fields are sufficient to distort the Coulomb potential,  $V(X)=-Ze^2/X$ , experienced by electrons at distance,  $X$ , from a nucleus of charge  $Z$ ,

$$V(X) = \frac{-Ze^2}{X} - e\mathbf{E}_0X \quad (2.6)$$

Electrons can then be freed from the reduced potential through two mechanisms, the first of which is known as ‘tunnelling ionisation’, and is illustrated in figure 2.3(c). Here, an electron does not have sufficient energy to overcome the potential barrier in a classical sense, but has some finite probability of tunnelling through the barrier in a quantum mechanical process. In the presence of a strong laser electric field, the barrier is suppressed according to the  $-e\mathbf{E}_0X$  term of equation 2.6, significantly increasing this probability, resulting in tunnelling ionisation. If the barrier is suppressed even further, electrons can escape from the severely reduced potential in a process known as ‘over-the-barrier’ ionisation, as illustrated in figure 2.3(d).

The laser intensity required to induce both of these barrier suppression ionisation (BSI) mechanisms can be derived from a parameter first proposed by Keldysh [162, 163],

$$\gamma_K \sim \sqrt{\frac{U_I}{\phi_{\text{pond}}}} \quad (2.7)$$

where  $\phi_{\text{pond}} = \frac{e^2\mathbf{E}_0^2}{4m_e\omega_L^2}$  is the ponderomotive potential, which characterises the energy acquired by an electron of mass,  $m_e$ , oscillating within a laser electric field, as described in detail in section 2.4.2. BSI becomes the dominant ionisation mechanism when  $\gamma_K \leq 1$  [162, 164], and so by making this substitution, and converting electric field to intensity using equation 2.4, the laser intensity at which these different ionisation mechanisms will dominate can be expressed as,

$$I_L \sim 2U_I \frac{m_e\omega_L^2\epsilon_0c}{\gamma_K^2e^2} \quad (2.8)$$

Again returning to the simplest case of ionising hydrogen ( $U_I=13.6 \text{ eV} \sim 2.2 \times 10^{-18} \text{ J}$ ), BSI becomes the dominant mechanism at an intensity of  $I_L \sim 1.4 \times 10^{14} \text{ Wcm}^{-2}$ , for a laser of central wavelength,  $\lambda_L=1054 \text{ nm}$ . In contrast, MPI is predominantly responsible for ionising hydrogen at intensities of  $I_L \sim 10^{10} \text{ Wcm}^{-2}$ , correspond-

ing to  $\gamma_K=100$ . As such, the Keldysh parameter is often employed to determine which ionisation mechanism will dominate at a given laser intensity, with  $\gamma_K \leq 1$  characterising high laser intensities and BSI, and  $\gamma_K \gg 1$  defining lower intensities which induce MPI.

Experiments have demonstrated excellent agreement with the theoretical description of BSI, not only for hydrogen, but also for a range of noble gases [165, 166], demonstrating that this is the dominant ionisation mechanism induced by a  $\lambda_L=1054$  nm,  $\tau_L \sim 1000$  fs, laser for intensities in the range  $I_L \sim (10^{13}-10^{17})$  Wcm<sup>-2</sup>. Importantly, a laser of this wavelength and approximate pulse duration was used in the experimental investigation presented in chapter 5, and so BSI is expected to be the most prevalent mechanism by which the front surface of target foils are ionised at the laser focus, in this study. Finally, collisional ionisation [167], where already freed electrons carry sufficient additional energy (according to equation 2.5) to ionise other atoms upon collision, becomes predominantly responsible for ionisation, rather than MPI or BSI, once these field ionisation mechanisms have already freed many electrons from their parent atoms [168]. Collisional ionisation is therefore important when electrons propagate into target foils, as discussed in section 2.5. Now that laser-induced ionisation of single atoms has been discussed, it is possible to describe plasma, formed when many atoms are ionised, and heated by an intense laser pulse.

## 2.3 Laser-induced plasma formation

As discussed in the previous section, photons contained within the ASE pedestal and rising edge of a high-power laser pulse are sufficiently energetic and numerous to free electrons from their atoms, forming plasma, which has unique properties compared to the other fundamental states of matter. The first of these unique properties occurs due to plasma containing charged particles which interact predominantly through the Coulomb force. This induces so-called collective behaviour, where one particle can influence many other particles in its vicinity, resulting in localised electric fields. This results in localised charged regions, in

contrast to the dynamics of a solid, liquid, or gas, where atoms interact individually through binary collisions, and charge neutrality is maintained over the entire material.

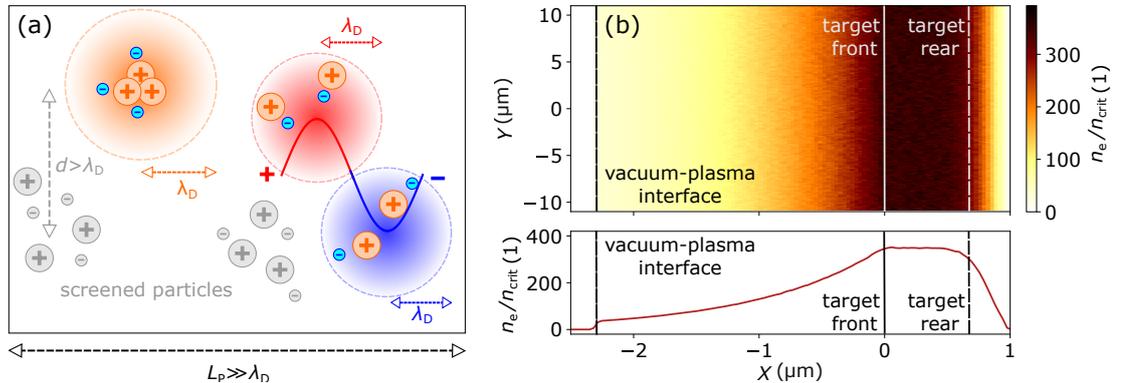


Figure 2.4: (a) Illustration of a plasma where electrons (cyan) arrange themselves to neutralise an immersed positive charge associated with a group of protons (orange), and the positive half cycle of an EM wave (red), screening particles at a distance greater than the Debye length,  $d > \lambda_D$ , from experiencing a force due to these charges (grey). This is also shown for protons arranging themselves to neutralise an immersed negative charge due to an EM wave (blue). (b) Illustration of a pre-plasma density profile induced on the front surface of a target foil.

Despite this, a plasma does maintain charge neutrality on a macroscopic scale, since there is an approximately equal amount of electrons and ions. One might expect this condition to be broken easily in the presence of an electric field, which would typically be considered to influence charged particles. However, on a macroscopic scale, this is not the case due to unique dynamics which occur within a plasma.

To understand this, imagine a charge immersed in a plasma of length,  $L_P$ . This is illustrated in figure 2.4(a), where the immersed charge can first be considered as that associated with a group of protons. Plasma electrons are collectively attracted to this charge through the Coulomb force, and effectively act to neutralise it so that charge neutrality is maintained in the overall plasma. More specifically, the electrons cause the Coulomb potential associated with the charge to attenuate exponentially as  $Qe^{-\lambda_D}$ , where  $\lambda_D$  is known as the Debye length and is related to the plasma electron density,  $n_e$ , temperature,  $T_e$ , and Boltzmann constant,  $k_B$ ,

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{e^2 n_e}} \quad (2.9)$$

Charged particles at a distance greater than the Debye length from the immersed charge are ‘screened’, feeling essentially no influence. To demonstrate this, the Coulomb potential experienced by a particle at distance,  $r$ , from a charge in a plasma is defined as,

$$\phi_r = \frac{Q}{4\pi\epsilon_0 r^2} \exp\left(\frac{-r}{\lambda_D}\right) \quad (2.10)$$

Plasma formed by low-intensity laser light preceding the main pulse varies in electron density as shown in figure 2.4(b), and in electron temperature as discussed in section 2.5, though for this example these values can be considered to be on the order of  $k_B T_e \sim 10$  keV, and  $n_e \sim 10^{28} \text{ m}^{-3}$ , on average. From equation 2.9 the plasma Debye length is then calculated as  $\lambda_D \sim 7$  nm. Substituting into equation 2.10, the Coulomb potential experienced by a particle at  $r \sim 100$  nm, outside the Debye sphere of an immersed charge is attenuated by  $\sim 10^{11}$ . The length scale of a typical plasma is much larger than the Debye length, on the order of microns, and so plasma is considered charge neutral on a macroscopic scale, assuming there are many particles within the Debye sphere, and the plasma frequency,  $\omega_P$ , is greater than the frequency of collisions,  $\omega_c$ .

The first of these assumptions is met in most plasma, since the electron density is typically on the order of  $n_e \sim 10^{28} \text{ m}^{-3}$ , as mentioned. The second assumption is crucial for a material to behave as a plasma, since it ensures that electron dynamics are governed by electric fields as described, rather than by hydrodynamic forces which determine electron behaviour in a material such as ionised gas, where charged particles frequently collide with neutral particles. Plasma frequency is related to electron density, and thus to electron temperature through equation 2.9,

$$\omega_P = \sqrt{\frac{n_e e^2}{\epsilon_0 m_e}} = \sqrt{\frac{k_B T_e}{m_e \lambda_D^2}} \quad (2.11)$$

The electron-ion collisional frequency has a weaker dependence on temperature [169],

$$\omega_c \propto n_e Z T_e^{-3/2} \quad (2.12)$$

and is therefore lower than the plasma frequency,  $\omega_c < \omega_P$ , for modest electron temperatures,  $k_B T_e \geq 1$  eV. Material ionised by the ASE pedestal and rising edge of the pulse temporal-intensity profile is generally orders of magnitude higher in temperature, and is therefore considered as plasma.

Now that the plasma state has been defined, the dynamics that occur when plasma is irradiated by an intense laser pulse can be discussed. Beyond creating plasma at the front surface of a target foil, light preceding a main pulse is also sufficiently intense so that it exerts significant thermal pressure on this plasma region, causing it to expand away from the target and into vacuum at approximately the ion acoustic velocity,

$$c_s = \sqrt{\frac{k_B(T_e + T_i)Z}{m_i}} \quad (2.13)$$

where  $k_B T_i$  defines the temperature of plasma ions, with mass,  $m_i$ . This is illustrated in figure 2.4(b), where the density of plasma electrons decreases exponentially with distance,  $X \leq 0$ , from a target foil of solid-density,  $n_{e0}$ , forming an expanding plasma according to,

$$n_e(X) = n_{e0} \exp(X/L) \quad (2.14)$$

where  $L$  is the scale length, which defines the distance over which the initial density decreases by a factor of  $1/e$ , where  $e$  is Euler's number. In this example, the plasma has a scale length of  $L=1$   $\mu\text{m}$ , and expands from a  $l=1$   $\mu\text{m}$  slab of solid-density CH plasma representing a CH target foil of this thickness. This is consistent with experimental measurements [170] and simulations [45] of the so-called pre-plasma density scale length, induced when intense light from sources such as the ASE pedestal, or rising edge of the temporal-intensity profile, irradiate a target foil before a main pulse of duration,  $\tau_L \sim 1$  ps, arrives. The scale length can also be estimated analytically from the speed of sound,  $c_s$ , in a material,

$$L = c_s \tau_L \quad (2.15)$$

and is thus smaller for shorter pulse durations, such as the  $\tau_L \sim (25-100)$  fs laser pulses used in the investigation presented in chapter 4, in which laser-plasma interaction dynamics, and resultant radiation production, is shown to change significantly when the pre-plasma density scale length is varied. This is attributed to changes in electron energy absorption and transport, which occurs through a range of mechanisms, described in section 2.4.3. Before getting to this discussion, it is important to first describe the dynamics of a single electron in response to an individual EM wave.

## 2.4 Response of electrons to intense laser light

In figure 2.4, pre-plasma induced by relatively low-intensity ASE light is illustrated. Due to their larger mass, ions within plasma are generally considered stationary on the timescale of the laser electric field cycle, for intensities (EM field strengths) covered in this thesis, whilst electrons do respond to the force associated with an EM wave according to the Lorentz equation,

$$\mathbf{F} = \frac{\partial \mathbf{p}}{\partial t} = -e(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \quad (2.16)$$

where  $\mathbf{E}$  and  $\mathbf{B}$  are the electric and magnetic field, and  $\mathbf{p} = \gamma m_e \mathbf{v}$  is the momentum acquired by an electron which is initially stationary, and therefore non-relativistic ( $\gamma = 1/\sqrt{1 - (\mathbf{v}^2/c^2)} = 1$ ). Electrons move throughout plasma according to this force, as will now be discussed.

### 2.4.1 Electron motion in response to an infinite planar electromagnetic wave

To begin, let us assume that the single photon from figure 2.4(a) is extended to be an infinite EM wave, the maximum amplitude of which never attenuates. The magnetic field,  $\mathbf{B}$ , is a factor  $\mathbf{v}/c$  lower than the electric field, and so, follow-

ing from equation 2.16, non-relativistic electrons ( $\mathbf{v} \ll c$ ) experience a negligible influence from the magnetic field. Instead they are accelerated primarily by the electric field, oscillating in the  $Y$  direction as  $\mathbf{E}_0 = \sin(\omega_L t - kY)$ , for a plane wave propagating along the  $X$  axis. Substituting this, and  $\mathbf{p} = m_e \mathbf{v}$  into equation 2.16, an initially motionless electron is accelerated upon interaction with a plane wave electric field according to,

$$m_e \frac{\partial \mathbf{v}}{\partial t} = -e \mathbf{E}_0 \sin(\omega_L t - kY) \quad (2.17)$$

Rearranging, and integrating with respect to time using the substitution  $u = \omega_L t$ , the so-called quiver velocity of an electron oscillating in the laser electric field is,

$$\mathbf{v}_{\text{osc}} = \frac{e \mathbf{E}_0}{m_e \omega_L} \cos(\omega_L t - kY) \quad (2.18)$$

reaching a maximum of  $\mathbf{v}_{\text{max}} = e \mathbf{E}_0 / m_e \omega_L$ , when the electric field peaks. If the electric field is sufficient that an electron gains energy greater than its rest mass energy,  $\epsilon_e \geq m_e c^2$  (0.511 MeV), in the first half wave cycle, then the quiver velocity approaches the speed of light, and the magnetic field term of equation 2.16 ceases to be negligible, instead accelerating electrons in the laser pulse propagation direction (perpendicular to both the electric field and magnetic field orientation). Returning to equation 2.16, making the substitution  $\mathbf{B} = \mathbf{v} / c \mathbf{E}$ , and integrating with respect to time, the complete equation for electron motion in an infinite planar EM wave can be expressed [54] as,

$$\mathbf{v}_{\text{tot}} = \frac{e^2 \mathbf{E}_0^2}{4m_e^2 \omega_L} \left[ \cos(\omega_L t - kY) + \frac{1}{2} \sin(2\omega_L t - kY) \right] \quad (2.19)$$

This is not an intuitive expression, but the important part to note is the oscillatory term,  $\sin(2\omega_L t - kY)$ , which indicates that electrons oscillate at twice the laser frequency in the propagation direction due to the magnetic field, in contrast to oscillating once per wavelength,  $\cos(\omega_L t - kY)$ , in the transverse direction due to the electric field, according to equation 2.18.

It is not practical to use these complicated equations regularly, and so a parameter known as the normalised light amplitude is often used to determine

whether plasma electrons are non-relativistic, only propagating in one dimension according to the laser electric field, or relativistic, oscillating in this way, but also propagating into higher density plasma regions, in a figure-of-eight like motion. This is important when considering how electrons absorb laser energy, as described in section 2.4.3, and so it is useful to define this parameter, as the ratio of the maximum quiver velocity (equation 2.18) to the speed of light,

$$a_0 = \frac{v_{\text{osc}}}{c} = \frac{e\mathbf{E}_0}{m_e\omega_L c} \quad (2.20)$$

Electrons are considered non-relativistic if they oscillate at a velocity much less than the speed of light,  $a_0 \ll 1$ , with energy less than their rest mass energy,  $\epsilon_e < 0.511$  MeV, and relativistic,  $a_0 \geq 1$ , if they oscillate at velocities approaching the speed of light, with energy,  $\epsilon_e \geq 0.511$  MeV. Relating the electric and magnetic field of an infinite plane wave to the time averaged Poynting vector, the optical intensity can be expressed through equation 2.4, for linear polarisation [54]. The intensity of light required to accelerate electrons to relativistic velocities can then be related to the normalised light amplitude by rearranging equation 2.20 into electric field terms and substituting,

$$I_L = \frac{a_0^2 m_e^2 4\pi^2 c^5 [\text{cm}] \epsilon_0 [\text{Fcm}^{-1}]}{\lambda_L^2 2e^2} \quad (2.21)$$

whilst expressing  $\epsilon_0$  and  $c$  in non-SI units so that the constant terms can be expressed as  $1.37 \times 10^{18} \text{ Wcm}^{-2}$ . The laser intensity required to accelerate electrons to a given  $a_0$  can then be expressed in these units,

$$I_L = \frac{a_0^2}{\lambda_L^2} 1.37 \times 10^{18} [\text{Wcm}^{-2}] \quad (2.22)$$

For the  $\lambda_L = 1054$  nm Vulcan-PW laser used in the results of chapter 5, electrons acquire relativistic velocities ( $a_0 = 1$ ) from light of intensity,  $I_L \sim 1 \times 10^{18} \text{ Wcm}^{-2}$ , tens of picoseconds before interacting with light at the peak intensity,  $I_L \sim 1 \times 10^{20} \text{ Wcm}^{-2}$ . Now that the motion of an individual electron interacting with an idealised, infinite plane wave is understood, the dynamics that occur when many plasma electrons interact with a realistic laser pulse can be discussed.

## 2.4.2 Electron motion in response to the laser ponderomotive force

In section 2.3, a realistic plasma density profile induced when laser light of intensity in the range  $I_L \sim (10^{10} - 10^{14}) \text{ Wcm}^{-2}$  interacts with a target foil was described. Light of this intensity precedes the main pulse which has an intensity on the order of  $I_L \sim (10^{18} - 10^{21}) \text{ Wcm}^{-2}$ , and is sufficient to induce relativistic effects for individual electrons, considering light of wavelength  $\lambda_L \sim 1000 \text{ nm}$ , as discussed in the context of an infinite plane wave with constant amplitude in the previous section. However, as illustrated in figure 2.1(b), a realistic laser pulse is not a plane wave propagating with constant amplitude, and instead focuses to a minimum area, and peak intensity, over some distance along the propagation ( $X$ ) axis. This is illustrated in figure 2.5, where the Gaussian propagation profile from figure 2.1(b) is overlaid onto the plasma profile from figure 2.4(b).

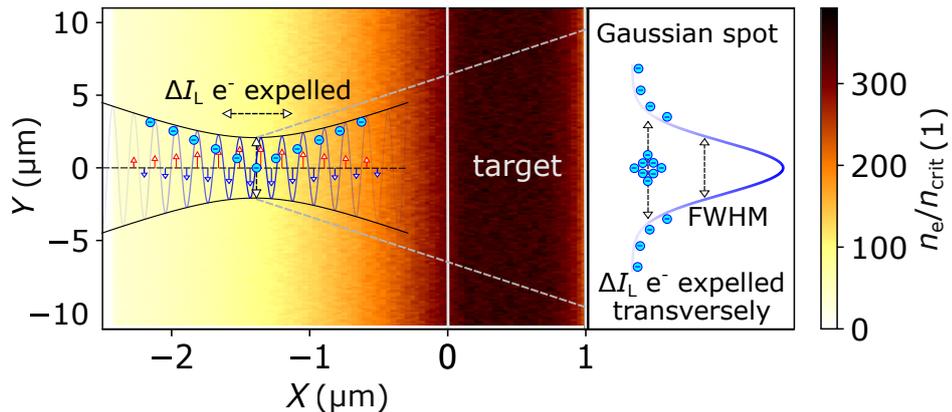


Figure 2.5: Illustration of an intense laser focusing in plasma, expelling electrons from high-intensity regions at the minimum beam size, to lower intensity regions, according to the ponderomotive force. This effect also occurs in the direction transverse to the laser pulse propagation, due to the Gaussian spatial-intensity profile of the focused laser.

In contrast to the case of an infinite plane wave with constant amplitude, the electric field amplitude decreases with distance from the focus position (herein referred to as best focus), as illustrated by a fading blue wave in figure 2.5. In the case of an infinite plane wave, an electron, initially at rest, is accelerated in one direction during the first half cycle of a single wavelength, gaining maximum velocity at the peak amplitude,  $\mathbf{E}_{\max}$ , according to equation 2.18, for non-relativistic

electrons. The electron is then accelerated in the opposite direction during the second half cycle with peak  $-\mathbf{E}_{\max}$ , effectively returning the electron velocity to zero at the end of the wave. However, if the electric field instead decreases in amplitude over some distance, an electron is accelerated to a peak velocity related to  $\mathbf{E}_{\max}$  on the first half cycle, but is decelerated according to a peak electric field of say  $0.9 \times \mathbf{E}_{\max}$ , thus retaining some velocity and gaining energy over one full laser cycle. Over many laser cycles this results in an electron gaining energy and moving from regions of high electric field to regions of lower electric field.

This is illustrated in figure 2.5, where two electrons move away from the highest intensity region at best focus, in one case towards the target foil, and in the other case away from it. This effect is even more extreme in the dimension parallel to the target, at best focus, since there is also a spatially Gaussian intensity profile in this dimension, which decreases over a smaller distance than in the propagation direction, on the order of microns, rather than tens of microns. Again, this expels electrons according to the intensity gradient, as illustrated in the right panel of figure 2.5. The force which governs this behaviour is derived [54, 168] from equation 2.16 as,

$$\mathbf{F}_{\text{pond}} = -\frac{e^2}{4m_e\omega_L^2}\nabla\mathbf{E}_0^2 = -\frac{e^2}{8\pi\epsilon_0m_e c^3}\nabla(I_L\lambda_L^2) \quad (2.23)$$

Known as the ponderomotive force, this determines the energy gained by non-relativistic electrons in the presence of an intense laser pulse. For relativistic electrons this can also be related to the normalised light amplitude defined in equation 2.20,

$$\mathbf{F}_{\text{pond}} = -m_e c^2 \nabla \gamma = -m_e c^2 \nabla \sqrt{1 + a_0^2/2} \quad (2.24)$$

The energy gained by an electron according to this force is then  $\epsilon_e = m_e c^2 (\gamma - 1)$ , which is  $\epsilon_e \sim 0.1$  MeV for a laser intensity, wavelength, and normalised light amplitude of  $I_L \sim 10^{18}$  Wcm<sup>-2</sup>,  $\lambda_L = 1054$  nm, and  $a_0 = 1$  respectively. Overall, the effect of the ponderomotive force over tens of microns in the propagation direction, and several microns in the transverse direction, around best focus,

results in distortion of the plasma density profile.

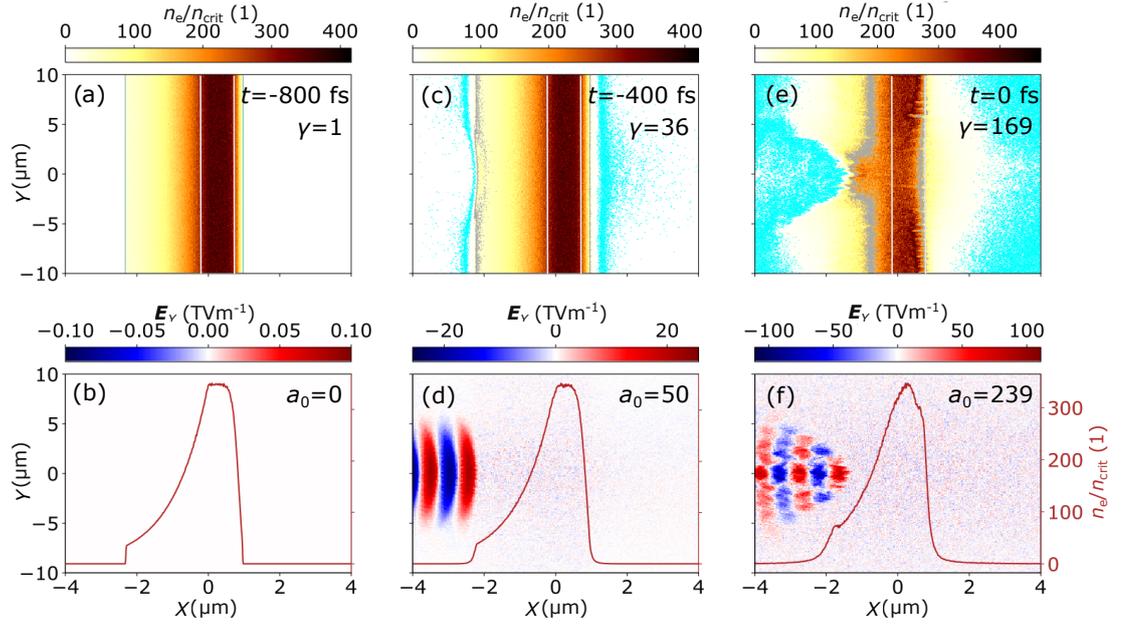


Figure 2.6: Results from a 2D PIC simulation in which an intense laser pulse interacts with pre-plasma on the front surface of a CH target. (a,c,e) shows the electron density as a colourmap, the gamma factor,  $\gamma$ , associated with the highest velocity electrons as text, and the classic,  $n_e=n_{\text{crit}}$ , and relativistically corrected,  $n_e=\gamma n_{\text{crit}}$  critical density surface, as cyan and grey contours. The colourmap of (b,d,f) shows the laser electric field in the polarisation ( $Y$ ) direction, with the associated peak  $a_0$  shown as text. The secondary (red) axis shows the average electron density in a region of  $Y=(-1 \rightarrow 1)$   $\mu\text{m}$ , around the laser axis.

To illustrate this, the interaction between a laser with the approximate parameters discussed so far,  $E_L=24$  J,  $\tau_L=400$  fs,  $\phi_L=5$   $\mu\text{m}$ ,  $I_L=3 \times 10^{20}$   $\text{Wcm}^{-2}$ , and a pre-plasma of scale length  $L=1$   $\mu\text{m}$ , in front of a CH target foil of  $l=1$   $\mu\text{m}$  thickness, and  $n_e=400n_{\text{crit}}$  electron density was simulated. This was achieved using EPOCH, which is a fully relativistic particle-in-cell (PIC) code [171], described in section 3.5, and used to simulate laser-target interactions over a range of conditions in chapters 4 and 5.

The initial pre-plasma density profile, shown in figure 2.6(a-b), at time,  $t=-800$  fs, before the main pulse arrives at  $t=0$  fs, is modulated due to the ponderomotive force associated with the rising edge of the laser temporal-intensity profile, as shown at time,  $t=-400$  fs, in figure 2.6(c-d). The laser pushes electrons towards the target in the  $X$  dimension, and away from the focal spot region in the  $Y$  dimension, and will propagate within the pre-plasma according to the dispersion relation,

$$\omega_L^2 = c^2 k^2 + \omega_P^2 \quad (2.25)$$

where  $k$  is the laser wave vector. The phase velocity at which peaks and troughs of an EM wave will travel in a plasma can then be derived from this equation as,

$$\mathbf{v}_{\text{phase}} = \frac{\omega}{k} = \sqrt{c^2 + \omega_P^2/k^2} \quad (2.26)$$

Meaning light with frequency greater than the plasma frequency,  $\omega_L > \omega_P$ , has a phase velocity greater than the speed of light. This does not violate the laws of relativity, which state that information cannot travel faster than the speed of light, given peaks and troughs of the wave do not carry information. Rather, information is carried by the pulse envelope, which travels at the group velocity,

$$\mathbf{v}_g = \frac{\partial \omega}{\partial k} = \frac{c^2}{\sqrt{c^2 + \omega_P^2/k^2}} \quad (2.27)$$

and is always less than the speed of light, even if  $\omega_L > \omega_P$ . The phase velocity at which light travels in a medium is characterised by the refractive index,  $\eta = c/\mathbf{v}_{\text{phase}}$ . Equating with the first relation in equation 2.26,  $\mathbf{v}_{\text{phase}} = \frac{c}{\eta} = \frac{\omega}{k}$ , and substituting for  $\omega_L$  in equation 2.25, the refractive index of a plasma is derived as,

$$\eta = \frac{c}{\mathbf{v}_{\text{phase}}} = \sqrt{1 - \left(\frac{\omega_P}{\omega_L}\right)^2} \quad (2.28)$$

If the laser frequency is less than the plasma frequency,  $\omega_L < \omega_P$ , the refractive index becomes imaginary, and laser light cannot propagate in the plasma. Relating the plasma frequency to electron density through equation 2.11, laser light is therefore shown to propagate within a plasma up to some ‘critical’ density, where  $\omega_L = \omega_P$ ,

$$n_{\text{crit}} = \frac{\omega_P^2 \epsilon_0 m_e}{e^2} \quad (2.29)$$

On figure 2.6(a-b), the so-called critical density surface, up to which the laser pulse can propagate, is illustrated by a cyan contour, and is initially located less

than  $X=-2 \mu\text{m}$  from the target foil front surface. As the main pulse begins to interact with the pre-plasma at  $t=-400 \text{ fs}$ , electrons are expelled from the high-intensity focal spot region according to the ponderomotive force. This pushes the critical density surface closer to the target foil front surface, allowing the laser pulse to propagate within a region of so-called underdense plasma,  $n_e < n_{\text{crit}}$ .

Looking closely at figure 2.6(c-d), the laser pulse actually propagates slightly beyond the critical density surface, up to a secondary surface marked by a grey contour. This is explained by electrons gaining energy from the laser pulse through the ponderomotive force, and through other mechanisms described in section 2.4.3. When the laser intensity is such that the normalised light amplitude,  $a_0$ , is greater than 1, electrons increase in mass by the relativistic Lorentz factor, and the critical density of equation 2.29 must then be considered as the relativistically corrected critical density,

$$\gamma n_{\text{crit}} = \frac{\gamma \omega_{\text{P}}^2 \epsilon_0 m_e}{e^2} \quad (2.30)$$

This effect becomes more extreme as the laser intensity rises towards its peak at time,  $t=0 \text{ fs}$ , in figure 2.6(e-f), at which point the laser pulse is seen to propagate well beyond the classical critical density. The laser fields then attenuate exponentially beyond the relativistically corrected critical density surface, over a distance known as the plasma skin depth. This is related to the plasma frequency, and hence density, by  $l_s=c/\omega_{\text{P}}$ , and defines the distance over which the laser field is attenuated to  $1/e$  of its peak value.

One final point of interest from figure 2.6 is the population of electrons in vacuum beyond the target rear surface. The laser pulse does not propagate through the target foil in this simulation, and so these electrons are not accelerated directly by the ponderomotive force as described so far. Instead, these electrons propagate through the target foil after gaining energy from the laser at the target front surface, through a variety of mechanisms which will now be discussed.

### 2.4.3 Laser-to-electron energy absorption

At the intensities relevant to the work presented in this thesis, the laser will transfer energy to electrons predominantly through collisionless mechanisms which will be discussed in this section. Before doing so however, it is important to touch on collisional energy transfer which occurs at relatively low-intensities, as found in light which temporally precedes, and exists in spatial wings of, a high-intensity pulse. At such intensities, the plasma temperature can be sufficiently cold that the collisional frequency of electrons,  $\omega_c \propto T_e^{-3/2}$  from equation 2.12, is higher than the plasma frequency,  $\omega_P \propto T_e^{1/2}$  from equation 2.11, meaning the effect of electron-ion collisions must be considered. In this regime, an electron which has gained some energy during the first half of an EM wave cycle through the ponderomotive force can transfer this energy to an ion with which it collides. The electron then loses coherence with the EM wave, and cannot return the energy in the second half of the cycle. As such, the electron absorbs energy from a laser photon in what is considered an inverse Bremsstrahlung mechanism [172, 173], since it is the opposite of Bremsstrahlung, or braking radiation, where an electron collides with an ion, slows down, and loses energy which is radiated as a photon [174, 175]. Despite inverse Bremsstrahlung contributing little to laser-to-electron energy absorption within the pre-plasma compared to collisionless processes, at the intensities relevant to this thesis, Bremsstrahlung will become extremely relevant when these energetic electrons propagate through the pre-plasma into a target foil, as will be discussed in section 2.5.

With that said, at the intensities relevant to this thesis, electrons predominantly absorb energy from a laser pulse through three collisionless mechanisms, which occur to different extents depending on the laser intensity, polarisation, incident angle, and pre-plasma scale length.

#### Resonance absorption

The first of these mechanisms is known as resonance absorption [72, 176, 177], and occurs when an intense, linearly polarised laser pulse propagates into a pre-

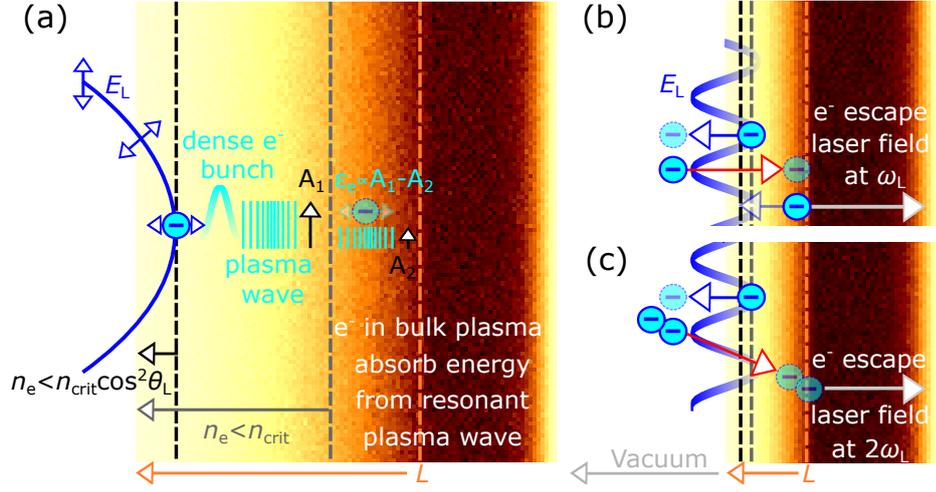


Figure 2.7: Illustration of laser-to-electron energy absorption in plasma of scale length (a) greater than the laser wavelength,  $L > \lambda_L$ , via resonance absorption and on the order of the laser wavelength,  $L \lesssim \lambda_L$ , via (b) vacuum absorption at laser intensities,  $I_L < 10^{18} \text{ Wcm}^{-2}$ , and (c)  $\mathbf{j} \times \mathbf{B}$  heating at higher intensities,  $I_L \geq 10^{18} \text{ Wcm}^{-2}$ .

plasma of relatively long scale length,  $L > \lambda_L$ , at an oblique angle of incidence. This causes light on one side of the pulse to experience a higher density, and thus a higher refractive index, than the other, according to the plasma density gradient in the propagation direction ( $X$ ) and equation 2.28. This effectively causes the pulse to ‘turn’, meaning it doesn’t propagate as far in underdense plasma as it would at normal incidence, instead reflecting from an angle adjusted critical density surface,  $n_{\text{crit}} \cos^2 \theta_L$ . As the laser pulse has turned in the plasma, electric field oscillations which were initially perpendicular to the plasma density gradient are now parallel to this gradient at the turning point, as illustrated in figure 2.7(a). These field oscillations drive electrons beyond the critical density surface, and into the plasma skin depth, at which point the field strength exponentially decays, as described in the previous section. Electrons are then accelerated across the critical density surface and into the plasma skin depth by a relatively strong field during the first half cycle of a laser wavelength, and are accelerated back across the surface in the opposite direction by a weaker field on the second half cycle.

As a result, electrons gain energy over a laser cycle, in a process similar to that induced by the ponderomotive force, though in this case due to a density, rather than an intensity gradient. Electrons can gain sufficient energy through this

process to travel beyond the plasma skin depth, at which point they experience no restoring force. Electrons therefore pile up in this region, inducing local charge imbalance with the stationary ions, resulting in a longitudinal electric field which acts on other electrons within the plasma. This behaviour can be considered as a longitudinal plasma wave propagating into higher density plasma with frequency equal to the driving laser frequency, as illustrated in figure 2.7(a). This electron bunching and subsequent wave propagation occurs over one laser cycle if the laser is very intense, or builds over many laser cycles at lower intensities. As the plasma wave propagates into denser plasma the leading edge experiences higher densities than the trailing edge, causing the trailing edge to overtake the leading edge. This causes the wave to break [178], damping the amplitude, which results in energy transfer to electrons within this region, as illustrated in figure 2.7(a).

The fraction of laser energy absorbed by plasma electrons through this mechanism has been modelled as a function of laser incidence angle [179], showing that optimal absorption occurs for oblique incidence angles [168, 178]. Experimental measurements replicate these models for large pre-plasma scale lengths, and for laser intensities in the range,  $I_L \sim (10^{12} - 10^{17}) \text{ Wcm}^{-2}$  [54], though they diverge at very small scale lengths,  $L < \lambda_L$ , and at higher intensities,  $I_L \geq 10^{18} \text{ Wcm}^{-2}$ , where electrons acquire relativistic velocities [140]. For these conditions the following mechanisms are predominantly responsible for laser-to-electron energy absorption.

### **Vacuum heating**

In plasma with a scale length shorter than the laser wavelength,  $L < \lambda_L$ , electrons absorb laser energy through a similar but distinct mechanism known as vacuum heating [180, 181]. This is illustrated in figure 2.7(b), where electrons initially located within the plasma skin depth are accelerated across the critical density surface, in this case into vacuum, since the density scale length is now smaller than the laser wavelength, and hence the distance over which electrons oscillate. The electrons then experience a much larger field which accelerates them back across the critical density surface and beyond the skin depth, at which point

they experience no restoring force on the next laser cycle, and can propagate with some velocity towards the target foil. This process can only occur over one laser cycle, in contrast to resonance absorption which can build over many cycles, and becomes the dominant laser-to-electron energy absorption mechanism at intensities,  $I_L \sim (10^{16}-10^{18}) \text{ Wcm}^{-2}$  [180].

### Relativistic $\mathbf{j} \times \mathbf{B}$ heating

The final mechanism to discuss drives electrons to relativistic velocities, and is therefore extremely relevant at intensities,  $I_L > 10^{18} \text{ Wcm}^{-2}$ , as found in the rising edge temporal-intensity profile of Vulcan-PW laser light at  $\lambda_L = 1054 \text{ nm}$ , shown in figure 2.2, and used for the results presented in chapter 5. At such intensities, the force experienced by electrons due to the laser magnetic field,  $\mathbf{v} \times \mathbf{B}$ , is no longer negligible, and is comparable to that experienced by electrons due to the electric field, as discussed in section 2.4.1. Electrons then acquire a velocity from both field components according to equation 2.19.

For a linearly polarised laser with the electric field oscillating perpendicular to the direction of the pre-plasma density gradient, electrons also oscillate perpendicular to this gradient due to the electric field, but are driven forward, across the density gradient by the  $\mathbf{j} \times \mathbf{B}$  force associated with the magnetic field. This is shown in figure 2.7(c), where electrons are driven across the critical density surface, and then experience a significantly reduced field, meaning they retain the transferred laser energy. This absorption mechanism is thus known as  $\mathbf{j} \times \mathbf{B}$  heating [182], and is distinct from the previous two mechanisms in that it drives electrons across the surface twice per laser cycle, due to the  $2\omega_L$  term of equation 2.19. This is illustrated by two electrons crossing the critical density surface in figure 2.7(b), rather than one electron for the other absorption mechanisms. Other distinctions are that  $\mathbf{j} \times \mathbf{B}$  heating is optimal for normal,  $\theta_L = 0$ , rather than oblique,  $\theta_L > 0$ , incidence, and that electrons are driven in the laser pulse propagation direction by the  $\mathbf{j} \times \mathbf{B}$  force [183, 184], rather than normal to the target, as for the other mechanisms. In addition, though each mechanism is most efficient when the laser is linearly polarised,  $\mathbf{j} \times \mathbf{B}$  heating will only occur for this

polarisation state, since it does not occur at all if the polarisation is circular, due to the magnetic field component of equation 2.19 averaging to zero in this case. This is highly relevant to the discussion of laser-driven ion acceleration in section 2.6. Having discussed the mechanisms by which plasma electrons gain energy from the incident laser and propagate into overdense plasma, it is useful to now consider their behaviour as they continue to propagate into a target foil.

## 2.5 Electron transport in overdense plasma

Up until this point, the discussion has focused predominantly on the energy gained by single electrons within an underdense plasma in response to laser EM fields. Those with sufficient energy were described to uncouple from the laser upon crossing the critical density surface, at which point they propagate in overdense plasma of increasing density, and into a target foil.

Before describing this propagation, it is first necessary to divert from considering single electrons, and to instead consider a population of many electrons which vary in energy, and number, according to intensity and density gradients within the pre-plasma. This results in an approximately Maxwellian distribution [185],

$$f(\epsilon_f) = N_f \sqrt{\frac{4\epsilon_f}{\pi(k_B T_f)^3}} \exp\left(-\frac{\epsilon_f}{k_B T_f}\right) \quad (2.31)$$

which characterises the number of ‘fast’ electrons,  $N_f$ , that is electrons which have acquired relativistic velocities from the laser, and therefore have energies greater than their rest mass,  $\epsilon_f \geq m_e c^2$  (0.511 MeV). If many electrons in the population have acquired these relativistic velocities, the distribution may be better represented by a two temperature Maxwell-Juttner distribution, rather than a single temperature Maxwellian distribution [186, 187]. This is the case for an example fast electron spectrum shown in figure 2.8, taken from a simulation of a laser-solid interaction presented in chapter 4. Here, two straight lines are fit over the range,  $\epsilon_e = (0.511-3)$  MeV, and  $\epsilon_e = (3-30)$  MeV, and the average fast electron temperature in these regions is calculated from the inverse gradient, as

$k_B T_f \sim 1$  MeV, and  $k_B T_f \sim 8$  MeV, respectively. This is an extremely important parameter to quantify, as it directly influences ion acceleration, the optimisation of which is the focus of this thesis.

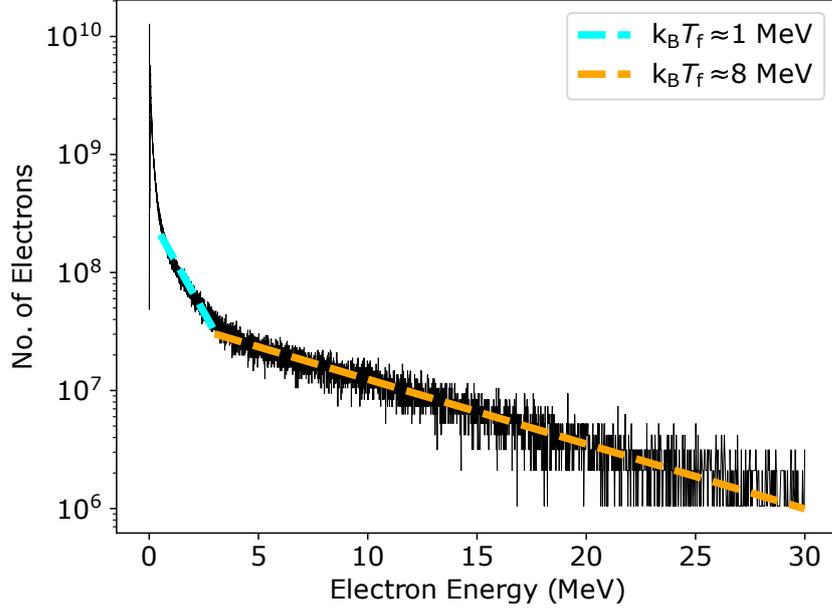


Figure 2.8: Two-temperature electron energy spectrum from a 2D PIC simulation.

As discussed, the energy acquired by electrons is directly related to the laser intensity, and the associated normalised light amplitude. The fast electron temperature has therefore been shown to scale with these parameters [72] as,

$$k_B T_f = m_e c^2 \left[ \left( 1 + \frac{a_0^2}{2} \right)^{1/2} - 1 \right] \quad (2.32)$$

for electrons which acquire relativistic velocities from a linearly polarised laser, incident at  $\theta_L = 0^\circ$ , with respect to target normal. Electrons are continually heated in the presence of circularly polarised light, rather than once per laser cycle for linearly polarised light, removing the denominator in the  $a_0^2$  term,

$$k_B T_f = m_e c^2 \left[ \left( 1 + a_0^2 \right)^{1/2} - 1 \right] \quad (2.33)$$

Laser-solid interactions are highly complex, and the fast electron temperature depends on many variables, often diverging from this scaling. Alternative temperature scalings have therefore been proposed [42, 188–190], with each agreeing the temperature scales to some degree with laser intensity, as is demonstrated for

simulations with a range of pulse durations in chapter 4.

With the properties of this fast electron population now described, the propagation of such electrons within a target foil can be discussed. To do so, it is useful to continue referring to the example electron population presented in figure 2.8, the bulk of which acquired an average temperature of  $k_B T_f \sim 8$  MeV, from a laser with energy,  $E_L = 15$  J. Using a simplified version of equation 2.31, the number of electrons in this population can be estimated as,  $N_f \sim A_{\text{frac}} E_L / k_B T_f \sim 10^{12}$ , assuming the fraction of laser energy absorbed by plasma electrons is  $A_{\text{frac}} = 30\%$ , though this parameter varies in the range (20-90)%, depending on parameters such as the laser incidence angle, intensity, focal spot spatial distribution, and target thickness [50, 183, 191–193]. These electrons are generated in a timescale on the order of the laser pulse duration, which for this example is  $\tau_L = 25$  fs. As such, the fast electron beam carries a current of  $I_f \sim (N_f e) / \tau_L \sim 6$  MA. Considering the electron beam as a cylinder with radius,  $r_L$ , and height,  $c\tau_L$ , the beam density,  $n_f$ , can be calculated as the number of electrons within the cylinder volume,  $V = \pi r_L^2 c\tau_L$  [194],

$$n_f = \frac{N_f}{\pi r_L^2 c\tau_L} \quad (2.34)$$

Using the estimated number of fast electrons,  $N_f \sim 10^{12}$ , and assuming they have a velocity of  $\mathbf{v}_f \sim c$ , the associated current density is calculated as  $\mathbf{j}_f = -en_f \mathbf{v}_f \sim 5 \times 10^{17}$  Am<sup>-2</sup>, and will induce a magnetic field according to [195],

$$\frac{\partial \mathbf{B}}{\partial t} = \frac{\rho \mathbf{j}_f}{r_L} \quad (2.35)$$

where  $\rho$  is the electrical resistivity of the target. Early studies demonstrated that this self-generated magnetic field will inhibit and reverse the electron beam motion [196], with electrons only continuing to propagate if they carry a current below the Alfvén [197] limit,

$$I_A \sim \frac{\beta \gamma m_e c^2}{e} = 17000 \beta \gamma \quad (2.36)$$

calculated as  $I_A \sim 0.1$  MA when  $\beta = v/c \sim 0.99$ , as is assumed for electrons acceler-

ated to relativistic velocities,  $\gamma \sim 6$ , by the  $I_L \sim 8 \times 10^{21} \text{ Wcm}^{-2}$ ,  $a_0 \sim 8$ , laser pulse considered in this example. A current of 6 MA exceeds this, and so fast electrons carrying this current should not propagate within a target foil. But they do, seemingly defying this limiting condition. This can be explained by considering the electric field which is also generated due to this fast electron beam, approximated in 2D as,

$$\frac{\partial \mathbf{E}}{\partial t} = -\frac{\mathbf{j}_f}{\epsilon_0} \quad (2.37)$$

The example electron beam of current density,  $\mathbf{j}_f \sim 5 \times 10^{17} \text{ Am}^{-2}$ , will induce an electric field of  $|\mathbf{E}| \sim 10^{15} \text{ Vm}^{-1}$ , on the order of the  $|\mathbf{E}| \sim 10^{14} \text{ Vm}^{-1}$  and  $|\mathbf{E}| \sim 10^{15} \text{ Vm}^{-1}$  required to ionise hydrogen and carbon, respectively [165, 166]. This means the electric field associated with the fast electron population can ionise atoms within the CH target foil used in the simulation illustrated in figure 2.6, and in the results chapters of this thesis. This creates a much colder population of ‘bulk’ electrons, in the same region as the fast electron beam, as illustrated in figure 2.9.

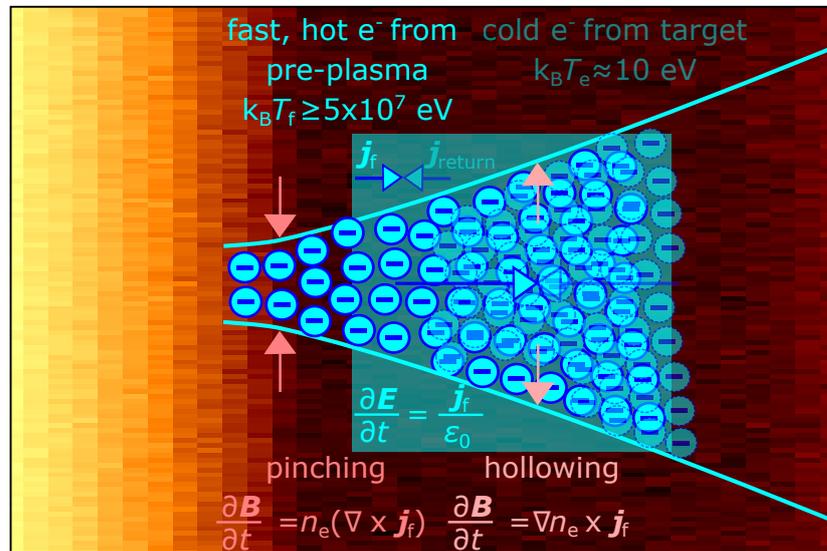


Figure 2.9: Illustration of a fast electron beam (opaque cyan) propagating into a target foil, and drawing a return current of cold electrons from the target bulk (transparent cyan), due to the electric field associated with the beam current. Magnetic fields are also associated with this current, and act to attract electrons towards, and push them away from, the central propagation axis, causing pinching and hollowing of the beam, respectively.

This induces an electric field, and thus ‘return’ current of density,  $\mathbf{j}_r$ , according

to equation 2.37, neutralising the forward going  $\mathbf{j}_f$  current, to a value around, or below the Alfvén limit. This allows fast electrons to propagate within a target foil. To explain this, the electric field that draws the return current can be related to the electrical resistivity of the target,  $\mathbf{E}=\rho\mathbf{j}_r$ . Using Faraday’s law, the associated magnetic field is then derived as,

$$\begin{aligned}\frac{\partial \mathbf{B}}{\partial t} &= -\nabla \times \mathbf{E} = -\nabla \times (\rho\mathbf{j}_r) \\ \nabla \times \mathbf{B} &= \mu_0(\mathbf{j}_f + \mathbf{j}_r)\end{aligned}\tag{2.38}$$

and is negligible, thus not inhibiting fast electron propagation, when the total current density,  $\mathbf{j}_f + \mathbf{j}_r$ , is very small. Imbalances in current density or resistivity do occur, inducing magnetic fields that are sufficient to affect fast electron beam propagation, in two major ways. The first effect results from the current density being higher along the centre of the beam and decreasing radially outward. This is a result of the Gaussian spatial-intensity profile of the laser focal spot, and causes ‘pinching’ of the beam, illustrated in figure 2.9, according to the first term on the RHS of the expanded equation 2.38, stated without higher order advection and diffusion terms for simplicity,

$$\begin{aligned}\frac{\partial \mathbf{B}}{\partial t} &= \nabla \times (\rho\mathbf{j}_f) \\ \frac{\partial \mathbf{B}}{\partial t} &= \rho(\nabla \times \mathbf{j}_f) + \nabla\rho \times \mathbf{j}_f\end{aligned}\tag{2.39}$$

This spatial gradient,  $\rho(\nabla \times \mathbf{j}_f)$ , in the current density is most prominent at the point where electrons are injected into the target foil, and can act to self-collimate the beam of fast electrons, opposing its inherent divergence [194, 195] which results from the focusing nature of the laser pulse, and is illustrated by light blue lines which diverge from the propagation axis at an angle of  $\theta_D \sim 20^\circ$ , in figure 2.9. Pinching also occurs to a lesser extent over smaller density variations throughout the target, acting to breakup the beam of fast electrons. This creates

filaments with more severe differences in density, in turn inducing further breakup of the beam of fast electrons, in a feedback process. Typically these structures are undesirable, and so filamentation is often considered as an instability [198–201].

The second important effect to consider when current neutrality is not achieved is that induced by resistivity gradients, described by the second term on the expanded RHS of equation 2.39,  $\nabla\rho \times \mathbf{j}_f$ . Resistivity is related to temperature, which varies spatially in the background plasma due to collisional effects associated with the cold return current. This induces a magnetic field which pushes electrons outwards from the beam centre, ‘hollowing’ the beam [202, 203], and effectively competing with the pinching effect resulting from density gradients. This effect can reverse in very hot plasma, approaching the Spitzer regime, so that the magnetic field draws electrons towards this central axis, complementing the pinching effect. In most cases the currents associated with these density and resistivity gradients are lower than the Alfvén limit, meaning fast electrons can propagate uninhibited to the target rear.

## 2.6 Laser-driven ion acceleration

The previous section discussed several mechanisms by which laser energy is transferred to plasma electrons, influencing their velocity, mass, and density, and causing them to propagate within plasma in various ways. Most importantly, the energy absorbed by electrons can subsequently be transferred to ions, which typically reach tens of MeV energies per nucleon, and have unique properties compared to ions accelerated by conventional radio frequency (RF) based accelerators. Among these unique properties are that ions are produced in highly laminar beams of small emittance, typically  $<4 \times 10^{-4}$  mm.mrad [32, 204],  $\sim 100 \times$  smaller than produced in RF-based accelerators [54], in a short duration. The exact properties depend on a variety of laser and target parameters which can be tailored to accelerate ions through various mechanisms, making laser-driven ion acceleration extremely flexible.

The investigations of this thesis focus on optimising and stabilising properties

of protons accelerated when short duration laser pulses interact with target foils. Target normal sheath acceleration (TNSA) was the first mechanism to produce tens of MeV protons from such an interaction, and is the most widely investigated mechanism in the field. The maximum energy of protons accelerated by TNSA is optimised in chapter 4, and the stability of TNSA proton energies is investigated in chapter 5. This mechanism is now described.

### 2.6.1 Target normal sheath acceleration

Results from the first experiments to demonstrate TNSA were published in the year 2000 [27–29], showing that protons were accelerated to maximum energies of up to  $\epsilon_{p\max} \sim 58$  MeV. Subsequent experiments varied laser and target parameters to better understand and optimise the mechanism [39, 41–44, 47, 49, 59, 205–207], with this, and improvements to laser technology [56, 57] and target manufacturing [89], resulting in TNSA maximum proton energies of up to  $\epsilon_{p\max} \sim 90$  MeV [62, 63].

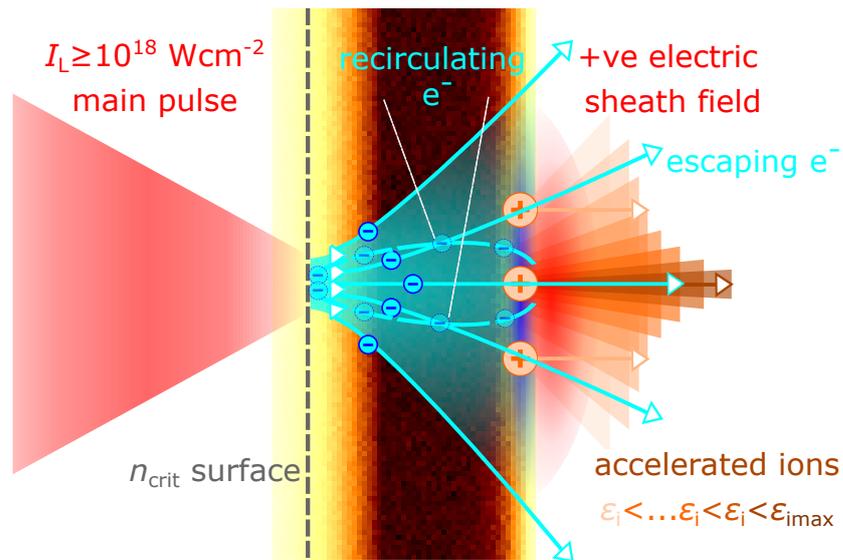


Figure 2.10: Illustration of target normal sheath acceleration.

Soon after the initial experiments in which it was demonstrated, TNSA was described to depend on much of the physics discussed in this chapter so far [33]. Initially, the front surface of a target foil is ionised by light preceding a  $I_L \geq 10^{18} \text{ Wcm}^{-2}$  main pulse, creating a pre-plasma, as described in sections 2.2 and 2.3. Electrons within this plasma then absorb energy from the main pulse, gaining

relativistic velocities at which they propagate into and through the target, as described in sections 2.4.3 and 2.5. The most energetic electrons can then escape the target, into vacuum, causing the target rear surface to become positively charged, as illustrated in figure 2.10. This results in charge separation between the escaping electrons and the now positively charged target rear, inducing a so-called electric sheath field which extends beyond the target rear, over a region of approximately the Debye length of electrons contained in the sheath, which can be calculated from equation 2.9 as  $\lambda_D=75$  nm for a population of sheath electrons with density,  $n_e\sim 10^{28}$  m<sup>-3</sup>, and temperature,  $k_B T_e\sim 1$  MeV. The sheath electric field strength can then be estimated from these parameters as  $|\mathbf{E}_{\text{sheath}}|\sim 10$  TVm<sup>-1</sup>, according to [33],

$$\mathbf{E}_{\text{sheath}} \sim \frac{k_B T_e}{e \lambda_D} \quad (2.40)$$

As discussed in section 2.2, this field strength is sufficient to ionise atoms in contaminant layers on the target rear surface [208]. These ions are then accelerated parallel to the target normal axis by the sheath field, with protons being accelerated more readily due to their higher charge to mass ratio,  $q/m_p$ , relative to heavier ions such as carbon. This process is demonstrated in figure 2.10, where ions are accelerated to the highest energies by the peak of the sheath field, reaching a maximum energy dependent on properties of the sheath [51, 209], and the driving laser pulse [40],

$$\epsilon_{\text{imax}} = 2Zk_B T_e [\ln(t_i + (t_i^2 + 1)^{1/2})]^2 \quad (2.41)$$

Here,  $k_B T_e$  is the temperature of electrons creating the sheath, often considered to be proportional to the square root of laser intensity,  $k_B T_e \propto I_L^{1/2}$  [72], as described in section 2.5. The maximum energy of ions accelerated by lasers of intensity,  $I_L \geq 10^{18}$  Wcm<sup>-2</sup>, via TNSA, is therefore also commonly expressed with this scaling,  $\epsilon_{\text{imax}} \propto k_B T_e \propto I_L^{1/2}$ . The term,  $t_i = \omega_i t_{\text{acc}} / \sqrt{2e}$ , in equation 2.41 is the time over which ions are accelerated, related to the ion plasma frequency,  $\omega_i$ , and a finite acceleration time,  $t_{\text{acc}}$ . The ion plasma frequency,  $\omega_i = \sqrt{n_e e^2 / \epsilon_0 m_e}$ , depends on

the sheath electron density, which can be estimated,  $n_e = N_f / c\tau_L S_{\text{sheath}}$ , from the number of fast electrons,  $N_f = E_L \times A_{\text{frac}} / k_B T_f$ , that absorb some fraction,  $A_{\text{frac}}$ , of the laser energy at the target front surface, in a time related to the pulse duration,  $\tau_L$ , and then propagate through the target foil of thickness,  $l$ , in a beam with a divergence angle,  $\theta_D$ , before escaping into vacuum, the majority then being confined to the sheath of transverse size,  $S_{\text{sheath}} = \pi(\phi_L + l \tan \theta_D)^2$ .

This model [40, 51] assumes sheath electrons to have a constant temperature throughout the interaction. However, the number, energy, and therefore temperature of fast electrons is not constant, and instead evolves over the course of an interaction. Without accounting for this, the model would overestimate the maximum ion energy compared to simulations and experimental results, but by defining a finite acceleration time,  $t_{\text{acc}} = \alpha(\tau_L + t_{\text{min}})$ , where  $\alpha = 1.3$  for  $I_L > 3 \times 10^{19} \text{ Wcm}^{-2}$ , and  $t_{\text{min}} = 60 \text{ fs}$  is the minimum time required for electrons to transfer energy to ions, the model can provide reasonable predictions of experimental and simulation results for intensities in the range,  $I_L = (1-6) \times 10^{19} \text{ Wcm}^{-2}$ , for pulse durations in the range  $\tau_L = (300-850) \text{ fs}$ , and for target foils of  $l \sim 20 \mu\text{m}$  thickness [40]. The model is less representative for different interaction conditions, and so modifications are typically made [41, 42, 210–213], and distinct models proposed [30], to accurately represent maximum ion energies over different, still relatively restricted parameter spaces. Machine learning techniques have recently been employed to model TNSA over a wider, more complete parameter space, accurately replicating experimental results from many different laser systems [214, 215].

Even the most accurate models only account for some, but not all, of the evolving TNSA dynamics in an interaction, not accounting particularly for ions accelerating into the sheath, decreasing charge separation and thus the accelerating field. Computational simulations provide a more complete account of evolving dynamics, as discussed in section 3.5, and so results from a 2D PIC simulation similar to that which illustrated laser-induced plasma and fast electron generation in figure 2.6 are now shown in figure 2.11. This expands on the simplified illustration of TNSA presented in figure 2.10, showing how these processes, and the following electron propagation discussed in section 2.5, evolve throughout a

laser-solid interaction, contributing to an evolving TNSA process. Figure 2.11(a-d) shows properties of the interaction at time,  $t=0$  fs, when the peak intensity,  $I_L=5\times 10^{20}$  Wcm $^{-2}$ , of a laser pulse with properties,  $E_L=1.5$  J,  $\tau_L=40$  fs, and  $\phi_L=3$   $\mu$ m, interacts with pre-plasma of scale length,  $L=300$  nm, at the front surface of a  $l=1$   $\mu$ m CH target foil.

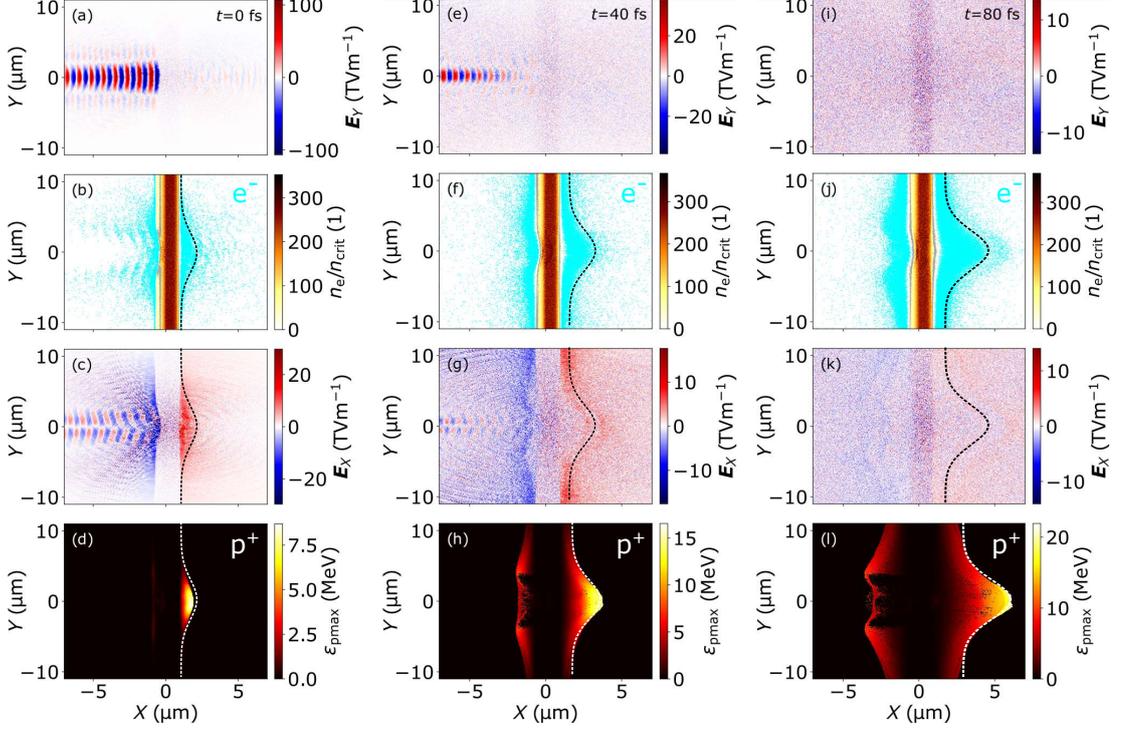


Figure 2.11: Illustration of TNSA in a 2D PIC simulation. (a) shows the laser electric field in the polarisation direction ( $Y$ ), at time,  $t=0$  fs, when the peak intensity interacts with (b) plasma of varying electron density (colourmap). Electrons at the classical (cyan),  $n_e=n_{\text{crit}}$ , and relativistically corrected (grey),  $n_e=\gamma n_{\text{crit}}$ , critical density are highlighted, showing electrons are confined in a sheath with a Gaussian like transverse distribution (black dashed line) at the target rear, resulting in a (c) longitudinal electric field which (d) accelerates protons. The same electric fields, particle densities, and energies are shown as the interaction continues at lower intensity in (e-h), and after the laser has diminished in (i-l).

As the laser intensity has risen to this peak over time,  $t=(-120\rightarrow 0)$  fs, electrons have been accelerated to relativistic velocities by the laser, and propagated through the target in a diverging beam, forming an electric sheath field with a Gaussian distribution in the transverse direction, which accelerates protons to the highest energies along the target normal axis. As the falling edge of the laser intensity profile continues to interact with the target in figure 2.11(e-h), fast electrons are continually generated at the front surface, and escape into vacuum to maintain the sheath field. In addition, electrons which were not initially

energetic enough to escape the target rear, and were reflected back to the front surface, as illustrated in figure 2.10, can absorb more energy from the laser as it continues to interact [50], meaning they now do have sufficient energy to escape the target rear and contribute to the sheath field. This so-called electron re-circulation results in transverse broadening ( $Y$  dimension), and replenishment, of sheath electrons, maintaining the field strength for longer, thus increasing the time over which protons are accelerated, and their resultant energy [38, 216]. Figure 2.11(i-l) shows that protons are continually accelerated for  $t \sim 40$  fs after the laser pulse diminishes, reaching a maximum energy of  $\epsilon_{p\max} \sim 20$  MeV in this example simulation.

The highest energy protons accelerated along the target normal axis can be preferentially selected whilst filtering out lower energy protons, producing a beam-like source of multi-MeV energies, with high laminarity, and an emittance on the order of  $10^{-4}$  mm.mrad, lower than produced in conventional RF-based accelerators [32]. In addition, this laser-driven proton beam is extremely short in duration at its source, being accelerated on the order of the laser pulse duration [40], and is compact in size, being generated over a transverse region of a few microns. A relatively high number of protons reach the highest energies,  $\sim 10^9$ , and more,  $\sim 10^{12}$ , are accelerated to lower, but still multi-MeV energies, from weaker regions of the sheath field [54]. Generating energetic protons on this ultra-short timescale, in this compact source size, is potentially useful for applications such as FLASH proton therapy [37]. Other applications require heavier ions, which can be accelerated from the target bulk when using different target materials such as aluminium, particularly if the hydrocarbon layer is removed by methods such as Ohmic heating [217], irradiation by an ion gun [218], or ablation with a preceding laser pulse of nanosecond duration [219].

Having now discussed the most studied laser-driven ion acceleration mechanism, the following sections will describe alternative mechanisms which can potentially accelerate ion beams with improved source properties compared to those produced via TNSA.

## 2.6.2 Radiation pressure acceleration

The next ion acceleration mechanism to discuss is that driven predominantly by the laser radiation pressure, which has been proposed to accelerate ions to maximum energies with a linear dependence on laser intensity,  $\epsilon_{\text{pmax}} \propto I_L$ , rather than with a weaker square root scaling,  $\epsilon_{\text{pmax}} \propto I_L^{1/2}$ , in TNSA [206]. Theoretical and simulation based studies have therefore predicted that radiation pressure acceleration (RPA) will produce higher maximum proton energies than can be achieved through TNSA, and will become the dominant ion acceleration mechanism in experiments using a new generation of lasers capable of delivering  $I_L \geq 10^{23} \text{ Wcm}^{-2}$  intensities [73]. At present, there are much fewer experimental demonstrations of RPA compared to TNSA, due to challenges which will be discussed in this section, though RPA has produced maximum proton energies of up to  $\epsilon_{\text{pmax}} \sim 93 \text{ MeV}$  in experiments [220], and has contributed to maximum proton energies of up to  $\epsilon_{\text{pmax}} \sim 96 \text{ MeV}$  [71], and  $\epsilon_{\text{pmax}} \sim 150 \text{ MeV}$  [31], in hybrid acceleration mechanisms enhanced by the onset of relativistic self-induced transparency (RSIT), the repeatability of which are investigated in chapter 5.

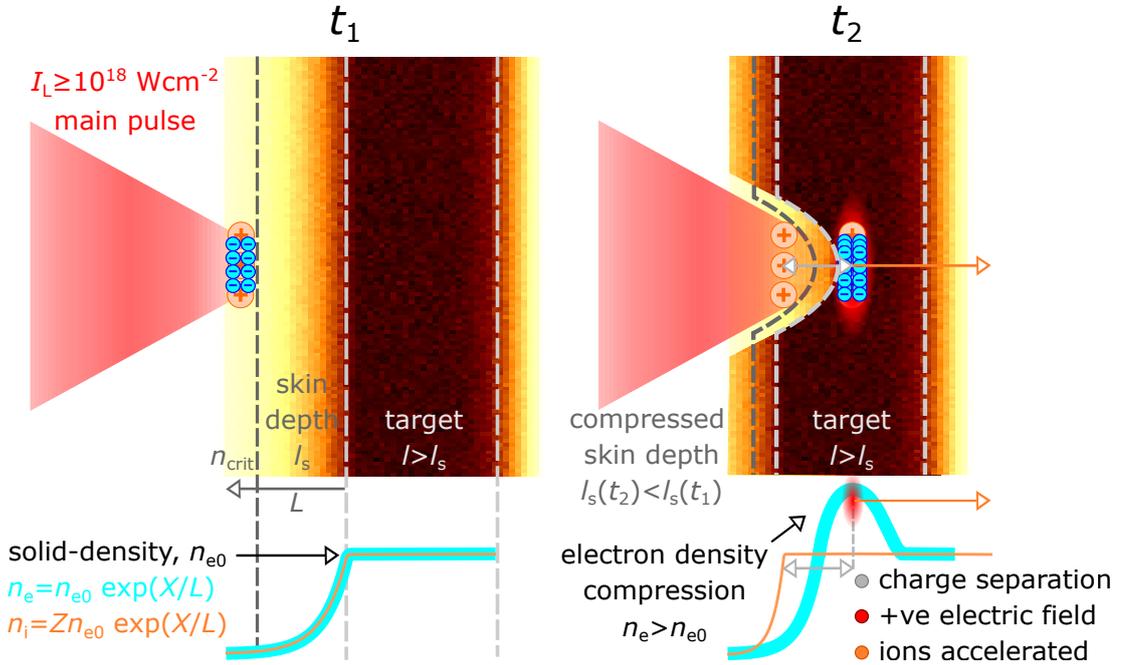


Figure 2.12: Illustration of radiation pressure acceleration occurring from time,  $t_1$ , when electron and ion densities are balanced in a pre-plasma, to time,  $t_2$ , when electrons separate from ions in the pre-plasma, creating a charge separation field which accelerates ions.

Similar to TNSA, protons are accelerated in RPA by a charge separation field, which in this case moves with the leading edge of the laser pulse, rather than being induced at the target rear-vacuum interface. This process begins with the laser radiation pressure, equivalent to the integral of the ponderomotive force over the pulse envelope, pushing electrons beyond the critical density surface, and into the plasma skin depth, causing them to accumulate and create a very dense, negatively charged region. If the radiation pressure is greater than the opposing thermal pressure of the expanding pre-plasma, the critical density surface will be pushed into the target foil [206], causing hole-boring of the front surface, as illustrated in figure 2.12. Ions in the region from which electrons have been ejected will remain there, since they experience negligible force due to their higher charge to mass ratio relative to electrons. This causes charge separation, creating a strong electric field that accelerates ions in the compressed, high density electron region, as illustrated in figure 2.12.

To estimate the maximum energy of ions accelerated in this so-called hole-boring regime (RPA-HB), it is useful to consider them as being contained within a mirror, the front surface of which is the critical density surface. Some fraction of the laser intensity is then reflected by,  $R$ , and transmitted through,  $T$ , the mirror, which moves forward due to the laser radiation pressure,

$$P_{\text{rad}} = (1 + R - T) \frac{I_{\text{L}}}{c} \quad (2.42)$$

As discussed in section 2.4.3, some fraction of the laser energy is also absorbed by plasma electrons at the critical density surface,  $A_{\text{frac}}$ , meaning the fraction of transmitted energy is  $T=1-R-A_{\text{frac}}$ . Rearranging, and substituting into equation 2.42, the laser radiation pressure can be expressed as,

$$P_{\text{rad}} = (2R + A_{\text{frac}}) \frac{I_{\text{L}}}{c} \quad (2.43)$$

and is maximised if no laser energy is absorbed by plasma electrons. Setting this condition for simplicity,  $A_{\text{frac}}=0$ , and equating the forces exerted due to the radiation pressure,  $\mathbf{F}_{\text{rad}}=P_{\text{rad}}S$ , and momentum,  $\mathbf{F}_{\text{rad}}=\partial\mathbf{p}/\partial t$ , of a moving

mirror of area,  $S$ , the velocity of the critical density surface, or the hole-boring velocity,  $\mathbf{v}_{\text{HB}}$ , can be expressed in acceleration terms [74],

$$\frac{\partial m \mathbf{v}_{\text{HB}}}{\partial t} = P_{\text{rad}} S \quad (2.44)$$

Equating the mirrors mass to density,  $m = \rho V$ , where the volume is related to the distance travelled by an initially stationary mirror over some time,  $V = S l_{\text{M}} = S \mathbf{v}_{\text{HB}} t$ , where  $l_{\text{M}}$  is the mirror length, the hole boring velocity is expressed as,

$$\mathbf{v}_{\text{HB}} = \sqrt{\frac{2I_{\text{L}}}{\rho c}} \quad (2.45)$$

Ions within the moving mirror then travel at this velocity, acquiring a maximum kinetic energy according to the peak laser intensity,

$$\epsilon_{\text{imax}} = \frac{m_{\text{i}}}{2} (2\mathbf{v}_{\text{HB}})^2 = \frac{m_{\text{i}} I_{\text{L}}}{\rho c} \quad (2.46)$$

Expressing the mass density as,  $\rho = A m_{\text{i}} n_{\text{i}} = A / Z m_{\text{i}} n_{\text{e}}$ , ions acquire a maximum energy per nucleon of,

$$\epsilon_{\text{imax}} = \frac{Z 2I_{\text{L}}}{A n_{\text{e}} c} \quad (2.47)$$

Theoretically, the maximum ion energy achieved in RPA is therefore linearly proportional to the laser intensity,  $\epsilon_{\text{imax}} \propto I_{\text{L}}$ , rather than having a square root dependence,  $\epsilon_{\text{imax}} \propto I_{\text{L}}^{1/2}$ , as for TNSA.

Achieving a perfectly reflecting critical density surface is difficult in practice, due to the numerous mechanisms by which electrons absorb (20-90)% of the incident laser energy at this surface [50, 183, 191–193], as described in section 2.4.3. This will diminish the RPA mechanism. However, this section also described that laser-to-electron energy absorption typically occurs through the  $\mathbf{j} \times \mathbf{B}$  mechanism at intensities relevant to these interactions,  $I_{\text{L}} \geq 10^{18} \text{ Wcm}^{-2}$ , and that the  $\mathbf{j} \times \mathbf{B}$  heating term vanishes when considering circularly polarised laser pulses. As such, circularly polarised light mitigates laser-to-electron energy absorption, and has been used to experimentally demonstrate RPA proton energies of up to  $\epsilon_{\text{pmax}} \sim 93$

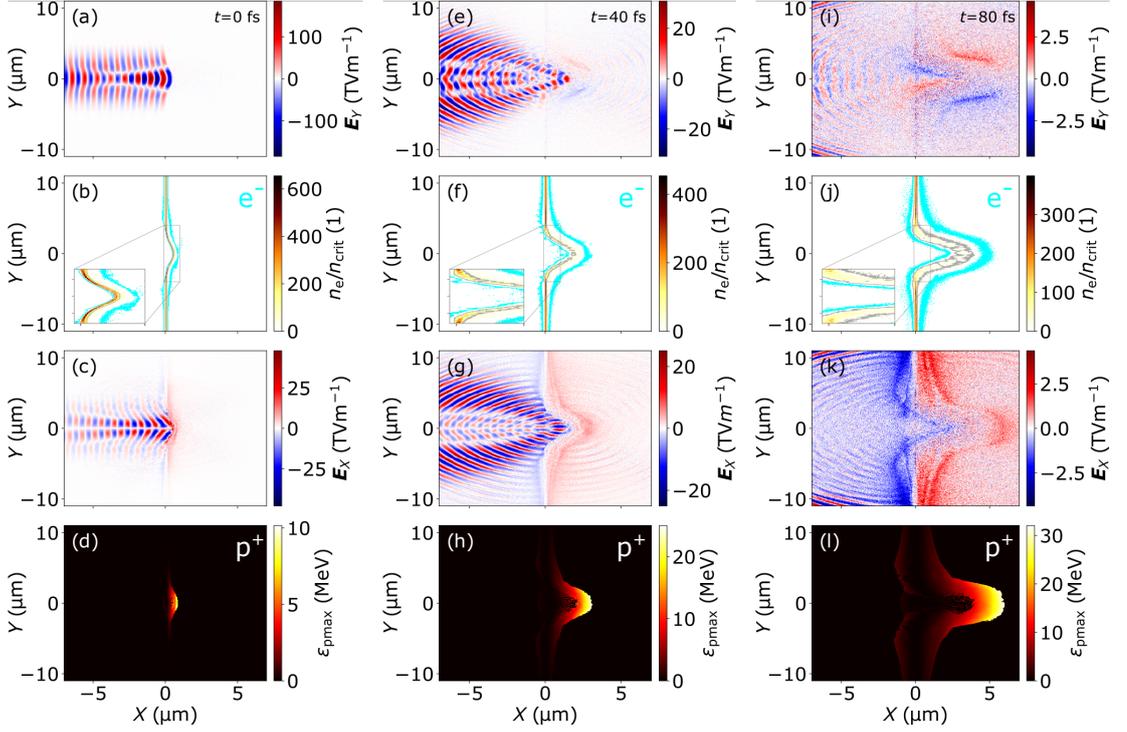


Figure 2.13: Illustration of RPA-HB in a 2D PIC simulation, where a circularly polarised laser pulse interacts with a CH target foil. (a) shows the electric field in the  $Y$  dimension, at time,  $t=0$  fs, when the peak intensity interacts with the (b) target. A cyan and grey contour indicate hole-boring of the classical,  $n_e=n_{\text{crit}}$ , and relativistically corrected,  $n_e=\gamma n_{\text{crit}}$ , critical density surface due to the laser radiation pressure. This generates a (c) longitudinal electric field which accelerates ions to (d) multi-MeV energies. The same electric fields, particle densities, and energies, are shown as the interaction continues at lower intensity in (e-h), and after the laser has diminished in (i-l).

MeV [220].

Having now described RPA-HB, it is useful to expand on the simplified illustration provided in figure 2.12, by showing results from a 2D PIC simulation with the same parameters used to produce TNSA protons in figure 2.11, the only differences being that the thickness of the CH target foil is reduced to  $l=100$  nm, and that the laser pulse is now circularly polarised, with an increased intensity of  $I_L=2\times 10^{21}$  Wcm $^{-2}$ . Results from this simulation are shown in figure 2.13(a-b), at time,  $t=0$  fs, when the peak intensity reflects from the critical density surface, boring a hole in the target, and doubling the peak electron density within the skin depth from the initial CH solid-density,  $n_e=300n_{\text{crit}}$ , to  $n_e\sim 600n_{\text{crit}}$ , where  $n_{\text{crit}}\sim 10^{28}$  m $^{-3}$  for the  $\lambda_L=800$  nm laser pulse. This causes charge separation from the stationary ions left behind, creating a longitudinal electric field of  $\sim 25$  TVm $^{-1}$ , which accelerates protons to multi-MeV energies, as shown in figure

2.13(c-d). As the interaction continues in figure 2.13(e-h) and figure 2.13(i-l), the laser radiation pressure pushes the critical density surface further, accelerating protons to maximum energies of  $\epsilon_{\text{pmax}} > 30$  MeV.

If the target foil is even thinner, on the order of tens of nanometres, the entire irradiated region can propagate with the laser pulse over tens of microns in vacuum, in a motion analogous to wind driving a sail [221]. This so-called RPA ‘light-sail’ (LS) mechanism [222–224] results in ions being accelerated more efficiently, gaining a velocity derived [225] as,

$$\mathbf{v}_{\text{LS}} = \frac{2I_{\text{L}}\tau_{\text{L}}}{l\rho c} \quad (2.48)$$

Again, assuming the laser pulse is perfectly reflected ( $R=1$ ,  $A_{\text{frac}}=0$ ,  $T=0$ ) for simplicity. Noting the inverse relationship to target foil thickness,  $l$ , the maximum kinetic energy of ions accelerated in this regime is,

$$\epsilon_{\text{imax}} = 2m_{\text{i}} \left( \frac{\tau_{\text{L}}I_{\text{L}}}{l\rho c} \right)^2 \quad (2.49)$$

Here, the maximum ion energy scales with peak intensity as  $\epsilon_{\text{imax}} \propto I_{\text{L}}^2$ , in contrast to the linear RPA-HB scaling,  $\epsilon_{\text{imax}} \propto I_{\text{L}}$  [74], and the square root scaling typically associated with TNSA,  $\epsilon_{\text{imax}} \propto I_{\text{L}}^{1/2}$  [33]. To illustrate this increased efficiency, the same simulation parameters used to illustrate RPA-HB in figure 2.13 are now used to simulate RPA-LS by reducing the CH target foil thickness to  $l=10$  nm, and returning to the lower laser intensity,  $I_{\text{L}}=5 \times 10^{20}$  Wcm $^{-2}$ , used in the simulation demonstrating TNSA, in figure 2.11.

Results from this simulation are shown in figure 2.14, and are initially very similar to those shown for a simulation of RPA-HB in figure 2.13, with the peak laser intensity driving hole-boring of the critical density surface, and accelerating protons to multi-MeV energies, as shown in figure 2.14(a-d). However, as the interaction continues in figure 2.14(e-f), the laser radiation pressure forces the critical density surface away from the initial target location, illustrated by a break in the relativistically corrected critical density surface (grey contour), exhibiting this light-sail behaviour which accelerates protons to energies of  $\epsilon_{\text{pmax}} > 20$  MeV.

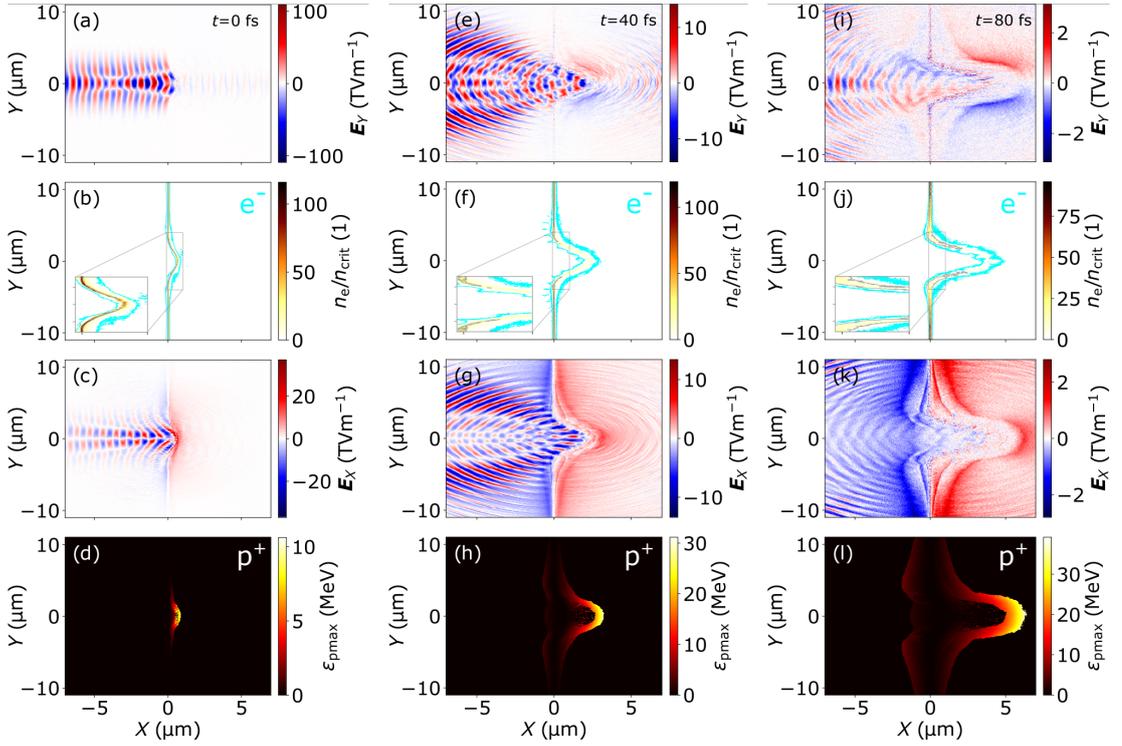


Figure 2.14: Illustration of RPA-LS in a 2D PIC simulation, where a circularly polarised laser pulse interacts with a CH target foil. (a) shows the electric field in the  $Y$  dimension, at time,  $t=0$  fs, when the peak intensity interacts with the (b) target. A cyan and grey contour indicate hole-boring of the classical,  $n_e=n_{\text{crit}}$ , and relativistically corrected,  $n_e=\gamma n_{\text{crit}}$ , critical density surface due to the laser radiation pressure. This generates a (c) longitudinal electric field which accelerates ions to (d) multi-MeV energies. The same electric fields, particle densities, and energies, are shown as the interaction continues at lower intensity in (e-h), with the laser now pushing the critical density surface several microns into vacuum, separating it from the target bulk, in a light-sail mechanism, which continues in (i-l).

The laser pulse continues to push this surface with a reduced intensity in figure 2.14(i-l), resulting in maximum proton energies of  $\epsilon_{\text{pmax}} > 35$  MeV. This is greater than the proton maximum energies driven by a linearly polarised laser pulse of the same intensity and duration via the TNSA mechanism, as shown in figure 2.11, and those driven by the same circularly polarised laser pulse with double the intensity through the RPA-HB mechanism, as shown in figure 2.13. As such, RPA-LS is extremely promising for increasing maximum proton energies beyond that which has currently been achieved in experiments.

Even more promising is that both radiation pressure based schemes are predicted to produce quasimonoenergetic ion sources [221, 225], since ions all originate from the same very thin layer, rather than from different regions in contaminant layers, and throughout much thicker targets, in TNSA. Furthermore, if

driven by a laser with a focal spot large enough to minimise hole-boring induced curvature of the target front surface, RPA can produce ion beams with lower divergence than TNSA. If realised, these beam properties would be useful for applications such as proton therapy, where it is desirable to deliver a high number of protons to a small tumorous area, at a given energy, whilst delivering minimal protons at different energies to healthy tissue surrounding the tumour [155].

As stated, experimental demonstrations of RPA are much less prolific than those demonstrating TNSA, because this is a very sensitive acceleration regime achieved over a relatively narrow range of interaction conditions, in contrast to TNSA which can be achieved for a wide range of laser pulse energy, duration, intensity, and target foil thickness values [30]. One particular challenge is that RPA simultaneously requires high laser intensities, but also that the target does not become transparent to laser light via RSIT, induced at such high intensities as described in the next section. If this occurs, the radiation pressure reduces to zero ( $R=0$ ,  $T=1$ ), according to equation 2.42, diminishing RPA. That being said, the onset of RSIT does not inherently diminish ion acceleration. On the contrary, it has been shown to enhance maximum proton energies to the highest recorded for a laser-driven source [31, 71], through a combination of TNSA, RPA, and other mechanisms which will now be discussed.

### **2.6.3 Acceleration enhanced by relativistic self-induced transparency**

Recapping the description provided in section 2.4.2, laser light can propagate in plasma where the electron density is lower than the critical density,  $n_e < n_{\text{crit}}$ . Termed underdense plasma, this was initially discussed to form as part of an expanding pre-plasma induced on the front surface of a relatively thick target foil. In this case an incident laser pulse propagates within underdense plasma before reflecting from the critical density surface, being unable to propagate in overdense plasma beyond this surface. However, if the laser radiation pressure drives the critical density surface into the target foil as discussed in section 2.6.2,

the laser pulse will essentially reflect from the foil front surface. This allows laser energy to be transferred to electrons deeper within the foil, increasing their mass by a relativistic gamma factor,  $\gamma$ , thus increasing the plasma critical density according to equation 2.29. The electron density in a CH target foil, used for examples so far, is  $n_e=300n_{\text{crit}}$ , where  $n_e\sim 1\times 10^{27}\text{ m}^{-3}$  and  $n_{\text{crit}}\sim 300\times 10^{27}\text{ m}^{-3}$  for a  $\lambda_L=800\text{ nm}$  laser pulse, according to equation 2.29, and equation 2.11. Stating  $a_0=\sqrt{2(\gamma^2-1)}$ , and relating to intensity through equation 2.22, a laser of this wavelength would require an intensity of  $I_L\geq 10^{23}\text{ Wcm}^{-2}$  to heat electrons to a velocity corresponding to  $\gamma\sim 300$ , so that an initially overdense target could become underdense,  $n_e < \uparrow\gamma n_{\text{crit}}$ , through so-called relativistic induced transparency (RIT). This is orders of magnitude higher than the intensity of laser pulses that have already been demonstrated to propagate through plasma [65, 69–71, 226], which can be explained by considering that laser-to-electron heating not only increases the electron mass, but also causes the plasma to expand, decreasing the electron density. As an example, a  $\lambda_L=800\text{ nm}$  laser pulse of intensity  $I_L\sim 4\times 10^{20}\text{ Wcm}^{-2}$ , will propagate in initially overdense plasma if it heats electrons such that their density reduces to  $n_e=5\times 10^{27}\text{ m}^{-3}$ , and their velocity increases to a value corresponding to  $a_0\sim 14$ , increasing the critical density by a factor,  $\gamma=10$ , to  $n_{\text{crit}}=20\times 10^{27}\text{ m}^{-3}$ . This combined effect,  $\downarrow n_e < \uparrow\gamma n_{\text{crit}}$  is known as relativistic self-induced transparency (RSIT), and has been utilised to generate high order modes of light [140], and to control electron and proton beam structures [69, 70].

The onset of RSIT can also significantly enhance laser-driven ion acceleration [31, 71], as now demonstrated in results from a 2D PIC simulation where a  $l=30\text{ nm}$  CH foil is irradiated by the same linearly polarised laser pulse of  $I_L=5\times 10^{20}\text{ Wcm}^{-2}$  intensity, used to drive TNSA in figure 2.11. Returning to a linearly polarised laser pulse, the  $\mathbf{j}\times\mathbf{B}$  absorption mechanism is no longer mitigated as required to induce RPA in the simulations presented in figure 2.13 and figure 2.14, and is instead instrumental in heating electrons, relativistically increasing their mass, and reducing electron density along the laser axis. Features of this simulation are shown in figure 2.15 for three timesteps that are slightly different from those shown for the previous acceleration mechanisms, now  $t=-20\text{ fs}$ ,  $t=0\text{ fs}$ ,

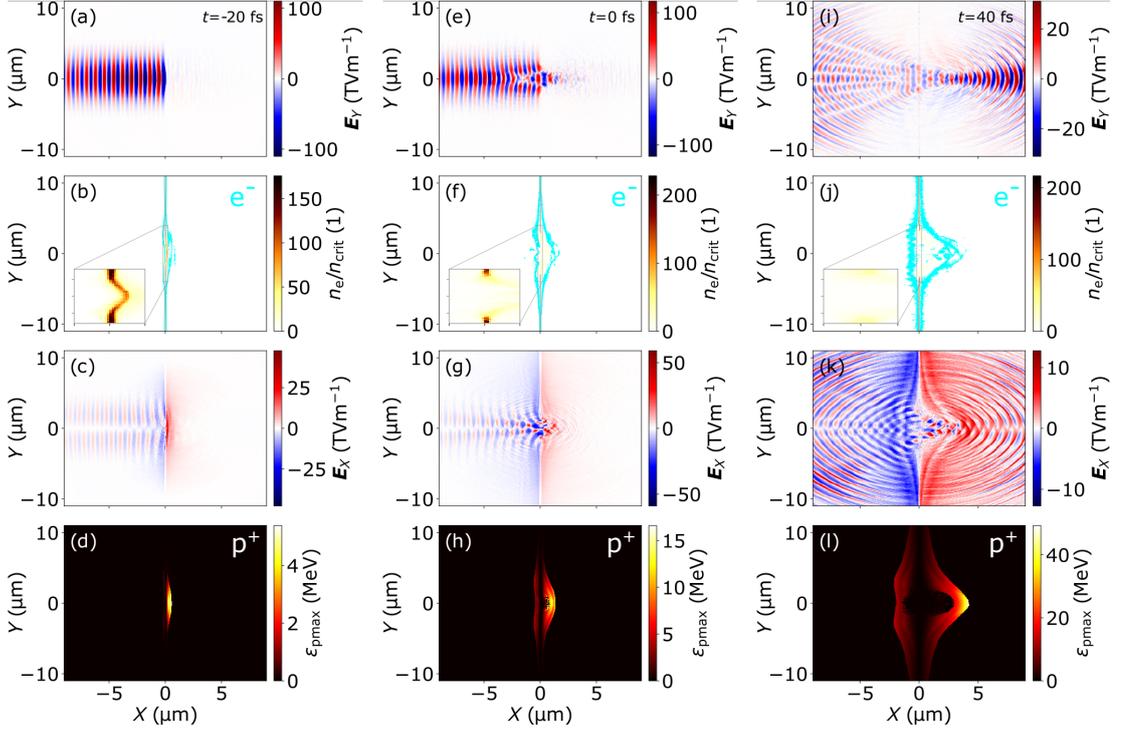


Figure 2.15: Illustration of RSIT-enhanced acceleration in a 2D PIC simulation, where a linearly-p polarised laser pulse interacts with a CH target foil. (a) shows the laser electric field in the polarisation direction ( $Y$ ), at time,  $t=-20$  fs, before the peak intensity interacts with the (b) target. A cyan and grey contour indicate not only hole-boring of the classical,  $n_e=n_{\text{crit}}$ , and relativistically corrected,  $n_e=\gamma n_{\text{crit}}$ , critical density surface, but also fast electrons which propagate to the target rear and into vacuum. This results in (c) longitudinal electric fields associated with the radiation pressure, and the sheath electrons at the target rear, with the combined fields (d) accelerating protons in an RPA-TNSA hybrid mechanism. When the peak intensity arrives in (e-h), the electron density in the focal spot region has decreased below the relativistically corrected critical density,  $n_e < \gamma n_{\text{crit}}$ , and the laser pulse propagates through the target, directly accelerating protons. This continues in (i-l).

and  $t=40$  fs, relative to the time at which the peak laser intensity interacts with the foil front surface. Figure 2.15(a-d) shows the laser pulse beginning to interact with the target, heating electrons to energies greater than their rest mass, and reducing the peak target density below  $n_e \sim 160 n_{\text{crit}}$ , even further in the focal spot region. This creates a fast electron population which propagates through the target, creating an electric sheath field which accelerates protons via the TNSA mechanism. At this time, the laser radiation pressure simultaneously drives hole-boring of the critical density surface, inducing a longitudinal electric field which is temporally separated from the TNSA field. The result is a temporally dual-peaked electric field [227], resulting in a hybrid acceleration mechanism, where both TNSA and RPA combine to accelerate protons to multi-MeV energies [65,

71].

By the time the peak laser intensity arrives in figure 2.15(e-h), electrons within the focal spot region have been reduced to a density on the order of the classical critical density, and have been accelerated to relativistic velocities, so that  $\downarrow n_e < \uparrow \gamma n_{\text{crit}}$ . At this point the laser pulse can propagate through the relativistically underdense region, as illustrated clearly for a later simulation time in figure 2.15(i-l). This enables extremely efficient laser-to-electron energy transfer via direct acceleration, increasing the electron temperature, thus boosting the longitudinal electric field, and the maximum energy of protons to  $\epsilon_{\text{pmax}} \sim 40$  MeV, above that which can be achieved purely through TNSA, RPA-HB, or RPA-LS, illustrated in figures 2.11, 2.13, and 2.14, respectively.

Publications [64, 75–77] have proposed this enhanced ion acceleration to result from a Buneman-like instability, with energy being transferred from electrons travelling at relativistic velocity to much slower protons, due to this inherent velocity difference. The energy lost by electrons is then regained as they continue to interact with the propagating laser light, growing the instability and boosting the energy of ions initially accelerated by TNSA and RPA in a so-called break out afterburner (BOA) process. Protons accelerated in this manner have been predicted to gain maximum energies which scale linearly with the laser electric field, and thus with the square root of intensity,  $\epsilon_{\text{pmax}} \propto I_L^{1/2}$  [228]. Other publications [65, 66, 78] agree that laser-to-electron energy transfer is very efficient in a target undergoing RSIT, and that this boosts the energy of protons accelerated initially by a hybrid RPA-TNSA hybrid mechanism, though disagree that this results from a Buneman-like instability. In this description the direct laser acceleration of electrons can also result in a plasma jet, with sufficient electron current density to induce a strong magnetic field perpendicular to the laser pulse propagation direction [65–67, 78]. This magnetic field can act as a pinching force on the electrons, in the same way as described in section 2.5, effectively confining the plasma jet to the laser pulse propagation axis. The electron current density can also induce a longitudinal electric field, enhancing the existing TNSA-RPA hybrid electric field. Ions can then be further accelerated by the enhanced elec-

tric field, and have been demonstrated to steer from the target normal axis, along which they are initially accelerated by TNSA and RPA, to the laser pulse propagation axis, as demonstrated in experiments and simulations where the target foil is angled with respect to the incident laser [65, 71].

Work to understand the subtle and complex physics that govern enhanced acceleration in the RSIT regime is ongoing [78], but what is clear is that the maximum ion energy achieved in an RSIT-enhanced regime depends strongly on the target foil thickness. This has been demonstrated experimentally [71], with maximum proton energies of up to  $\epsilon_{\text{pmax}} \sim 96$  MeV being produced at an optimal target foil thickness, a then record energy for laser-driven protons. This optimisation was replicated in simulations, which demonstrated proton energy enhancement to be most significant when target foils underwent RSIT at the same time as the peak laser intensity arrived at the front surface, with the maximum proton energy reducing if RSIT occurred too early, in thinner foils, or too late, in thicker foils [79]. In a subsequent experiment, RSIT was induced in a similar target foil using a laser with higher temporal-intensity contrast [45], resulting in maximum proton energies of up to  $\epsilon_{\text{pmax}} \sim 150$  MeV [31], by far the highest recorded for a laser-driven source.

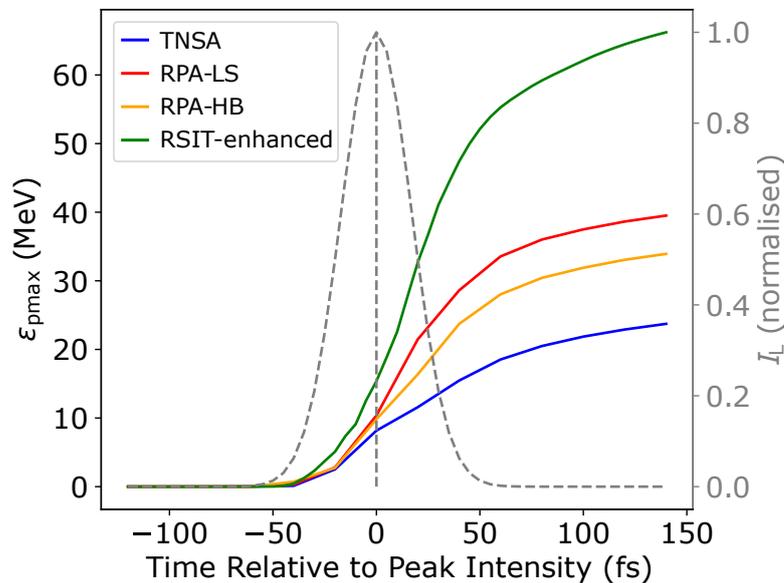


Figure 2.16: Maximum energy of protons accelerated by TNSA, RPA-HB, RPA-LS, and an RSIT-enhanced mechanism, for each timestep of the simulations presented in figures 2.11, 2.13, 2.14, 2.15, respectively. The temporal-intensity profile of the laser is the same in each case, and is illustrated as it rises to a normalised intensity.

RSIT-enhanced acceleration is then promising for increasing proton energies towards that required for applications such as radiotherapy [155]. However, there are indications that this regime does not produce high maximum proton energies with the same repeatability as a pure TNSA mechanism. This is investigated in chapter 5, so that RSIT-enhanced acceleration can be understood and controlled to consistently produce maximum proton energies beyond that achievable using TNSA or RPA alone. As a final indication of the potential of RSIT-enhanced acceleration, figure 2.16 compares the maximum proton energy at each timestep of the simulation illustrated in figure 2.15 to that of the simulations demonstrating TNSA, RPA-HB, and RPA-LS in figure 2.11, figure 2.13, and figure 2.14, respectively. The Gaussian temporal-intensity profile of the laser pulse is the same in each case, rising to a peak of  $I_L=5\times 10^{20}$  Wcm<sup>-2</sup> in the RSIT-enhanced, TNSA, and RPA-LS simulations, and  $I_L=2\times 10^{21}$  Wcm<sup>-2</sup> in the RPA-HB simulation. It should be noted that the RSIT-enhanced mechanism was optimised by selecting a target foil thickness of  $l=30$  nm, so that the target underwent RSIT as the peak laser intensity arrived, and that there are ways to optimise the other mechanisms. This is demonstrated in chapter 4, where a Bayesian optimisation algorithm selected an optimal pre-plasma density profile which increased laser-to-electron energy absorption, in turn increasing a TNSA sheath field, and the maximum energy of protons accelerated by this mechanism.

## 2.7 Summary

To summarise, this chapter has provided a brief overview of the physics underpinning laser-driven acceleration of ions to multi-MeV energies. This began by discussing the fundamental properties of intense,  $I_L>10^{18}$  Wcm<sup>-2</sup>, short duration,  $\tau_L=(30-1000)$  fs, laser pulses in section 2.1, and continued by describing the mechanisms by which this laser light creates plasma by ionising the front surface of solid target foils, in sections 2.2 and 2.3. Electrons within this plasma were then described to absorb laser energy through a variety of mechanisms in section 2.4, accelerating them to relativistic velocities which they carry into overdense

plasma as described in section 2.5, setting up electric fields which accelerate ions to multi-MeV energies through several mechanisms discussed in section 2.6. This provides a foundation for investigating the optimisation and stabilisation of laser-driven ion acceleration in chapters 4 and 5, using methodologies which will now be described in chapter 3.

# Chapter 3

## Methodology

Having discussed the fundamental physics that occurs when short duration, high-power laser pulses interact with target foils in chapter 2, it is now necessary to describe how these pulses are created and amplified to the intensities required to drive the physics of interest to this thesis. This is presented in the context of the Vulcan-PW laser system, since this was used for the experimental investigation detailed in chapter 5, and shares many processes with high-power laser systems across the world, such as optical parametric, chirped pulse, and regenerative amplification. Methods used to characterise the laser energy, pulse duration, temporal-intensity contrast, and focal spot quality, are described, as are components used to improve these properties, such as plasma mirrors, adaptive optics, and wavefront sensors. Diagnostics used to characterise radiation produced in a laser-solid interaction are also covered, and the chapter closes by describing methods underpinning particle-in-cell (PIC) simulations, which inform and complement experimental investigations, and are employed in chapters 2, 4, and 5.

### 3.1 Creation and amplification of short duration, high-power laser pulses at the Vulcan-PW facility

The first component in the Vulcan-PW laser system is a passively mode locked oscillator consisting of a Ti:Sapphire gain medium, pumped by a continuous wave (CW), frequency doubled,  $2\omega$ , Nd:YAG laser [58]. This creates a train of pulses, each of which has a central wavelength of  $\lambda_L=1053$  nm, and contains  $E_L\sim 5$  nJ of energy, in a duration of  $\tau_L\sim 120$  fs. One pulse is selected from this train and pre-amplified using the optical parametric chirped pulse amplification (OPCPA) technique [229, 230], illustrated in the inset of figure 3.1(a). Here, the pulse selected from the oscillator is separated into two pulses, which are transmitted through, and reflected by, a pellicle beamsplitter (PBS) with 70%, and 30%, of the initial energy, respectively. The transmitted pulse then acts as a seed pulse for a regenerative amplification process, illustrated in figure 3.1(b). The incoming seed pulse for the regenerative amplifier is reflected by a polariser according to the pulse polarisation, in this case assumed to be linear p-polarisation. The pulse then passes through a half waveplate, which rotates the polarisation by  $90^\circ$  to linear s-polarisation, and then through a Faraday rotator, which can adjust the polarisation more finely according to the strength of an externally applied magnetic field.

A second polariser is then set up to reflect the linearly s-polarised pulse into a Pockels cell, which behaves as a quarter waveplate when supplied a certain voltage, acting to rotate the pulse polarisation to circular. The pulse then reflects from a cavity mirror, reversing the circular polarisation direction, causing the pulse polarisation to rotate by  $90^\circ$  back to linear p-polarisation, upon a second pass through the Pockels cell. As the second polariser was set up to reflect linear s-polarisation, it transmits the now linearly p-polarised pulse, which propagates into an optically pumped Nd:YLF gain medium, reflects from a second cavity mirror, and returns through the gain medium, gaining energy on both passes.

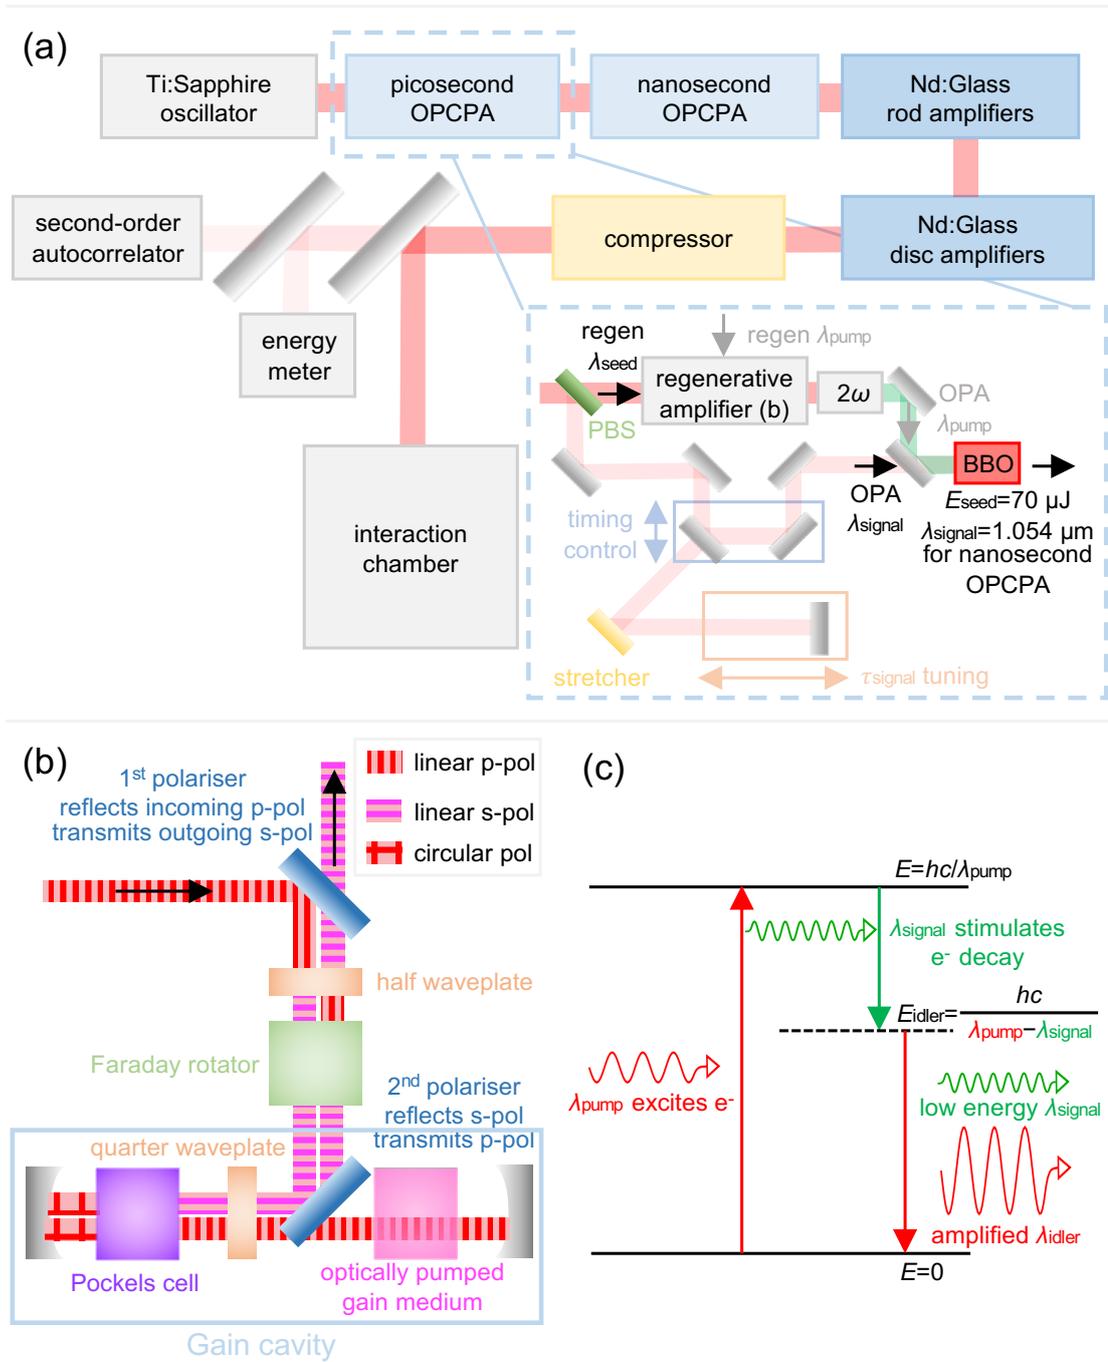


Figure 3.1: (a) Illustration of pulse amplification and temporal compression in stages of the Vulcan-PW laser system, resulting in a final pulse characterised by an energy meter and second-order autocorrelator, and propagated into an interaction chamber. A detailed illustration of the picosecond OPCPA stage is shown in the inset, and includes a regenerative amplifier which operates as illustrated in (b), creating a pump pulse for the OPA process illustrated in (c).

The slightly amplified pulse is then transmitted back through the second polariser and Pockels cell, reflecting from the first cavity mirror, and back to the gain medium. The pulse is therefore trapped in this gain cavity, gaining energy over

many round trips, until the voltage supplied to the Pockels cell is switched off, causing the pulse to become linearly s-polarised, so that it can reflect from the second polariser, and out of the cavity. A magnetic field is then applied to the Faraday rotator such that the pulse becomes linearly p-polarised, rotating to become linearly s-polarised upon transmission through the half waveplate, allowing it to transmit through the first polariser.

Through this process, the initially low energy,  $E \sim 5$  nJ, seed pulse is amplified to an energy of  $E \sim 500$   $\mu$ J. This pulse is then frequency doubled, and used as a pump pulse for a picosecond OPCPA process, as shown in the inset of figure 3.1(a), which also shows 30% of the initial oscillator pulse being reflected by the pellicle beamsplitter, to be used as a signal pulse for this process.

OPA is illustrated in figure 3.1(c), and requires the signal and pump pulse profiles to match so that energy is extracted efficiently. The pump pulse duration increases to  $\tau_{pump} \sim 15$  ps due to gain narrowing in the regenerative amplification process, and so the signal pulse is stretched using a grating, and tuned to a duration of  $\tau_{signal} \sim 3$  ps, using a moveable mirror to adjust the signal pulse path, as illustrated in the inset of figure 3.1(a). This process means the signal and pump pulse are temporally overlapped upon entering a  $\beta$ -barium borate (BBO) crystal which is birefringent, ensuring the pump, and signal pulse, at different wavelengths, are phase matched upon passing through the material.

This facilitates OPA, which begins with the high energy, longer wavelength, pump pulse,  $\lambda_{pump} \sim 1054$  nm, exciting electrons in the BBO crystal to a higher energy state,  $E = hc/\lambda_{pump}$ . The lower energy seed pulse then stimulates these electrons to decay to a lower energy intermediate state, converting the pump pulse to a low energy signal pulse, with the difference in energy,  $E_{idler} = hc/(\lambda_{pump} - \lambda_{signal})$ , being radiated as a so-called idler photon. Of the initial  $E_{pump} \sim 500$   $\mu$ J pump energy, only  $\sim 15\%$  is converted to the idler photon, resulting in an output pulse of energy,  $E_{idler} \sim 70$   $\mu$ J, which seeds the next, nanosecond OPCPA stage of the amplification process, as shown on figure 3.1(a). This seems counter-intuitive to the goal of amplifying the pulse energy, but this stage is included to improve the temporal-intensity profile of the final high energy pulse [58].

In OPA, the pump pulse energy is almost completely transferred to the signal and idler pulse in a stimulated process, minimising amplified spontaneous emission (ASE), since there is essentially no additional heating of the gain medium, unlike in regenerative amplification. As discussed in section 2.1, ASE light precedes the main Vulcan-PW pulse and, if not compensated for, can significantly pre-expand or even completely destroy a target foil before the main pulse arrives, by driving a shock through the target to induce plasma formation on the rear surface, diminishing ion acceleration mechanisms [43, 44]. The nanosecond OPCPA stage was initially designed to mitigate this issue [230, 231], though later work showed significant ASE was still resulting from this stage [58].

The other benefit of OPA over other amplification techniques is that any additional heating that does occur is only stored in the gain medium over the pump pulse duration. ASE therefore only occurs on this timescale, rather than being generated when parasitic pulses reflect through the gain medium at later times in, for example, a regenerative amplifier. Using a picosecond OPCPA stage to generate a signal pulse for the nanosecond OPCPA process thus ensures the temporal-intensity profile of this pulse is improved compared to other approaches, and allows the pulse duration to be finely adjusted to that desired in this next stage.

However, before the output pulse from the picosecond OPCPA stage can be used as a signal pulse in the nanosecond OPCPA stage, it must first be stretched to a duration of  $\tau \sim 5$  ns, to avoid damaging optics as the pulse is amplified. This was not a concern in the picosecond OPCPA stage, as optics damage according to the energy delivered in a given time, otherwise known as the optical fluence,

$$F = \frac{E_L}{S} \frac{1}{\sqrt{\tau_L}} \quad (3.1)$$

The pulse created in the picosecond OPCPA process has an energy of  $E = 70 \mu\text{J}$ , a duration of  $\tau \sim 15$  ps, and a beam area of  $S = \pi r^2$ , where  $r \sim 0.05$  cm, resulting in a fluence of,  $F \sim 2 \times 10^{-3} \text{ Jcm}^{-2}\text{ps}^{-1/2}$ , safely below the threshold of  $F_{DT} \sim 30 \times 10^{-3} \text{ Jcm}^{-2}\text{ps}^{-1/2}$  at which optics begin to damage for pulses of this duration.

However, the goal of the nanosecond OPCPA stage is to amplify the energy to

$E \sim 1$  mJ, which, if all other parameters of equation 3.1 are kept constant, would increase the fluence to  $F \sim 33 \times 10^{-3} \text{ Jcm}^{-2}\text{ps}^{-1/2}$ , inducing damage to optics due to the ionisation mechanisms discussed in section 2.2 [232]. The fluence could be reduced below this so-called damage threshold (DT) by tripling the beam radius, though optics would then have to be nine times larger to avoid clipping the beam, increasing their cost. This may be feasible in the nanosecond OPCPA stage, but would be completely unfeasible in the later stages illustrated in figure 3.1(a), where the energy is amplified to hundreds of Joules. At this stage the optical intensity would also be  $I_L \geq 10^{16} \text{ Wcm}^{-2}$ , if the spot size and pulse duration were kept constant, which for most optical materials [233, 234] means the intensity dependent refractive index, described by the optical Kerr effect [235], becomes important. Due to the Gaussian spatial-intensity profile of the pulse, high-intensity light at the centre of the focal spot will experience a stronger refractive index when propagating in optics compared to lower intensity light in the pulse wings, effectively creating a focusing lens which reduces the focal spot size, increasing the pulse intensity, resulting in more significant self-focusing in a feedback process. The Gaussian temporal-intensity profile of the pulse causes a similar effect, inducing self-phase modulation which acts to broaden the pulse bandwidth [236], thus decreasing its duration, and increasing intensity, whilst also introducing structures to the intensity spectrum, causing deviation from a Gaussian-like profile. To mitigate these effects, the pulse duration is stretched to  $\tau \sim 5$  ns using a pair of diffraction gratings to induce different path lengths for the different wavelengths present in the pulse bandwidth, so that the fluence damage threshold is increased to  $F_{\text{DT}} \sim 0.5 \text{ Jcm}^{-2}\text{ps}^{-1/2}$ , and the pulse intensity is reduced, allowing pulses to carry energies of up to  $E_L = 250$  mJ without damaging optics or inducing intensity dependent nonlinear effects. This is the chirped pulse amplification (CPA) aspect of the OPCPA process, which was initially demonstrated in 1985 [26], and shared the Nobel prize for physics in 2018, due to the influence of this technique in facilitating intensities beyond  $I_L \sim 10^{18} \text{ Wcm}^{-2}$  for the first time, giving rise to many different avenues of physics research, including laser-solid interactions.

The stretched,  $\tau \sim 5$  ns, output of the picosecond OPCPA stage is then used as a signal pulse for the nanosecond OPCPA stage, which uses the same OPA technique described before in a slightly different configuration, now using a frequency doubled Nd:YAG laser, operating at 10 Hz, to deliver pump energies on the order of tens of milliJoules to 3 separate BBO crystals [58, 230]. Again,  $\sim 15\%$  of the pump energy is converted to idler photons produced in this process, resulting in an energy on the order of milliJoules, contained in a temporally clean pulse, without significant ASE. This pulse is then propagated into the first main amplification stage of the Vulcan system, which consists of a series of Nd:Glass rod amplifiers pumped by white light flash lamps [58]. As the beam is amplified through this system, it is expanded from an initial radius of  $r=0.9$  cm, to a final radius of  $r=10.4$  cm, so that the optical fluence is reduced below the damage threshold, supporting an output pulse energy of  $E \sim 85$  J. This pulse is then amplified further by a series of Nd:Glass disc amplifiers, reaching a final maximum energy of  $E \sim 500$  J.

The fully amplified pulse is then propagated through a spatial filter to remove noise from the pulse wings and enhance the spatial beam quality, before being transported to a compressor comprising diffraction gratings with opposite sign dispersion to the gratings used to temporally stretch the pulse before the nanosecond OPCPA stage. The pulse is then re-compressed to a final duration of  $\tau_L \sim 1000$  fs with  $\sim 60\%$  efficiency, reducing the energy to  $E=300$  J. The beam is expanded to a radius of  $r=30$  cm for this final stage to reduce fluence on the compressor grating, and on an off-axis parabolic mirror (OAP) of  $f$ -number,  $F_{\#}=3$ , which focuses the beam to a spot size of FWHM,  $\phi_L \sim 4$   $\mu\text{m}$ , producing an intensity of  $I_L \geq 10^{20}$   $\text{Wcm}^{-2}$ .

The pulse energy is measured on each full power Vulcan-PW shot using an energy meter which absorbs 1% of the beam, transmitted through a final turning mirror, as illustrated in figure 3.1(a). The duration of the main pulse is also measured from this transmitted beam using a single-shot, second-order autocorrelator [56, 237], which splits the beam into two pulses that are focused into a crystal which exhibits a  $\chi^{(2)}$  nonlinearity, thus creating an output pulse of

shorter wavelength, in a process similar to OPA. The output light spreads due to the cylindrical lens used for focusing, and is then captured by a camera, with the position of light on the detector corresponding to different spatial positions in the crystal, which can be reconstructed into the pulse temporal-intensity profile.

By adjusting the path length of one of the pulses, the timing delay between the pulses can be increased, so that the intensity can be measured at a different point on the temporal-intensity profile. Doing this over many shots, the entire temporal-intensity profile can be reconstructed over a relatively high dynamic range, in a scanning third-order autocorrelator technique, an example output of which is shown in figure 2.2 of section 2.1, for the Vulcan-PW laser. Despite significantly reducing ASE by employing OPCPA over two amplifying stages, it is still apparent in this pulse temporal-intensity profile. Also shown is a pre-pulse  $\sim 200$  ps before the main pulse arrives, which typically originates from a reflection on the back surface of an optic within the laser system. The front surface of these optics will typically be  $\geq 99\%$  reflective, only transmitting  $\leq 1\%$  of the beam to be reflected from this back surface, explaining why this pre-pulse is much less intense than the main pulse. However, after amplification, this pre-pulse is still intense enough to pre-ionise a solid-density target. Wedged optics are used throughout the Vulcan-PW system to eliminate these pre-pulses [58], though it is very difficult to eliminate every one. Post-pulses highlighted on figure 2.2 originate from unwanted reflections in the same way, and can actually generate identical pre-pulses through intensity dependent non-linear refractive index effects [82]. The final feature highlighted on figure 2.2 is the rising edge of the intensity profile. This results from certain wavelengths of the main pulse not being optimally compressed due to imperfections on gratings in the compressor stage, or slight misalignment of the beam onto these gratings, resulting in so-called uncompensated dispersion [57]. Research to further reduce these typically undesirable features of the temporal-intensity profile is ongoing, though it is likely impossible to achieve a completely Gaussian temporal-intensity profile at such high intensities.

The importance of the laser temporal-intensity contrast has been known for

a long time, and so methods have been developed to gate the temporal-intensity profile, so that much of the lower intensity features don't pre-expand a target foil before the main pulse arrives. One method by which this is achieved is by frequency doubling the laser light in an intensity dependent process, and then using dichroic mirrors to only reflect the highest intensity light which has been converted to the shorter wavelength. This scheme does come with some difficulties, the first being that the doubling crystals must be very thin, large, and therefore expensive to manufacture, to minimise nonlinear phase issues and optical damage [82]. More importantly,  $\sim 50\%$  of the incident laser energy is lost in this process. In addition, electron heating is proportional to the laser wavelength as described in chapter 2, and so light at half the wavelength will induce much lower electron temperatures, which is detrimental to most applications of laser driven radiation sources. As such, the most common method by which the laser-temporal intensity profile is 'cleaned' is by utilising so-called plasma mirrors.

## 3.2 Plasma mirrors

Figure 3.2 shows a plasma mirror utilised on an experiment at the Gemini laser facility [238], alongside an illustration of how this device operates. The plasma mirror shown consists of a glass substrate with an anti-reflective coating, such that it is  $\geq 96\%$  transmissive to light at the laser wavelength, in this case  $\lambda_L = 800$  nm. Laser light is focused onto the plasma mirror with a diameter on the order of a few millimetres, and initial ASE light with intensity below the ionisation threshold,  $I_L \sim 10^{11} \text{ Wcm}^{-2}$ , is transmitted through the surface. The light intensity then increases as the interaction continues, at some point reaching  $I_L \sim 10^{11} \text{ Wcm}^{-2}$ , which is sufficient to ionise the plasma mirror surface, according to the mechanisms described in section 2.2. This creates an overdense plasma, which reflects laser light at the critical density surface, as described in section 2.4.2. Some laser energy is absorbed by electrons in the overdense plasma according to the mechanisms described in section 2.4.3, though the remaining  $\sim (60-85)\%$  is reflected, and continues focusing onto a target foil. Importantly, it is therefore

predominantly the rising edge intensity profile which pre-ionises the target foil for a few picoseconds before the main pulse arrives, rather than predominantly ASE light over hundreds of picoseconds, or even nanoseconds. As a result, plasma mirrors have been shown to improve the laser temporal-intensity contrast, the ratio between the peak intensity and preceding light intensity, by up to two orders of magnitude. Plasma mirrors are often therefore utilised to improve the temporal-intensity profile for laser-solid interactions, and are mandatory in some experiments to protect optics in the laser system from damage due to laser light being reflected from the target foil, since by this point in the interaction overdense plasma on the mirror has expanded inhomogeneously due to the spatial-intensity profile of the main pulse, causing back-reflected light to be diffusely scattered, rather than efficiently reflected back into the laser system.

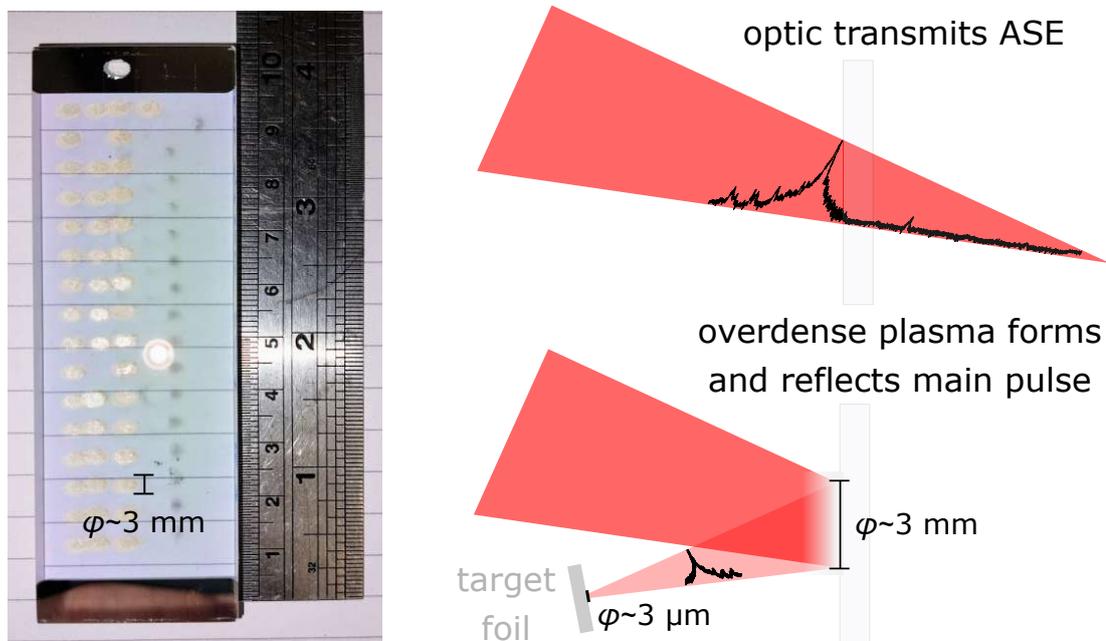


Figure 3.2: Photograph and operating principle of a plasma mirror utilised to improve the laser temporal-intensity profile before irradiation of a target foil.

Many instabilities can arise in plasma due to this inhomogeneous expansion, and so it is crucial that ionisation only occurs a few picoseconds before the main pulse arrives, minimising the time in which the overdense plasma can expand, and these instabilities can grow, to degrade the wavefront quality of the main pulse upon reflection [54]. This is extremely important, as wavefront aberrations

can significantly reduce the quality of the final focal spot. For this reason, the beam wavefront is commonly measured before being focused onto a PM and/or target foil, using a Shack-Hartmann wavefront sensor [239], and is then improved using an adaptive optic (AO) if necessary. This is the final step in the Vulcan-PW laser system before the beam is sent into the experimental vacuum chamber, and will now be discussed.

### 3.3 Laser focal spot improvement and characterisation

Before the final high energy output beam of the Vulcan-PW laser is sent into a vacuum chamber to interact with a target foil, it is reflected from an AO, and imaged by a Shack-Hartman wavefront sensor [239], typically placed beyond the intended interaction plane, which contains an array of microlenses that form images of different regions on the beam profile.

If the incident wavefront is perfectly flat, with no aberrations, each beam sample is focused along the optical axis of the respective lens, forming an equally spaced, grid like pattern when imaged by a camera. If, however, there are aberrations in the wavefront, the focusing of some beam samples will shift by an amount relative to the severity of the aberration. This relative shift is characterised by fitting high-order Zernike polynomials to the phase reconstructed from the centroids of the offset spots, to measure deviation from a flat wavefront induced by different types of aberration, in terms of wavelength. Each Zernike coefficient is then input into software which determines how a deformable mirror on the AO should contort to compensate for these deviations from a flat wavefront, and it then contorts in this way using an array of piezoelectric actuators, to correct the measured aberrations. This ensures the pulsed beam has a relatively flat wavefront as it enters the experimental vacuum chamber, where it is then directed to a target foil using a turning mirror and an  $F_{\#}=3$  OAP, as shown in figure 5.1(a) for the experimental set up used in chapter 5.

This relatively flat wavefront makes it possible to focus the initially  $r=30$  cm

beam to a spot with a Gaussian-like spatial intensity profile and a diameter close to the minimum achievable given the diffraction limit. The beam alignment and focal spot spatial intensity profile are adjusted before every shot by monitoring the focal spot of a relatively low power ( $\sim$ mW) CW beam that is representative of the full energy pulsed beam. This is done using a so-called focal spot camera with objective lens focusing to magnify the beam by  $20\times$ , placed in the position where a target foil will be placed. Once optimised, an image of the focused CW spot is saved so that the focal spot diameter, and the energy encircled within this region, can be measured as demonstrated in figure 3.3.

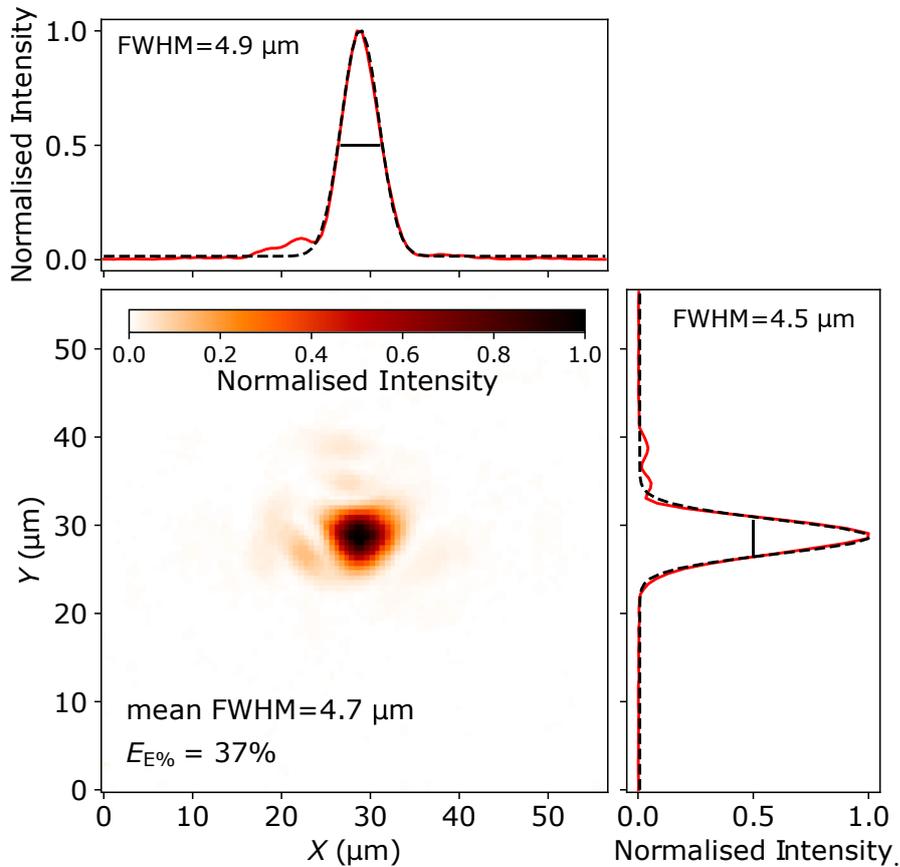


Figure 3.3: The central window shows a focal spot achieved by focusing a low energy CW equivalent of the Vulcan-PW laser beam onto a CCD camera using an  $F_{\#}=3$  OAP. The focal spot diameter is calculated by fitting a Gaussian profile across the centre of the spot in the  $X$  (top) and  $Y$  (right) dimensions and taking an average. The percentage of energy encircled within this focal spot is calculated as the ratio of the summed signal within one FWHM of the spot centre in both dimensions, compared to the summed signal in the total image.

Here, an example focal spot image taken before one of the experimental interactions presented in chapter 5 is shown alongside line-outs of the intensity from

(-25→25)  $\mu\text{m}$  around the focal spot centre in the  $X$  and  $Y$  dimensions, respectively. A Gaussian profile is fit to each of these lineouts, and the associated full width at half maximum (FWHM) values are averaged to provide a focal spot size of  $\phi_L=4.7 \mu\text{m}$ . The encircled energy within the focal spot is then calculated by summing the intensity within one FWHM of the spot centre in the  $X$  and  $Y$  dimensions respectively, then dividing by the total energy delivered to the OAP. In this case, 37% of the laser energy delivered to the OAP is contained within  $\phi_L=4.7 \mu\text{m}$ . The code used to make this measurement contains features which make it relatively robust to signal anomalies and to changes in the focal spot camera and laser system being used. It was developed for the work presented in this thesis and is available as part of the the laser-plasma interaction python library (LPI-Py), which contains various python codes to generate, manage, and analyse experimental data, and to model laser-plasma interactions [101].

After optimising the focal spot in this way, the software that monitors the beam quality with the Shack-Hartmann wavefront sensor is run again, and the AO is adjusted to correct for the final measured distortions. This is extremely important, as it was previously discovered that the defocus aberration of the Vulcan-PW laser beam increases significantly after each full power shot due to thermal lensing resulting from the amplification stages, taking  $\sim 1$  hour to return to the previous minimised value. Even then, the beam defocus was shown to oscillate due to temperature gradients in the laser laboratory [55]. The Zernike coefficient associated with defocus,  $Z_3$ , was therefore measured by the Shack-Hartmann wavefront sensor as close as possible to each shot being fired on the experiment detailed in chapter 5, so that the shift in focus position from best focus could be calculated [54] as,

$$\Delta_{\text{focus}} = \frac{4Z_3\lambda_L}{\text{NA}^2} \quad (3.2)$$

where NA is the numerical aperture of the OAP and is related to the  $f$ -number as  $\text{NA} = 1/2F_{\#}$ . The focal spot size was measured by the focal spot camera at best focus, and is therefore increased due to this shift in focus, now being calculated by substituting  $\Delta_{\text{focus}}$  for  $X$  in equation 2.1,

$$W_{0\Delta_{\text{focus}}} = W_0 \sqrt{1 + \left(\frac{\Delta_{\text{focus}}}{Z_R}\right)^2} \quad (3.3)$$

Multiplying by  $\sqrt{2\ln(2)}$ , the true focal spot size can be expressed in FWHM terms. For the example focal spot shown in figure 3.3, the defocus distortion measured before the interaction was  $Z_3=0.08\lambda_L$ . Accounting for this distortion, the focal spot size increases from  $\phi_L=4.7 \mu\text{m}$ , to  $\phi_L=4.9 \mu\text{m}$ . This is a relatively minor distortion, and is close to the average across the experimental data presented in chapter 5, though much larger values of up to  $Z_3=0.55\lambda_L$  were recorded during this experiment, substantially increasing the focal spot size.

Now that the entire beam path to target is understood, and the methods used to measure the laser energy, pulse duration, and focal spot size have been described, the simple calculation of optical intensity presented in equation 2.2 can be expanded to determine the peak optical intensity of the Vulcan-PW pulsed beam as,

$$I_L = \frac{E_{\text{precomp}} \times \text{comp}_{\%} \times \text{PM}_{\%} \times E_{\text{E}\%}}{\tau_L} \frac{1}{\pi \times (\phi_L/2)^2} \quad (3.4)$$

where  $\text{comp}_{\%}$  and  $\text{PM}_{\%}$  are the compressor and PM energy throughput efficiency, respectively. For the experimental data presented in chapter 5, the peak optical intensity ranges from  $I_L=(2-9)\times 10^{20} \text{ Wcm}^{-2}$  due to fluctuations in the energy, pulse duration and focal spot size. Variations in each of these parameters can significantly change the dynamics of a laser-solid interaction, influencing the heating and propagation of plasma electrons, and thus properties of accelerated ions, as described throughout chapter 2. A number of diagnostics are utilised to measure changes to these particle properties as laser parameters and other interaction parameters vary throughout an experiment, and will now be described.

### 3.4 Experimental diagnostics

Before beginning, it is important to note that many diagnostics are utilised to diagnose experimental laser-solid interactions. However, the investigation of chap-

ter 4 is simulation based, and the investigation of chapter 5 only presents experimental results from one diagnostic. A detailed description of the many diagnostics used in experiments is not therefore required to understand the results presented in this thesis, though several were used on the experiment detailed in chapter 5, and more were used on experiments undertaken alongside the work presented in this thesis. It would therefore be misleading to only discuss one experimental diagnostic, and so this section briefly covers those diagnostics which are commonly utilised on laser-solid interaction experiments, beginning by illustrating them as they are commonly placed around a target foil in figure 3.4.

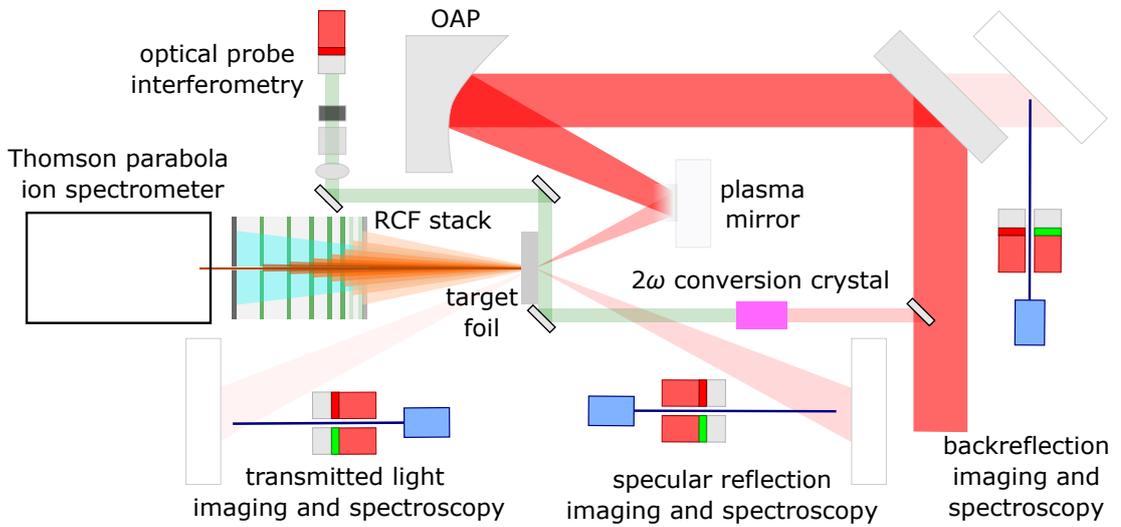


Figure 3.4: Suite of diagnostics typically employed to characterise laser light, pre-plasma expansion, and high energy particles produced during a laser-solid interaction.

### 3.4.1 Optical diagnostics

The first diagnostics to cover are those that measure properties of the laser. Section 2.6 discusses the importance of knowing the fraction of laser energy reflected, absorbed into plasma electrons, and transmitted through the target foil, since this defines whether ion acceleration occurs predominantly due to RPA, TNSA, or both, in an RSIT-enhanced hybrid mechanism. As such, these properties are diagnosed on most experiments involving thin target foils. As shown on figure 3.4, light which is specularly reflected from a target foil can be collected using a PTFE material which acts as a Lambertian surface, diffusely scattering light

in all directions. The light incident on this screen can then be captured by a telephoto lenses placed at any viewing angle, since the light intensity is the same in all directions, and focused onto a camera fitted with interference or bandpass filters that only transmit light at the fundamental laser wavelength. The spatial profile of the reflected beam can then be characterised, and by taking calibration shots where a beam of known energy is directed onto the screen, so too can the fraction of light reflected from a target foil. Typically, a second, almost identical, imaging system is fitted with filters so that it detects specularly reflected light at harmonic wavelengths, generated in the interaction region where the laser intensity is highest. By pointing the end of an optical fibre towards this screen, and relaying the specularly reflected light to an optical spectrometer, distinct spectral information is detected for many wavelengths.

Another diffuse scatter screen is commonly placed behind an optic in the beamline to measure a leak of the light reflected back along the laser incident axis, and a third screen can be placed after the target foil to diagnose transmitted light. This last screen is a crucial diagnostic in experiments involving investigation of RSIT dependent radiation production. Ion acceleration, for example, depends strongly on the time at which a target foil undergoes RSIT [71], which can be inferred by measuring the fraction of laser energy transmitted through the target, after calibrating by taking shots with no target. Experiments involving precise measurements of the RSIT onset time will diagnose the transmitted light using techniques such as spectral interferometry [240]. Another active area of research in the RSIT regime is that of plasma optics, where targets that undergo RSIT will alter the incident laser polarisation and generate light of higher order modes, allowing typically beam-like multi-MeV electrons and protons accelerated during an interaction to be manipulated into these higher order structures [69, 70]. A device known as a Stokes polarimeter is used on these experiments to characterise the transmitted laser polarisation, and spatial profile, to diagnose changes to the incident polarisation, and the generation of higher order modes [68, 140].

By measuring light specularly reflected from a target foil, as well as that transmitted, the fraction of laser energy absorbed by plasma electrons can be inferred.

This is a crucial parameter in defining radiation produced during laser-solid interactions as described in section 2.4.3, and investigated in detail to optimise ion acceleration in chapter 4. Some light is also scattered in other directions during an interaction, and will therefore not be captured by these diagnostics. As such, experimental investigations that aim to accurately measure the total fraction of laser energy absorbed will surround the target foil with an integrating sphere [183, 241] coated in this diffusely scattering material. All light not absorbed by plasma electrons can then be measured using optical fibres placed in small holes in the sphere, which relay the light to spectrometers as mentioned [50].

Many other optical diagnostics are used in laser-solid interactions, with optical probes being one of the most common. The simplest form of this diagnostic involves propagating a small,  $r \sim 1$  mm beam, transversely across the front surface of a target foil, with this beam either being sub-sampled from the main collimated beam before its focused, or originating from a separate laser. A translation stage with moveable mirrors is used to increase or decrease the distance travelled by this probe beam before it reaches the target foil, so that it passes across the target at a chosen time relative to the arrival of the main pulse. Typically this is  $t=0$  fs relative to the main pulse, so that the density gradient of the pre-plasma formed by ASE light can be characterised. This characterisation often involves Nomarski interferometry, where probe light is imaged by a lens and separated into two beams of equal intensity, but orthogonal polarisation, by a Wollaston prism [242].

Typically, the separation angle is  $\sim 3^\circ$ , so that the beams are mostly separated, but slightly overlap when imaged by a camera. Light which was propagated across the target foil, and the laser-induced pre-plasma on the front surface, from the bottom of the first image, then overlaps with light which has propagated in vacuum, from the top of the second image. By placing a polariser in front of the camera, the two beams are modulated to have the same polarisation, so that they interfere in the overlapping region, creating fringes. Light from the first beam which has propagated through vacuum will produce straight fringes of fixed separation, but that which has propagated through underdense plasma

will experience a phase shift compared to light from the second beam which has propagated in vacuum, bending fringes in this region. This phase shift can then be used to calculate the refractive index, and thus the electron density of the underdense plasma through which the light propagated before being detected [242, 243].

A common difficulty with this technique is that light scattered and generated from underdense plasma is much more intense than the probe light, and is collected by the lens, essentially blinding the diagnostic to probe light. To avoid this, probe light is often frequency doubled, and filters are placed in front of the camera so that only light at this wavelength is captured, with light at the laser fundamental wavelength being attenuated or reflected. This is also useful because frequency doubled light propagates through denser plasma regions, according to equation 2.28, providing more information on the plasma expansion. More complicated probe designs split light into several components of different wavelengths and polarisations, then using separate timing slides to induce different path lengths, so that the beams probe the foil target front surface at distinct times, providing information on the temporal evolution of the pre-plasma density profile [49, 244]. Optical probes have also been propagated across the rear surface of a target foil, to measure the density of escaping fast electrons [48].

Now that optical diagnostics commonly used in laser-solid interactions have been described, the next diagnostics to consider are those which measure properties of particles accelerated during these interactions.

### 3.4.2 Particle diagnostics

Investigations presented in chapters 4 and 5 focus on improving properties of ions accelerated during laser-solid interactions. Diagnostics used to characterise these properties will be described presently, along with those used to diagnose properties of fast electrons, which facilitate ion acceleration as described in chapter 2.

## Thomson parabola ion spectrometer

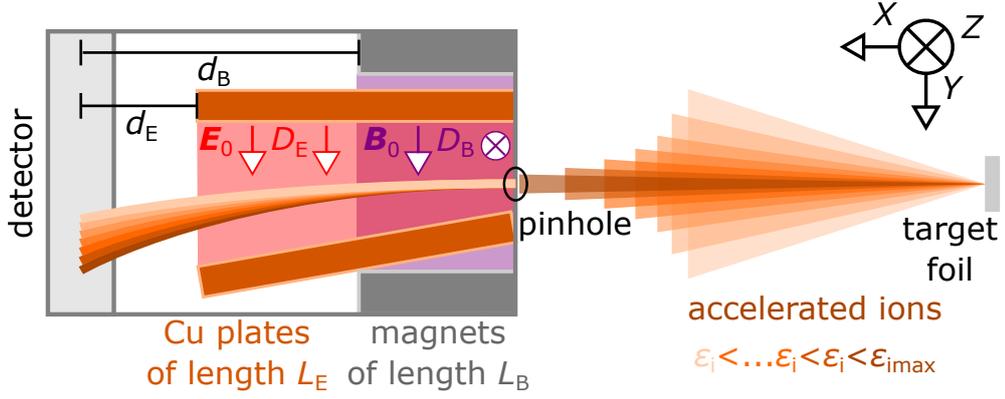


Figure 3.5: Operating principle of a Thomson parabola ion spectrometer diagnostic

The first diagnostic to consider is a Thomson parabola spectrometer (TPS), which separates ions based on their energy, and charge to mass ratio,  $q_i/m_i$ , using a design [245] such as that illustrated in figure 3.5. Here, ions pass through two permanent magnets of length,  $L_B$ , and are deflected by the associated magnetic field,  $\mathbf{B}_0$ , according to their mass and energy, as described by equation 2.16. In the coordinate system of figure 3.5, ions propagating in the  $X$  dimension are then displaced by a distance,

$$D_B(Z) = \frac{q_i L_B \mathbf{B}_0(Y)}{m_i \mathbf{v}_i(X)} (0.5 L_B + d_B) \quad (3.5)$$

in the  $Z$  dimension, relative to unperturbed propagation, over a distance,  $d_B$ , from the end of the magnets to a detector. Simultaneously, a pair of typically copper plates are supplied with a voltage to induce an electric field, deflecting ions according to their charge, mass, and energy, in the  $Y$  dimension relative to the coordinate system of figure 3.5, with a stronger, velocity squared, dependence,

$$D_E(Y) = \frac{q_i L_E \mathbf{E}_0(Y)}{m_i \mathbf{v}_i^2(X)} (0.5 L_E + d_E) \quad (3.6)$$

A pinhole at the front of the spectrometer ensures only ions within a small angle are accepted, reducing flux to avoid saturating a detector, and improving the energy resolution of the diagnostic. Different detectors are used with this diagnostic, one of which is a multi-channel plate (MCP), which converts incident ions into

many electrons, that are subsequently converted to photons by a phosphorescent screen. A lens collects these photons and relays them to a camera, creating an image with multiple so-called tracks of different ion species. Separation of these tracks determines which species are present, with the length of individual tracks corresponding to the energy range of a single ion species.

Another detector often used with a TPS is image plate (IP), which also converts ionising radiation to photons through a phosphorous material, but through a different process. Here, energetic particles incident on the IP ionise the phosphor layer, exciting electrons into the phosphor conduction band, where they stay for up to a few hours, due to the lattice structure of the material. This allows the IP to be removed from the vacuum chamber, and scanned by a laser, which stimulates emission of these conduction band electrons, resulting in emission of photons into a photomultiplier tube, which amplifies the signal and converts it to a complete image using an analogue to digital converter. Using a piece of IP which has previously been irradiated by a source of particles with known energy, the signal on this image can be used to calculate the energy and number of incident particles. The disadvantage of this technique is that IP must be removed from the chamber to be scanned in this way, limiting the number of shots which can be taken, in contrast to using an MCP which remains in the chamber for the entire experiment. The advantage however, is that IP provides much higher spatial, and therefore, energy resolution.

The final detector which has historically been used with a TPS is California-Resin 39 (CR-39), which is a polymer based material that damages according to individual ions propagating some distance into the surface. After an interaction, CR-39 is removed from the experimental chamber, and the surface is etched away by a sodium-hydroxide solution, revealing individual pits created by individual ions. These pits are then counted along each track corresponding to the different ion species, using a microscope and automated detection software, resulting in an energy spectrum for each species. There are several disadvantages to this approach, in that the maximum ion energy that can be resolved is limited by the material thickness, that pits overlap if the ion flux is very high, and can no

longer be individually resolved, and that etching the material and counting the pits can be very time consuming. For these reasons, a TPS will typically be fitted with an MCP or IP to diagnose ions, though CR-39 can still be useful to absolutely calibrate the diagnostic. Regardless of which detector is used, a TPS cannot characterise the spatial profile of an ion beam. This information can be very useful, and is therefore often characterised, alongside the energy and number of ions, using a stack of film which is sensitive to ionising radiation, placed a few centimetres behind a target foil. This diagnostic was used to measure the energy of protons accelerated from laser-solid interactions for the investigation presented in chapter 5, and is now described.

### Radiochromic film stack

Radiochromic film (RCF) contains layers of plastic sandwiching an organic material which discolours when exposed to ionising radiation, altering the optical density (OD) of this layer in proportion to the radiation dose. A piece of RCF can then be placed in an optical scanner, which will measure the intensity of incident light,  $I_{\text{incident}}$ , which is transmitted,  $I_{\text{transmitted}}$ , through the altered optical density,

$$\text{OD} = -\log_{10}\left(\frac{I_{\text{incident}}}{I_{\text{transmitted}}}\right) \quad (3.7)$$

By first applying this technique to RCF exposed to radiation of known energy, the optical density can be converted to the energy of incident radiation. This technique is particularly useful for detecting protons, since they deposit most of their energy at a specific thickness in a material, according to the Bragg peak [246, 247], in contrast to electrons and photons which deposit energy much more broadly. By stacking many layers of RCF together, protons over a large range of energies can be detected, as shown in figure 3.6.

By interspersing layers of RCF with mylar and iron filters of known thickness, the energy resolution of the stack can be tuned so that different layers of RCF will detect protons of a specific energy. An aluminium filter is also typically added to

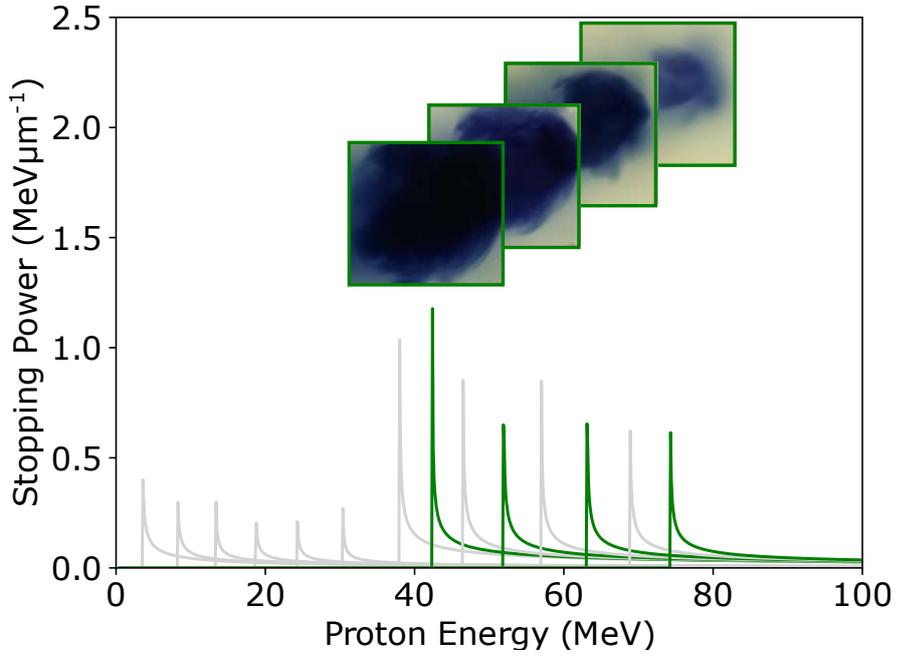


Figure 3.6: Energy dependent stopping of protons in an RCF stack. Four layers from the stack show the proton beam size decreasing at higher energies.

the front of the stack, to protect the RCF from damage due to any transmitted laser light, but also to stop heavier ions, so that only protons are detected by the RCF layers [248]. Figure 3.6 also shows four layers of RCF exposed to protons that were accelerated to energies of  $\sim(40-75)$  MeV in the experiment presented in the investigation of chapter 5. Each layer has an active area of several centimetres, and can therefore capture spatial information of the proton beam. As discussed in section 2.6.1, and illustrated in figure 3.4, a TNSA proton beam is broader at lower energies, since higher energy protons are produced at the peak electric sheath field. This effect is clearly observed in the RCF layers shown, where the proton beam is initially quite large, and has a high dose at energies of  $\epsilon_p \sim 40$  MeV, but is much smaller, with less dose, at  $\epsilon_p \sim 75$  MeV. The stack design used in this experiment did include more RCF layers after the final layer shown on figure 3.6, though this was the last layer on which protons were detected, and so  $\epsilon_{p\max} \sim 75$  MeV was determined as the maximum energy to which protons were accelerated from this particular laser-solid interaction.

The spatial information provided by an RCF stack is extremely useful in diagnosing the physics involved in an interaction, particularly in those where protons

are accelerated by multiple mechanisms. By irradiating the target foil at an angle, the target normal axis and laser axis are offset, so that TNSA and RPA protons will separate over the several centimetres between the rear surface of the target foil and the RCF stack, allowing the mechanism responsible for the highest energy, or highest number of protons to be inferred. Structures in the proton spatial profile can similarly provide valuable information about the acceleration mechanism, with protons accelerated in an RSIT-enhanced mechanism often exhibiting ring like patterns, rather than the more uniform beam shown in figure 3.6. Similarly, early experiments determined that rear surface TNSA was primarily responsible for the highest energy protons in a laser-solid interaction, rather than front surface acceleration, by implementing structures on the rear surface, which were then detected on layers of the stack corresponding to the highest proton energies [32]. Spatial information provided by an RCF stack thus gives it an advantage over a TPS, though it has a disadvantage in that it limits the data acquisition rate, since RCF must be removed from the chamber and replaced either on every laser shot, or every few shots, if several stacks are placed on a wheel and rotated into place behind the target on consecutive shots. As such, scintillator-based diagnostics have been used to measure the energy and spatial profile in a similar fashion to RCF [94, 95], but by using a camera to image photons emitted from proton induced fluorescence of the material, allowing operation at much high repetition rates, required to diagnose radiation produced in experiments using high-power lasers which operate at  $\gtrsim 1$  Hz [80–84]. Electron and x-ray diagnostics are similarly being developed to operate at this previously unattainable repetition rate for high-power lasers [97, 249].

Now that the methods used to diagnose experimental laser-solid interactions have been discussed, the final important methodology used throughout this thesis is that associated with numerical simulations. This will now be discussed.

## 3.5 Particle-in-cell simulations

In the previous section, diagnostics used to measure important properties of laser-solid interactions in experiments were discussed. These diagnostics provide vital information, though they are limited in the information which they can provide. In particular, measurements made by these diagnostics are almost always temporally integrated, meaning they cannot detect how important features of the fast electron and proton population evolve over time. These, and other interaction dynamics, change over few to tens of femtoseconds in interactions, and significantly influence source properties at the end of the entire interaction [65, 71, 250].

To overcome this limitation, numerical simulations are typically performed after an experiment to approximate the interaction conditions, and to diagnose important features at tens or hundreds of femtosecond increments throughout the interaction. This is the case in chapter 5, which uses simulations to comprehensively investigate the susceptibility of RSIT-enhanced proton acceleration to fluctuating laser pulse parameters, something which was only possible to a limited extent on experiment, given the low repetition rate of the laser. Furthermore, simulations are useful in deciding which parameters should be varied on experiment, by identifying interesting interaction conditions ahead of time. Simulations are therefore crucial in investigating laser-solid interactions, though they too have limitations which should be understood. Several numerical simulation approaches are utilised to investigate laser-solid, and more widely, laser-plasma interactions [251–253], though this thesis exclusively uses a PIC code known as EPOCH [171]. This section will therefore describe the general operation of PIC codes, other implementations of which exist in addition to EPOCH [254, 255].

Completely modelling an experimental laser-solid interaction would require a full accounting of the dynamics of every particle involved, and the EM fields they induce. This would require coupled equations of motion to be solved for each of these particles [256], though this is completely unfeasible given the computational resource and time required to do so, even using modern supercomputers. Fortunately, plasma dynamics are well described by considering many particles within

a given region as a collective entity characterised by a distribution function which describes the number, position, and velocity of particles within this region, at a given time, as described in sections 2.3 and 2.4. This so-called kinetic description of a plasma can approximate dynamics observed in experiments very accurately, with less computational resource than would be required to treat particles individually, though solving the Vlasov equation which underpins this description still requires significant computational resource, since it includes three position ( $X, Y, Z$ ), and three velocity co-ordinates ( $\mathbf{v}_X, \mathbf{v}_Y, \mathbf{v}_Z$ ).

PIC codes overcome this issue by mapping the distribution function for these particles into a single so-called macroparticle [257] with the same charge to mass ratio as each of the individual particles being represented. Macroparticles therefore move according to the Lorentz force as described for single particles in section 2.4, with their associated current density generating electric and magnetic fields according to equations 2.37 and 2.38. A laser-solid interaction can then be simulated by placing macroparticles into individual cells on a large grid, each with an associated density, temperature, and energy. Macroparticles are also weighted to characterise the number of individual particles each represents. This enables more computational resource to be applied to grid regions with high macroparticle weights, and less to regions of lower weights, such as vacuum.

This mapping of macroparticles to individual cells is illustrated in figure 3.7, and is the first step of PIC codes, which proceed by using the flux of particles in each cell to determine an associated current density. This is used to calculate the electric and magnetic fields generated due to the motion of each macroparticle using a finite-difference time-domain (FDTD) technique, so that both fields can be calculated between simulation timesteps,

$$\frac{\mathbf{E}_{n+1/2} - \mathbf{E}_n}{\Delta t/2} = c^2 \nabla \times \mathbf{B}_n - \frac{\mathbf{j}_n}{\epsilon_0} \quad (3.8)$$

$$\frac{\mathbf{B}_{n+1/2} - \mathbf{B}_n}{\Delta t/2} = -\nabla \times \mathbf{E}_{n+1/2} \quad (3.9)$$

Here,  $\mathbf{E}_{n+1/2}$  and  $\mathbf{B}_{n+1/2}$  are the electric and magnetic fields calculated between

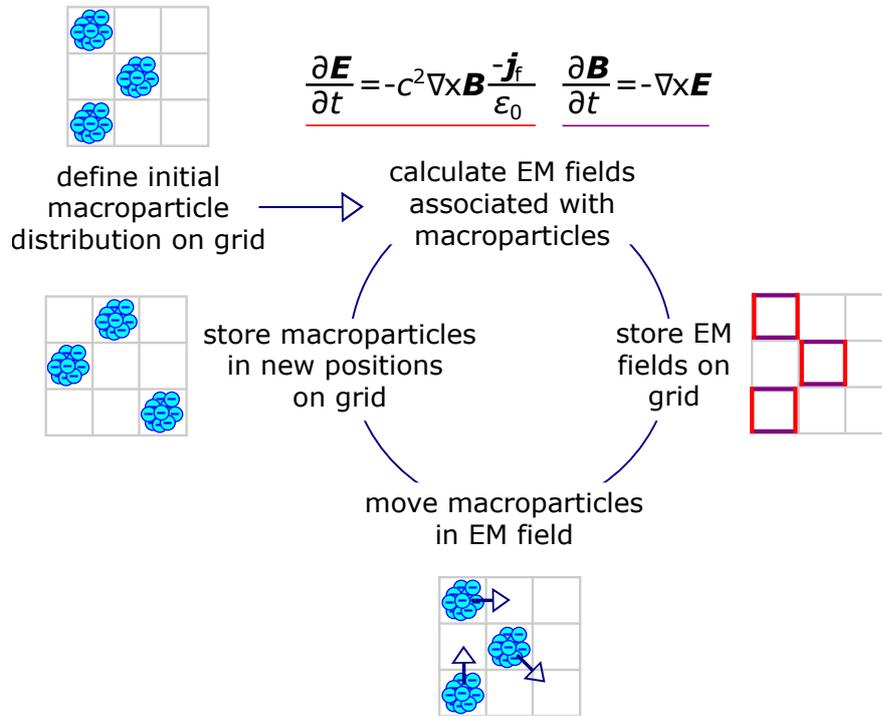


Figure 3.7: Diagram illustrating the PIC simulation method. Ten macroparticles (cyan), each representing the position and velocity of many individual particles, are placed in cells on a grid, and induce electric and magnetic fields which are stored on the edges and vertices of each cell. Macroparticles are then moved by these fields, taking up new positions on the grid. This process repeats for as many simulation timesteps as specified.

simulation timesteps,  $t_n$  and  $t_{n+1}$ , due to the motion, and current density,  $\mathbf{j}_n$ , associated with particles at  $t_n$ . These fields are then stored on the vertices and edges of cells in a Yee grid-like implementation [258] as illustrated in figure 3.7, and are substituted back into the Lorentz force equation to move macroparticles to new positions at the next simulation timestep,  $t_{n+1}$ . The macroparticle velocity in each cell is centred using a Boris algorithm [259], and the particle flux induces a current density at this timestep, inducing EM fields which are solved on the next half timestep of the code  $t_{n+3/2}$ , and move macroparticles to new positions at the next timestep,  $t_{n+2}$ , resulting in a circular process which repeats for each simulation timestep.

Three dimensional PIC simulations are required to model experimental interactions as accurately as possible, by fully accounting for the 3D nature of the laser pulse, target foil, plasma expansion, and particle acceleration. However, running these simulations takes significant computational time, on many CPUs

or even GPUs, limiting the number which can be performed, or even prohibiting the use of 3D PIC simulations altogether. As such, PIC simulations are often performed in 2D or even 1D, to provide a good, but less complete, approximation of experimental interactions, using considerably less resource than equivalent 3D simulations. 2D PIC simulations are therefore used throughout this thesis, each with laser and target parameter values that were chosen to approximate realistic interaction conditions as accurately as possible in a reasonable time frame, given the computational resources available. These parameter values are described in detail for each simulation presented, as are values chosen for parameters of the simulation grid and macroparticles themselves, which are equally important. A third order B-spline macroparticle shape is used in each 2D simulation, meaning the position of each macroparticle is interpolated over three cells in the  $X$  and  $Y$  dimension, as are the induced electric and magnetic fields, reducing statistical noise compared to mapping macroparticles to the closest individual cell.

Grid and macroparticle parameter values are typically selected by performing test simulations, features of which are monitored and improved if required, to ensure a physically accurate simulation, in a so-called convergence testing process. The major parameters considered are the grid size, the number of cells on the grid, and the number of macroparticles per cell, to ensure realistic plasma behaviour, and that particles don't leave the simulation box prematurely. To explain this, it is useful to discuss the values chosen for the 2D PIC simulation employed to demonstrate TNSA in figure 2.11. The grid size for this simulation was selected so that the laser pulse could propagate in pre-plasma at the front surface of the target foil, rather than reflecting from the critical density surface close to the edge of the grid, and so that high energy particles accelerated from the target foil rear surface did not leave the grid before their acceleration had saturated. The grid dimensions chosen were  $X=(-20\rightarrow20)$   $\mu\text{m}$ , and  $Y=(-20\rightarrow20)$   $\mu\text{m}$ , and there were 4000, and 3456 cells on the grid, so that each cell has a resolution of 10 nm/cell and 12 nm/cell, in the  $X$  and  $Y$  dimensions, respectively. The initial temperature and density of electron macroparticles was selected as  $n_e\sim 300n_{\text{crit}}=5\times 10^{29}$   $\text{m}^{-3}$ , and  $k_{\text{B}}T_e=1$  keV, to approximate a CH target foil, and its initially non-relativistic

constituent electrons, which are characterised by a Debye length,  $\lambda_D=0.3$  nm, in each dimension, according to equation 2.9.

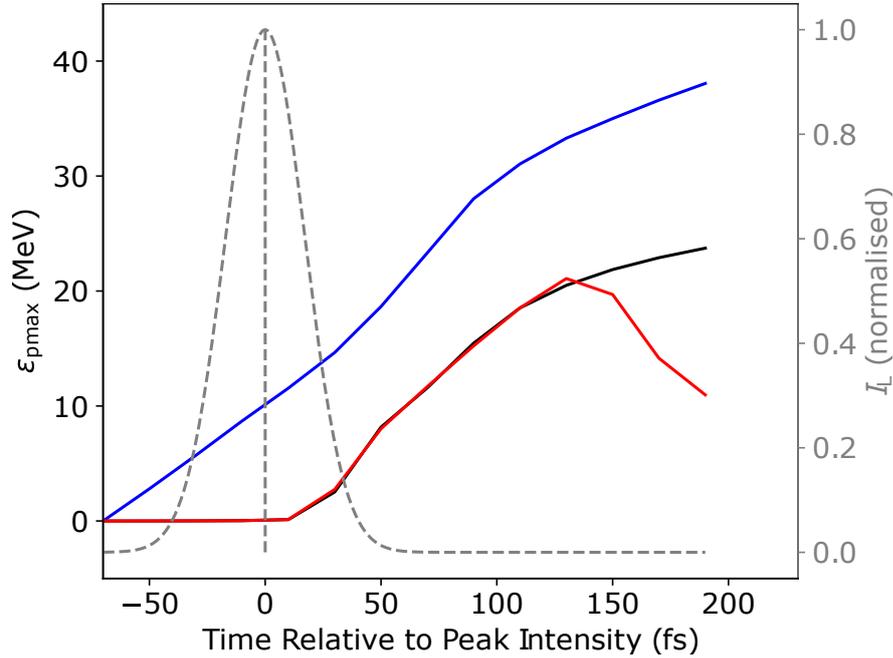


Figure 3.8: Maximum energy of protons accelerated during a laser-solid interaction, in 2D PIC simulations where the cell resolution is sufficient/insufficient to approximate physical plasma behaviour (black/blue), and where the simulation box is too small so that protons leave before their acceleration has saturated, meaning their final energy will be underestimated (red).

Setting the cell resolution to be comparable to the Debye length ensures realistic plasma behaviour, by ensuring fields associated with macroparticles only very weakly influence macroparticles at a distance greater than the Debye length, as described in section 2.3. If this is not the case, electrons will experience nonphysical oscillation and heating. In context of the investigations presented in chapters 4 and 5, this means many more electrons will gain energies greater than their rest mass energy than is physically correct, resulting in the acceleration of protons to significantly higher energies than would be expected in a realistic interaction. This is shown in figure 3.8, where the maximum energy of protons is shown as a function of time for the simulation used to demonstrate TNSA in figure 2.11, and for a second simulation with identical parameters, but for the number of cells which has been reduced by  $4\times$  in both dimensions, worsening the resolution by  $4\times$ , to 40 nm/cell and 48 nm/cell, respectively. In the second simulation, moderate laser intensities induce nonphysical electron heating, resulting in proton

acceleration to multi-MeV energies well before the peak laser intensity interacts with the target foil, unlike in the well-resolved simulation. This nonphysical electron heating will also occur if the number of macroparticles per cell is too low. Also shown on figure 3.8 is a third simulation with the same parameters as the first, well-resolved simulation, but for the grid  $X$  dimension, which was reduced to  $X=(-20 \rightarrow 7)$   $\mu\text{m}$ . In this case the simulation grid is too small so that the highest energy protons leave before their acceleration has saturated, resulting in a sharp decrease in maximum proton energy. This highlights the importance of correctly specifying not only laser and target parameters, but also parameters of the simulation itself, to ensure realistic interaction conditions are approximated.

### 3.6 Summary

This chapter has described methods used to generate high-power laser pulses, and to characterise their energy and duration, as well as the focal spot size which can be achieved by focusing these pulses with an OAP. Furthermore, deviations from an idealised laser temporal-intensity profile and spatial-intensity profile were discussed, and methods to mitigate these deviations were described in the form of plasma mirrors, and a coupled adaptive optic and Shack-Hartmann wavefront sensor system. Diagnostics used to characterise important features of experimental laser-solid interactions, such as those employed in the experimental investigation presented in chapter 5, were then discussed. So too were methods underpinning PIC simulations, and the considerations which must be made for these simulations to provide accurate insights into experimental interaction conditions, as investigated in chapters 4 and 5.

# Chapter 4

## Multi-parameter Bayesian optimisation of laser-driven ion acceleration in particle-in-cell simulations

### 4.1 Introduction

Chapter 1 introduced the potential to utilise the unique properties of ions accelerated during laser-solid interactions in important physics research and societal challenges. For this to be realised, these properties must be improved beyond what has currently been achieved, with the energy of laser-driven protons being discussed as a key property which must be improved for applications such as radiotherapy [37, 155]. Progress towards this goal has historically come through optimising the target normal sheath acceleration (TNSA) mechanism [33], resulting in an increase of maximum achievable proton energies from  $\sim 58$  MeV in 2000 [27–29] to  $\sim 90$  MeV today [62, 63]. This has resulted from investigations which have sought to understand the dependence of proton maximum energies on interaction parameters such as the foil target thickness [38, 39], the laser energy, pulse duration [40, 41], focal spot size [42], and temporal-intensity contrast [43–45], the properties of beams of fast electrons generated and transported within plasma

[46–50], and the accelerating field they induce [40, 51]. However, optimal values presented for these parameters in the literature are specific to the laser system and experimental set-up used. Laser parameter values vary unpredictably, and are often co-dependent, meaning setting one parameter to its optimal value may change the optimal value of another. As an example, the optimal target foil thickness for TNSA is dependent on the laser temporal-intensity contrast [43], and is thus dependent on the laser system being used [45]. Investigations seeking to optimise just one or two parameters whilst keeping all others constant are therefore not likely to result in the absolute maximum proton energies which can be achieved for a given laser system. Ideally, maximum proton energies and other important beam properties would be optimised for all laser and target parameters upon which they are known to depend. This is challenging, given humans cannot easily interpret functions in more than three dimensions, and demands a solution that can replicate and optimise multi-dimensional functions.

As introduced in chapter 1, machine learning (ML) techniques are well suited to such challenges, and have therefore been applied in a number of experimental and numerical simulation-based studies of laser-plasma interactions [129, 133–135, 215, 260–263]. Various ML algorithms have been used in these studies, and there are advantages and disadvantages associated with each, making it important to select an appropriate algorithm for the task at hand. Bayesian optimisation [264] is a ML technique that is particularly useful for optimising in regimes where limited data is available and where the outputs to be optimised are subject to uncertainty. This is the case in experiments using low repetition rate lasers, where the overall shot numbers are low and are often subject to shot-to-shot fluctuations, as demonstrated in chapter 5. It is also the case in numerical investigations involving high-fidelity particle-in-cell (PIC) simulations, where the number of simulations that can be run is limited by the large computational and storage resources required. By contrast, ML approaches based on neural networks [215] and genetic algorithms [133] require large quantities of data to make accurate predictions, and direct optimisation algorithms can struggle to optimise noisy functions. Owing to these advantages, Bayesian optimisation has been applied

to laser-driven electron [134, 260, 261] and ion [135] acceleration experiments. Bayesian inference has also been used to model laser-driven ion acceleration [214].

In this chapter, Bayesian optimisation is used to optimise the maximum proton energy achieved via TNSA in PIC simulations. Using this technique, the maximum proton energy was optimised with up to  $200\times$  fewer simulations than would have been required to optimise by linearly varying key laser and target parameters in a conventional grid-search approach. This corresponds to a reduction of  $\sim 48$  days in simulation time, meaning it would have been unfeasible to optimise the maximum proton energy via conventional grid-search for the parameter space covered. As such, the work presented in this chapter demonstrates the potential to apply Bayesian optimisation to laser-driven ion acceleration in situations where limited data is available and where multi-parameter optimisation would not otherwise be possible.

The chapter will also report on the development of a code called BISHOP, which was undertaken as part of the work of this thesis. BISHOP automatically generates simulation inputs, submitting them to a remote high-performance computing (HPC) cluster, and analysing the resultant data. Using BISHOP significantly reduces the time required to produce and prepare the simulation data and enables the automated feedback loop inherent to the Bayesian optimisation process. As such, BISHOP represents a newly developed tool that can be used to generate large scale data-sets of simulation results and facilitate the application of various ML algorithms in future studies. The results presented in this chapter have been published in *New Journal of Physics* [100].

## 4.2 Development and operation of the BISHOP Code

Research within the field of laser-solid interactions is rapidly transitioning to studies of large data-sets facilitated by high repetition rate lasers [80–84], targets [85–88, 90, 91], and online diagnostics [92–96], which can be analysed using statistical and ML-based techniques [129]. This will undoubtedly lead to tech-

nological advances and new physics insights as discussed in chapter 1. However, for this type of research to become widespread, existing infrastructure must be developed, and new infrastructures must be created to facilitate data generation and handling on such large scales [102–104]. A good example is in handling data generation from numerical simulation codes such as EPOCH [171].

Here, researchers are required to submit an ‘input.deck’ containing a list of variables which determine features of the simulation such as the size of the mesh grid, the run time, and the composition of the incident laser and target plasma. Researchers typically vary simulation parameters manually, which works well when generating tens of simulations, but quickly becomes inefficient when generating hundreds or thousands of simulations, as is required for the application of statistical analysis and ML. At this scale, manually updating simulation variables becomes extremely time consuming and prone to human errors.

These issues aside, ML techniques such as Bayesian optimisation self-generate new parameter values in an iterative process. In this case, the researcher would have to analyse simulation(s) at each iteration of the ML process, and manually submit input.deck(s) for the next iteration. This would require a researcher to be on hand 24/7, lest they slow down the data generation process by failing to play their part in the cycle for a few hours, diminishing one of the major benefits of using ML. Automating the process of generating numerical simulation data would provide an opportunity to take advantage of these ML techniques and allow much larger simulation-based data-sets to be studied. To this end a code called ‘BISHOP’ has been developed in python, and has been used to generate thousands of PIC simulations to investigate the optimisation of maximum proton energy in the TNSA regime.

BISHOP automates the process of generating numerical simulation data so that large scale data-sets can be produced quickly, and with minimal user input. Before initiating BISHOP, users must provide four files which determine how it will generate simulations. A summary of these files, and some of the details contained within them, is illustrated in figure 4.1. First, users provide a base ‘input.deck’ and ‘runscript.txt’ for the PIC code, as they would when running

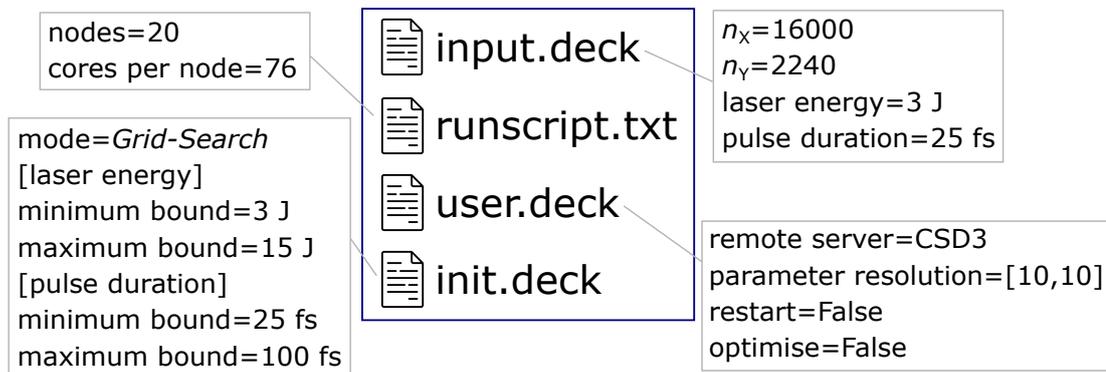


Figure 4.1: Files used by BISHOP to automatically generate simulation `input.decks` are outlined in blue. The `input.deck` file contains simulation conditions to be kept constant, and to be varied. The `runscript.txt` file contains details such as the number of nodes, and CPU cores, used to run simulations on a remote computer, itself defined in the `user.deck` file along with details such as the number of increments by which to vary each parameter, whether BISHOP should restart when a simulation fails, and whether Bayesian optimisation should be used to select simulation inputs between bounds specified in the `init.deck` (`optimise=True`), or whether this should be done by another operating mode (`optimise=False`), in this case *Grid-Search*.

a simulation manually. A ‘`user.deck`’ should then be provided to define whether the simulations will run on a local computer (`remote server=False` on figure 4.1) or on a remote system such as the Cambridge Service for Data-Driven Discovery (CSD3) [265]. For the latter, login details should also be provided so that BISHOP can connect to the remote system via Secure Shell Protocol (SSH).

Also contained in the `user.deck` are a number of other variables that are important to BISHOP’s operation. This includes the number of simulation output files to be generated and information on whether to restart the sequence of simulations in the case of a simulation failing. Finally, the user will create an ‘`init.deck`’ file which determines the parameters that BISHOP should vary, and the mode by which they should be varied from those detailed in table 4.1.

When operating BISHOP in the *Grid-Search* or *Random* modes described in table 4.1, the user should include the laser/target parameter name(s) and the respective minimum and maximum values that each parameter is permitted to take in the `init.deck`. An example of this is shown in figure 4.1 where the laser energy and pulse duration are to be varied by *Grid-Search*. When BISHOP is executed, some number of `input.decks` are created where the parameters are varied in either linear increments, or randomly, between the specified minimum and maximum bounds. The number of `input.decks` generated depends on the

Operating Mode	Parameter Definition	Execution
<i>Grid-Search</i>	Min= $x_{\min}$ Max= $x_{\max}$	Parameter(s) are varied in linear increments between the defined minimum and maximum values
<i>Random</i>	Min= $x_{\min}$ Max= $x_{\max}$	Parameter(s) are randomly varied between the defined minimum and maximum values
<i>Vector</i>	Array= $[x_1, x_2, \dots x_n]$	Parameter(s) are set to $x_1, x_2$ and $\dots x_n$ for $n$ individual simulations
<i>Optimise</i>	Min= $x_{\min}$ Max= $x_{\max}$	Parameter(s) are set between the defined minimum and maximum values as decided by a Bayesian optimisation algorithm

Table 4.1: Operating modes of BISHOP.

number of parameters to be varied, and the parameter resolution specified in the user.deck. For the example illustrated in figure 4.1, two parameters are to be varied in ten linear increments between the respective minimum and maximum bound, making for a data-set of 100 simulations which have identical set-ups, but for these parameters. This was the set-up for the results detailed in section 4.6.1.

Having generated these input.decks, BISHOP then runs the first of these simulations ( $E_L=3$  J,  $\tau_L=25$  fs) using the EPOCH PIC code. BISHOP periodically checks whether this simulation has finished running and, once it has, moves the data from the remote computing cluster to the local directory from which the code was executed. At this point the next simulation in the sequence ( $E_L=4.3$  J,  $\tau_L=25$  fs) is run, and the process repeated until the data from all 100 simulations is stored in the local directory. The entire process is illustrated in figure 4.2. Importantly, it is also possible to have different resolutions for each parameter such that a data-set could be created where the laser energy and pulse duration are varied over 5 and 10 increments between the respective minimum and maximum bounds, making for a 50 simulation data-set.

The *Vector* mode detailed in table 4.1 operates slightly different in that it generates input.decks as specified explicitly by the user. As an example, one may wish to generate simulations in which the laser energy and pulse duration should be exactly ( $E_L=3.7$  J,  $\tau_L=25$  fs), ( $E_L=6.7$  J,  $\tau_L=35$  fs), ( $E_L=9.9$  J,  $\tau_L=90.5$  fs), in which case three input.decks would be generated containing these parame-

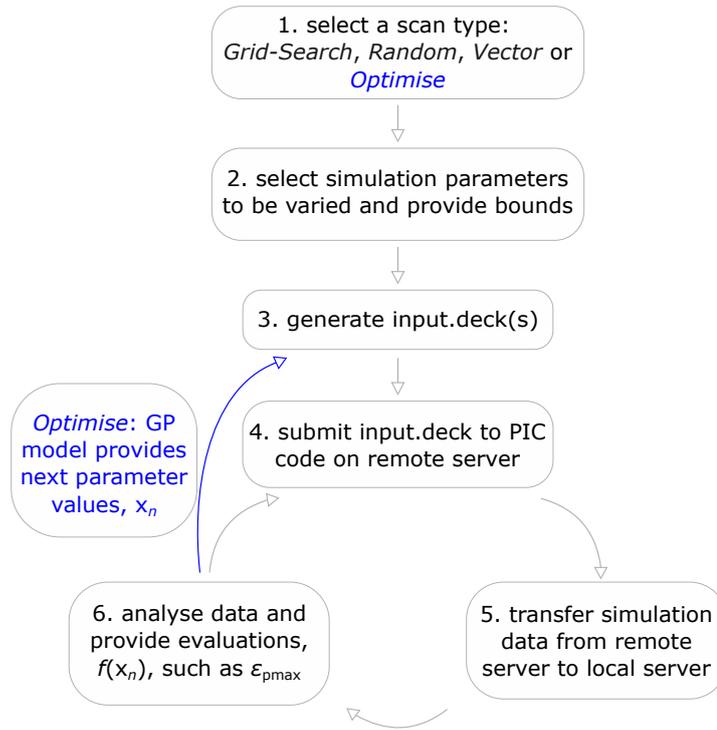


Figure 4.2: Operational steps of the BISHOP code, where simulation input.deck(s) are automatically generated and submitted to a PIC code on a HPC cluster according to the selected scan type.

ter values. BISHOP would then operate exactly as described for the previous modes, as illustrated in figure 4.2. The *Vector* mode was utilised to generate PIC simulations for the investigation presented in chapter 5.

As shown in this figure, BISHOP operates slightly differently in the *Optimise* mode, selected by setting `optimise=True` in the ‘user.deck’, as shown in figure 4.1. Here, an integrated Bayesian optimisation algorithm [266] guides the specified parameter values for a number of simulations using a Gaussian process regression model. This process is described in detail in section 4.5.

### 4.3 Selecting a machine learning-based optimisation technique

In recent years ML techniques have become accessible to non-experts as high-performance computing has become more readily available and codes have been made open source. This provides researchers with a toolbox of algorithms to

model complex multi-dimensional functions and assist with tasks such as regression, classification, optimisation and feature detection.

As with any toolbox, some tools are better suited to certain tasks than others, and it is important to know which algorithm is best suited to the task at hand. In the field of laser-driven ion acceleration, researchers are often interested in modelling or optimising a source property which fluctuates significantly from one interaction to another, using relatively small data-sets as limited by the low repetition rates of certain high-power lasers (see chapter 5), or the computational resource required to generate high-fidelity 2D PIC simulations, as done in this chapter. Bayesian optimisation is a ML technique that is expected to optimise efficiently under these limitations, and it has been applied to laser-driven ion [135] and electron [134, 260, 261] acceleration experiments as such. In each of these studies, source properties were automatically optimised with little to no human input. This makes Bayesian optimisation particularly interesting for applications, since it facilitates the possible use of laser-plasma accelerators by non experts, who have simply to turn on the accelerator and check that the output source property is as required for their use case.

However, before applying Bayesian optimisation in this chapter, it was important to first determine that it was indeed the best ML algorithm for optimising laser-driven ion acceleration in a data-limited simulation environment. To this end, the Bayesian optimisation algorithm from scikit-learn was applied to the so-called ‘branin’ function [267] alongside a variety of other optimisation algorithms included in this python library [266]. This function is commonly used for benchmarking the performance of optimisation algorithms since it contains multiple optima, making it difficult for algorithms to identify the global optimum. For this benchmarking process each algorithm was tasked with optimising the function 20 times and was permitted to evaluate the function a maximum of 150 times. The average convergence to the global optimum, and the number of iterations until 95% convergence to the global optimum, is shown for each case alongside the time taken to perform the 20 optimisations in table 4.2.

Gradient descent-based algorithms (Conjugate Gradient [268], Powell [269],

Method	Median Convergence to Optimum (%)	Median $n_{\text{iters}}$ until $\geq 95\%$ Convergence	Optimisation Time: Branin (s)	Optimisation Time: Simulations (hr)
Bayesian	$100 \pm 0$	$30 \pm 3$	3000	10
Powell	$100 \pm 10$	$70 \pm 20$	0.1	30
Nelder-Mead	$100 \pm 30$	$90 \pm 30$	0.1	30
Conjugate Gradient	$100 \pm 0$	$50 \pm 10$	0.1	20
Random Forest	$80 \pm 20$	$150 \pm 40$	400	60
Genetic Algorithm	$100 \pm 10$	$150 \pm 3$	0.2	60
Random Sampling	$60 \pm 20$	$150 \pm 0$	1	60

Table 4.2: Median convergence, and median number of iterations ( $n_{\text{iters}}$ ) until convergence, to the branin function global optimum by different optimisation algorithms. The total time required for each optimisation is shown alongside the equivalent time which would be required to optimise the maximum proton energy in simulations presented in section 4.6.1, if the median  $n_{\text{iters}}$  simulations were run.

[270], Nelder-Mead [271]) identified the global optimum on most repetitions ( $\sim 100\%$  convergence) using very little computational resource, taking only milliseconds to complete 150 function evaluations on each repetition. However, the Nelder-Mead and Powell methods are susceptible to becoming stuck in local optima and thus missing the global optimum, hence the high standard deviation in the median convergence. These methods also require many iterations to achieve  $\geq 95\%$  convergence. The random forest method performs poorly on all measures, as does random sampling which is a dummy operation where no optimisation is actually taking place, highlighting why optimisation methods are necessary. The genetic algorithm converges to the global optimum on most of the 20 repetitions, though it always requires the entire 150 function evaluations to do so. When not limited in this way, the algorithm converges on every repetition (average of 99% convergence in  $500 \pm 70$  iterations), requiring little computational resource and time to do so.

For the optimisation results presented in table 4.2 the function is already known and so a large number of evaluations is possible. However, for the results presented in section 4.6.1 this is not the case, since each evaluation will be a simulation which takes  $\sim 22$  minutes to run on 1520 CPU cores. As such, the most important consideration is not the speed of the algorithm in searching

the parameter space (optimisation time), as might be the case on experiments where the laser operates at  $\gtrsim 1$  Hz [82]. Instead, it is most important to choose an algorithm which identifies the global optimum consistently, in few iterations, minimising the number of computationally expensive simulations which must be run. In this regard, the Bayesian optimisation algorithm included in scikit-learn is clearly the most effective, as demonstrated in table 4.2, which shows the total time that each algorithm would require to optimise the branin function if every evaluation was a simulation. This was calculated by dividing the total optimisation time by the median  $n_{\text{iters}}$  until convergence and adding the total time required to perform  $n_{\text{iters}}$  simulations. As a result, the disadvantage of an increased optimisation time for Bayesian optimisation compared to the other ML techniques in table 4.2 when optimising the branin function is not relevant for the optimisation problem presented in this chapter.

Another potential disadvantage of Bayesian optimisation is that the optimisation time can scale poorly as the number of parameters to be optimised increases [272]. However, this is also a potential problem for most ML-based optimisation techniques, and depends on the specific optimisation problem. To ensure this wasn't an issue for these results, the Bayesian and Nelder-Mead/Powell algorithms were applied to optimise the 6 dimensional so-called Hartmann function [273, 274], and the time required to optimise increased by  $\sim 2.5\times$ , and  $\sim 3.4\times$  over that required to optimise the branin function for the respective algorithms. As such, there is no significant increase in Bayesian optimisation time with an increasing number of parameters for the results presented. Furthermore, the optimisation time can be improved by running the algorithm over multiple CPU cores in parallel [275], and optimising for too many parameters can actually diminish performance, according to the Hughes effect [276].

Given these benchmarking results, Bayesian optimisation was selected as the best candidate to optimise the maximum proton energy achieved via TNSA in sections 4.6.2 and 4.6.3.

## 4.4 Bayes theory and Gaussian process regression

Now that Bayesian optimisation has been selected to optimise ion acceleration in the TNSA regime it is important to describe how this ML technique works. Fundamentally, Bayesian optimisation follows the principles of Bayes theorem [277] which describes the probability of some condition A occurring given that condition B has occurred,  $P(A|B)$ .

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (4.1)$$

This is known as the *posterior* probability, and it depends on the *prior* probability of condition A occurring,  $P(A)$ , the probability of condition B occurring (which normalises the sum of the probability distribution to 1), and the probability of condition B occurring given that condition A has occurred,  $P(B|A)$ . In other words, the *prior* probability is an initial belief that condition A will occur which is updated to a more accurate posterior probability as more information or data is gathered (probability of condition B occurring). As such, Bayes theorem allows us to update our belief that a condition will occur as we receive new information, meaning we can make predictions with more confidence as new data is gathered. This is explained intuitively in reference [278]. Bayes theory is the principle at the heart of the Bayesian optimisation process, which begins by building a model of an objective function from some initial data.

For the work presented in sections 4.6.2 and 4.6.3, the objective function is the maximum proton energy generated during TNSA with several key laser and target parameters as the variables. Here, the objective function was modelled using Gaussian process regression (GPR), which fits multiple Gaussian functions through some initial data, and integrates over all of these functions to produce a mean function which approximates the objective function [264]. This is visualised in figure 4.3 where GPR is applied to an example 1D objective function.

At first there is no data available for the Gaussian functions to fit to and the

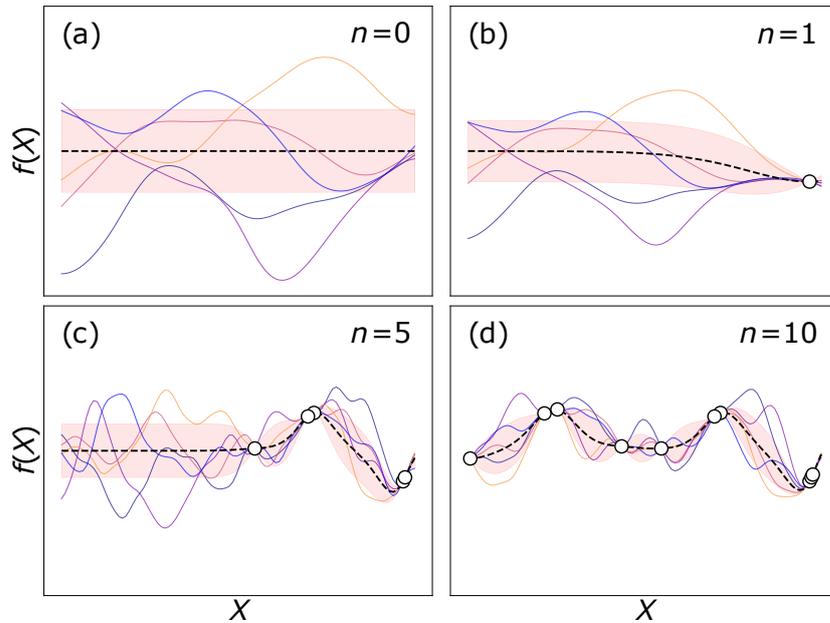


Figure 4.3: Example of Gaussian process regression on a 1D objective function. In (a) there is no data for the five Gaussian functions (coloured lines) to fit through and the mean (black dashed line) of the Gaussian process model is zero across the parameter space. As  $n$  evaluations,  $f(X)$ , of the function are provided in (b)-(d) the five Gaussian functions fit to the data, the mean becomes non-zero across the parameter space, and the uncertainty ( $\pm 1$  standard deviation shown in shaded red) reduces in regions where evaluations have been made.

mean evaluation,  $\mu$ , of the function is zero across the parameter space,  $X$ , as shown in figure 4.3(a). In other words, the *prior* probability is zero, since there is no data upon which to make predictions. As this is a Bayesian approach, the prior probability will be updated as more data is collected. This is seen in figure 4.3(b)-(d), where the model begins to replicate the objective function with increasing accuracy as up to 10 evaluations are made. To achieve this, structures within the data are approximated using a kernel,  $k$ , which computes the covariance,  $\text{Cov}$ , between points in the parameter space.

$$\text{Cov}[f(X), f(X')] = k(X, X') \quad (4.2)$$

Many different kernels have been proposed to model certain structures within a data-set [264, 279]. Choosing the correct kernel to model structures that are expected in a given objective function is perhaps the most important factor to consider before applying Bayesian optimisation. For this example and the results presented in sections 4.6.2 and 4.6.3, a Matérn kernel [279] was selected which is

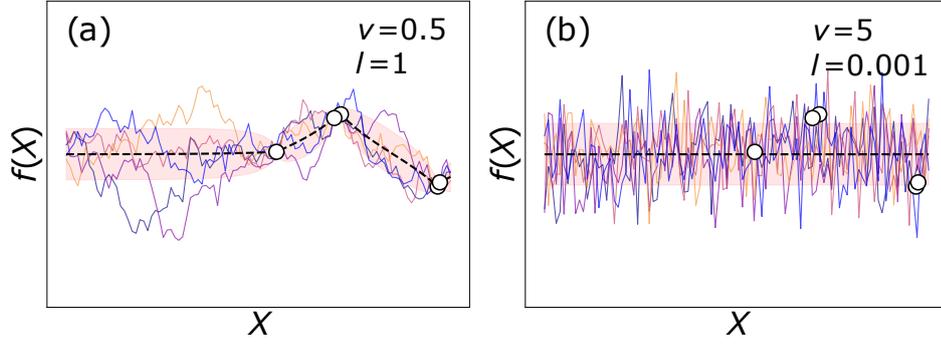


Figure 4.4: Creating a GPR model using the same data as figure 4.3 (c) but varying the (a)  $v$  and (b)  $l$  parameter in the Matérn kernel.

defined as,

$$k(X_i, X_j) = \frac{1}{\Gamma(v)2^{v-1}} \left( \frac{\sqrt{2v}}{l} d(X_i, X_j) \right)^v K_v \left( \frac{\sqrt{2v}}{l} d(X_i, X_j) \right) \quad (4.3)$$

where  $K_v$  is a modified Bessel function of the second kind,  $l$  is the Euclidean distance between points  $X_i$  and  $X_j$ , and  $l$  is the length scale of the kernel which determines the variability of the modelled function between  $X_i$  and  $X_j$ . The final parameter,  $\Gamma$ , represents a so-called gamma distribution, which acts as a prior probability distribution, of ‘shape’ and mean defined by parameters  $\alpha$  and  $\beta$  respectively [264]. The Matérn kernel is an extension of the commonly used radial basis function (RBF) kernel, including an additional parameter,  $v$ , which determines how *smooth* the modelled function will be.

For small values of  $v$ , the kernel will produce a function which fluctuates significantly between points which are close together (not *smooth*), whilst as  $v \rightarrow \infty$ , the Matérn kernel tends to the RBF kernel and the resulting function will be *smoother*. This can be visualised in figure 4.4(a), where the same data and Matérn kernel have been used as in figure 4.3(c) but for the  $v$  parameter which has been changed from 5 to 0.5 ( $l=1$  in both cases), resulting in a much less *smooth* function. This tunability makes the Matérn kernel more adaptable to modelling different structures within a data-set compared to other kernels.

The length scale parameter,  $l$ , is similarly influential in determining the modelled function, since it determines the range of the parameter space over which

covariance between data points is expected. For large  $l$  values, points  $X_i$  and  $X_j$  which are far from one another (large euclidean distance) have a high covariance, whilst for small  $l$  values more variability is expected in the function between these points. As  $l \rightarrow 0$  there is essentially no relationship between points in the parameter space, and the kernel becomes a *white noise* (WN) kernel as shown in figure 4.4(b) which uses the same data and Matérn kernel as the GPR model in figure 4.3(c) but for the length scale which has been changed from  $l=1$  to  $l=0.001$  ( $v=5$  in both cases). Whilst not very useful as the main kernel of a model, adding a white noise kernel to the main kernel allows uncertainty in the objective function to be modelled. For this reason, a white noise kernel is added to the Matérn kernel for the results presented in sections 4.6.2 and 4.6.3

Parameters such as  $l$  and  $v$  are commonly referred to as *hyper-parameters*, and must be specified carefully for a GPR model to accurately model an objective function and the structure(s) present within it. Misspecification of these hyper-parameters can lead to the GPR model failing to replicate the objective function [264, 279], though the sci-kit algorithm automatically optimises each hyper-parameter to mitigate against this [266].

Perhaps more crucial is choosing a kernel that can accurately model the objective function in the first place. As an example, data-sets that include complicated periodic or step like variations, and that contain multiple different structures, may require several different kernels to be combined as described in references [264] and [279]. For the results presented in this chapter the objective function to be optimised is the maximum energy of TNSA protons, the variation of which is expected to be smooth and continuous over the parameter ranges considered. As such, a combination of the Matérn kernel and a WN kernel were used to model this objective function during the Bayesian optimisation results presented in sections 4.6.2 and 4.6.3.

## 4.5 Bayesian optimisation

Bayesian optimisation is a ML technique that uses Gaussian process (GP) regression to model and optimise a given objective function. An example optimisation process is illustrated in figure 4.5, where the Bayesian optimisation algorithm from scikit-learn [266] was used to maximise a simple 1D objective function. To begin this process, a number of function inputs,  $X_n$ , are randomly generated within the limits of the parameter space and are evaluated by the objective function,  $f(X_n)$ , as shown in figure 4.5(a). These evaluations are then used to create a model of the objective function using GPR as detailed in section 4.4. Next, the probability of finding an improved evaluation of the objective function, compared to those that have already been made, is calculated for thousands of potential input values by combining the mean and standard deviation of the GP model to create an acquisition function. The input value which maximises the acquisition function is then evaluated by the objective function on the next iteration of the Bayesian optimisation process, as shown in figure 4.5(b). This new information is used to update the GP model, and another input value is suggested to be evaluated by the objective function, resulting in an iterative feedback loop which is repeated for  $n$  iterations. In this example, the objective function optimum is identified within 10 iterations as shown in figure 4.5(c). After 25 iterations the GP model is confident that the function has been optimised, and the acquisition function only suggests sampling the optimal conditions, as shown in figure 4.5(d).

Just as different kernels can be used for GPR, several acquisition functions are available to be used in the Bayesian optimisation process. The Upper Confidence Bound (UCB), Probability of Improvement (PI), and Expected Improvement (EI) are the most common, and are constructed slightly differently [280]. Typically, only one of these functions is used throughout the optimisation process, and the convergence of the algorithm to the objective function's optimum may depend on which is selected. For the results presented in sections 4.6.2 and 4.6.3, candidate input values are sampled by each of these acquisition functions at every iteration, rather than by just one function throughout the entire process. This technique

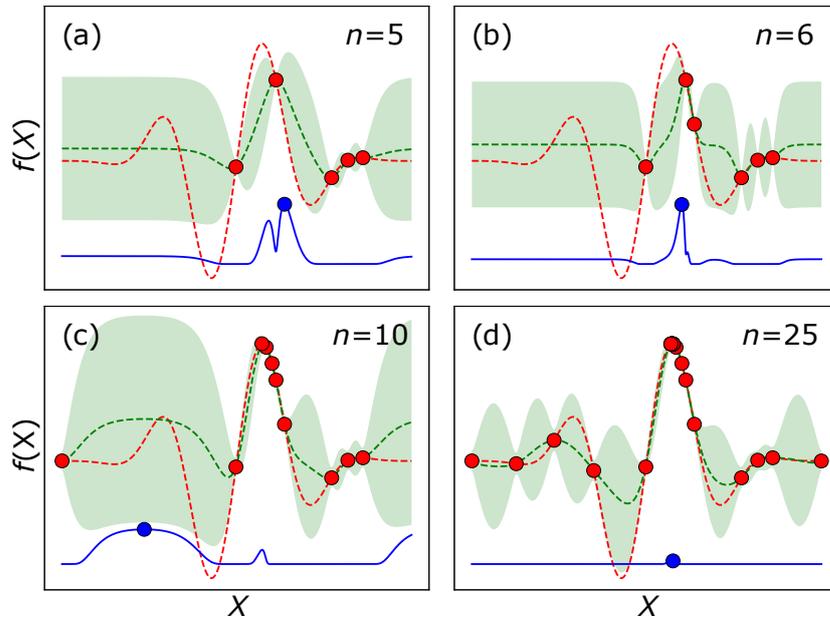


Figure 4.5: Illustration of a typical Bayesian optimisation procedure, where a GP model of mean,  $\mu(X)$ , (dashed green) and uncertainty,  $\sigma_{\text{GP}}(X)$ , (green fill) is built from evaluations,  $f(X)$ , (red dots) of an objective function (dashed red). An acquisition function (blue) is derived from the model after each evaluation and used to determine the input,  $X_{n+1}$ , which gives the highest probability of finding an improved  $f(X)$  compared to previous evaluations (blue dot). This point is then evaluated by the objective function and the process is repeated for  $n$  iterations.

is referred to as *hedging*, and has been shown to maximise the performance of a Bayesian optimisation algorithm [281]. In addition, hedging eliminates the need to choose an optimal acquisition function from the outset, reducing the number of variables for the user to consider before applying Bayesian optimisation. Hedging does increase the computational time required to select new input values between iterations of the optimisation process, but the increase is negligible compared to the time required to perform individual simulations in sections 4.6.2 and 4.6.3.

When an acquisition function is selected, candidate input values are sampled using the mean,  $\mu_{\text{GP}}$ , and uncertainty,  $\sigma_{\text{GP}}(X)$ , of the GP model. Each acquisition function contains a parameter that determines how much weight should be given to the mean and uncertainty. For example, the UCB is constructed as  $\text{UCB}(X) = \mu_{\text{GP}} + \kappa\sigma_{\text{GP}}(X)$ , where small values of the *kappa* parameter,  $\kappa$ , will cause the acquisition function to preferentially select input values in areas of the parameter space where the mean of the GP model is high, whilst penalising those in regions where the model uncertainty is high. This is known as favouring *exploitation* of the parameter space - focusing on regions where the model

predicts the highest probability of finding the optimum of the objective function - rather than *exploration* of areas which have been sparsely sampled and so the model uncertainty is high. The other acquisition functions mentioned, EI and PI, contain a parameter,  $\xi$ , which is analogous to  $\kappa$  in UCB and works in a similar fashion. Determining the balance between *exploitation* and *exploration* of the parameter space is another important variable to consider when setting up a Bayesian optimisation process and is explored in section 4.6.2.

It is clear then, that the configuration of a Bayesian optimisation algorithm should be considered carefully before it is applied to an objective function. For the results presented in sections 4.6.2 and 4.6.3, the Bayesian optimisation algorithm within scikit-learn was used with a Matérn kernel combined with a WN kernel since the objective function is expected to be relatively smooth but return noisy evaluations. The  $\nu$  parameter of the Matérn kernel is set to 5 and the length scale is automatically tuned at each iteration to improve performance, and to reduce the need to specify a length scale in advance. Similarly, hedging is used to select the optimal acquisition function at each iteration to improve the Bayesian optimisation performance by identifying the objective function optimum more quickly. When UCB or EI/PI are selected,  $\kappa=1.96$  and  $\xi=0.01$  are used, respectively. These are similar to the values used in Hoffman *et al.* [281].

## 4.6 Results

### 4.6.1 Automated grid-search of laser parameters to optimise TNSA

The focus of this chapter is on optimising the maximum energy to which protons are accelerated during TNSA in as few simulations as possible by applying machine learning (ML). In section 4.3, several ML techniques were applied to optimise an example 2D function with known optima and Bayesian optimisation was shown to consistently identify the global optimum in the least iterations.

This testing gave a good indication that Bayesian optimisation could effec-

tively identify optimal parameters as required for the results in this chapter, though this was for a known function where the optimisation results could be compared directly to a known global optimum. For the results in section 4.6.3, the algorithm will be applied to optimise the maximum energy of protons accelerated during a complex multi-dimensional laser-solid interaction, for which no prior optimum is known. As such, it was prudent to first generate a reference data-set of simulation data so that the performance of the Bayesian optimisation algorithm could be benchmarked in this regime in a similar fashion to the testing of section 4.3. This avoids the possibility of wasting significant computational resource by applying an algorithm that is not able to successfully optimise in this regime. It also ensures that the optimal laser and target conditions identified were physically meaningful, since Bayesian optimisation is inherently a ‘black box’ approach with no physical model for the data being optimised. In other words, the algorithm may still be able to find an optimal solution to a function which is not physically correct, so it is important to ensure that the maximum proton energy scales with the input laser and target parameters as expected for TNSA in these simulations to avoid this possibility. To mitigate these risks, a reference data-set of one hundred 2D PIC simulations, upon which the algorithm could be tested, was generated using BISHOP’s *Grid-Search* mode as described in section 4.2.

Each simulation took  $\sim 22$  minutes to run using the fully relativistic 2D EPOCH PIC code [171] on 1520 of the skylake CPUs provided by CSD3 [265]. Grid dimensions of  $16,000 \times 2,400$  cells were defined, corresponding to a domain of  $(-10 \rightarrow 50)$   $\mu\text{m}$  and  $(-20 \rightarrow 20)$   $\mu\text{m}$ , and a resolution of  $\sim 4$  nm/cell and  $\sim 17$  nm/cell in the  $X$  and  $Y$  dimensions respectively, and there was an initial average of 10 particles per cell, per species. The target was defined over a 1  $\mu\text{m}$  region in  $X$  containing  $\text{H}^+$  and  $\text{C}^{6+}$  ions with a starting temperature of 0.04 keV and electrons with a starting temperature of 10 keV. The initial electron density was selected as  $203n_{\text{crit}}$ , where  $n_{\text{crit}}$  is the plasma critical density, defined as  $n_{\text{crit}} = \frac{m_e \epsilon_0 \omega_L^2}{e^2}$ , and  $m_e$ ,  $\epsilon_0$ ,  $\omega_L$ , and  $e$  are the electron mass, free-space permittivity, laser angular frequency, and the elementary charge, respectively. This initial electron density

was chosen to approximate a CH plastic target foil by neutralising an equal mix of  $H^+$  and  $C^{6+}$  ions.

When considering the laser and plasma parameters to be varied it was important to choose those which have been extensively studied and shown to have a significant influence in defining ion acceleration in the TNSA regime. To this end the laser energy,  $E_L$ , and pulse duration,  $\tau_L$ , were selected as key parameters to be varied, since numerous studies have modelled their influence in defining TNSA [40, 41, 51]. BISHOP was provided with bounds of  $E_L=(0.3\rightarrow 15.0)$  J and  $\tau_L=(25\rightarrow 100)$  fs, for these parameters respectively, and the laser wavelength and focal spot size (FWHM) were fixed at  $\lambda_L=800$  nm and  $\phi_L=3$   $\mu\text{m}$ . In simulations the entire defined laser energy is contained within the  $\phi_L=3$   $\mu\text{m}$  focal spot, in contrast to experiments where only  $\sim 40\%$  of the laser energy delivered to an OAP is focused into a spot of equivalent size. As such, an upper bound of  $E_L=15$  J is chosen to replicate an experimental value of  $E_L\sim 40$  J, so that the maximum proton energy is optimised over a range of conditions covering those available with Ti:sapphire high-power lasers which are currently available for experiments [238], as well as those currently being commissioned for experiments in the near future [80–83]. The laser energy and pulse duration were varied linearly, in equally spaced values, between the bounds provided, over 100 simulations.

In each simulation, the target front surface is irradiated at the focal position of the laser beam, defined as  $X=0$   $\mu\text{m}$ . The peak laser intensity arrives at this position at  $t=0$  fs. As the laser interacts with the target it transfers energy to plasma electrons, the most energetic of which propagate through the target and form an electric sheath field at the rear side, setting up the TNSA mechanism [51]. This fast (kinetic energy  $\geq 0.511$  MeV) electron population has a broad exponential energy distribution of the form  $\frac{\partial N}{\partial E} \propto \exp(-\frac{\epsilon_f}{k_B T_f})$ , where  $\epsilon_f$  is the electron energy and  $T_f$  is the temperature of the population. Studies have shown that the fast electron temperature increases with the laser irradiance [72, 188, 189, 282], and that it is one of the key properties defining the maximum energy ion energy achieved via TNSA [40, 41, 51].

Given its importance in the TNSA mechanism, the fast electron temperature

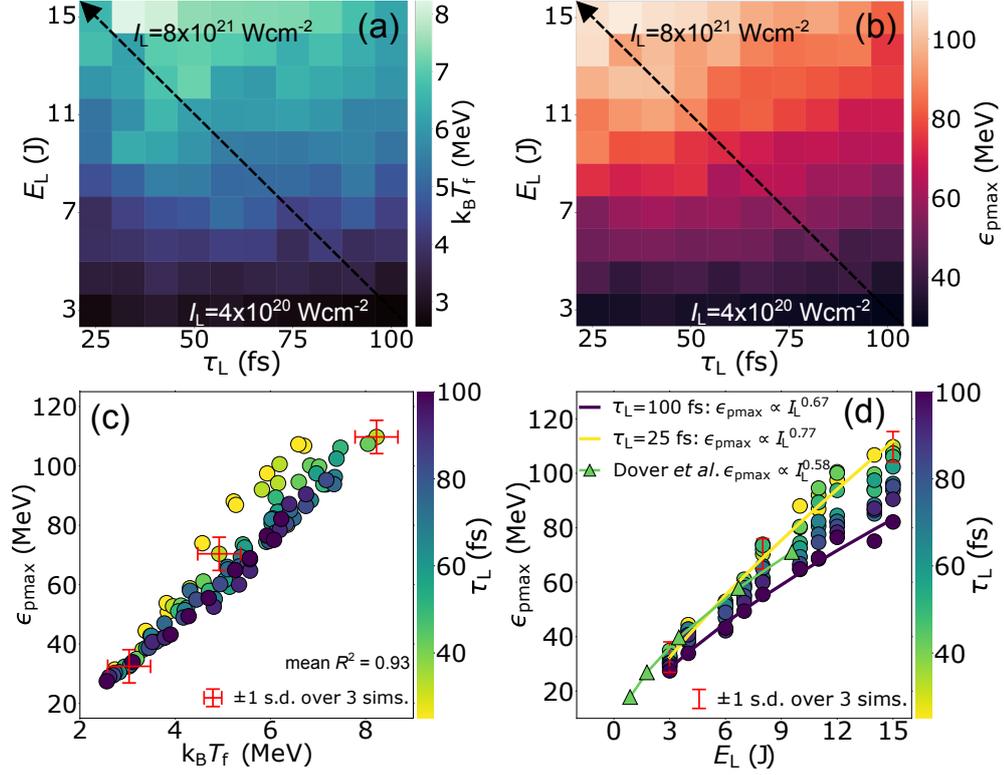


Figure 4.6: (a) Fast electron temperature,  $k_B T_f$ , and (b) maximum proton energy,  $\epsilon_{pmax}$ , as a function of laser pulse duration,  $\tau_L$ , and energy,  $E_L$ , which were varied systematically by BISHOP over 100 2D PIC simulations. The corresponding minimum and maximum laser intensities are given. (c)  $\epsilon_{pmax}$  as a function of  $k_B T_f$ . Linear regression was performed at each pulse duration for data points with increasing laser energy, and the mean  $R^2$  value is calculated to be 0.93. Error bars indicate  $\pm 1$  standard deviation in  $k_B T_f$  and  $\epsilon_{pmax}$ , calculated from three repeat simulations. (d)  $\epsilon_{pmax}$  as a function of laser pulse energy, with given power law scaling for the shortest and longest pulse duration simulated (circles). Experimental  $\epsilon_{pmax}$  data from Dover *et al.* [42] is multiplied by 2.2 (triangles) and compared to the 2D simulation results.

was measured in each of the 100 simulations and is shown as a function of the laser conditions in figure 4.6(a). For fixed pulse duration the laser intensity is increased by increasing energy, and a power scaling of the form  $k_B T_f = a \cdot I_L^b$  with  $b=0.58 \pm 0.03$  is found, in good agreement with previously published results [72]. The maximum temperature is obtained for  $E_L=15$  J and  $\tau_L=33$  fs.

In each case, the fast electron temperature was determined by fitting to the high energy component of the electron energy distribution, from 0.511 MeV to 90% of the maximum electron energy, as shown in figure 4.7(a). This ensured that only relativistic electrons were considered and reduced the influence of significant statistical noise at the highest energies due to a relatively low number of electrons. This method provides an accurate linear fit as shown in figure 4.7(b), where the

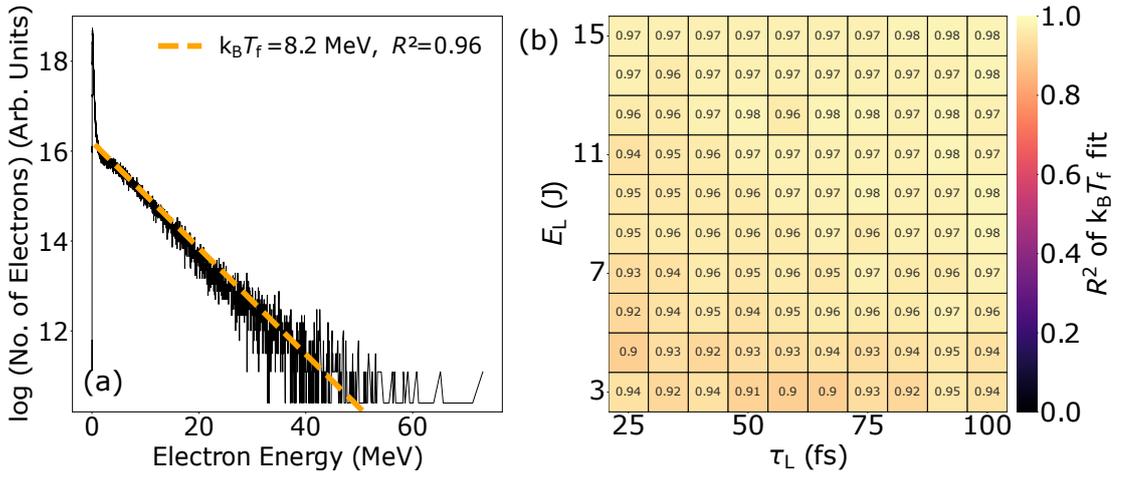


Figure 4.7: (a) Example electron spectrum for a simulation with  $E_L=15$  J and  $\tau_L=25$  fs. The fast electron temperature is calculated by fitting from 0.511 MeV up to 90% of the maximum electron energy (dashed orange). The calculated temperature and coefficient of determination,  $R^2$ , are shown. (b)  $R^2$  value of the electron temperature fit for 100 simulations with the same target, where the laser energy and pulse duration were linearly varied.

coefficient of determination is  $\geq 0.9$  in every case, making this a viable method for determining the electron temperature.

As with the electron temperature, the highest maximum proton energy ( $\epsilon_{pmax}=109$  MeV) is obtained for  $E_L=15$  J and  $\tau_L=33$  fs as shown in figure 4.6(b). This is consistent with a strong overall correlation between  $\epsilon_{pmax}$  and  $k_B T_f$  as observed in figure 4.6(c), and between  $\epsilon_{pmax}$  and  $E_L$  (thus  $I_L$ ), as shown in figure 4.6(d). This measurement was made by considering one proton with the highest energy in each simulation. It should be noted that the statistical noise shown at the highest energies of the electron spectrum in figure 4.7(a) can also affect the high energy proton population. Considering the single highest energy proton could therefore result in unreliable measurements of the maximum proton energy, though this property was later measured as the mean of the 1% most energetic protons (thousands of protons), and there was negligible difference compared to the absolute maximum proton energy, the metric used in this study. The scaling of maximum proton energy with electron temperature and laser intensity also remained the same, confirming the number of protons at the maximum energy is sufficiently large so that the absolute maximum energy of a single proton is a reliable metric.

It is also noted that the absolute electron temperatures and maximum proton energies are higher than would be expected from 3D PIC simulations and

experiments. This is consistent with previous studies using 2D PIC simulations where the maximum ion energy was up to 3 times higher than in equivalent 3D simulations [283–285]. This is attributed to artificially greater electron heating and slower sheath expansion as a result of the reduced dimensions, resulting in a stronger electric sheath field and increased acceleration time. To enable comparison with experimental results, maximum proton energies from Dover *et al.* [42] are multiplied by 2.2 and included in figure 4.6(d). These experimental results were obtained for similar laser conditions as used in the present study with  $I_L$  in the range  $(0.3 \rightarrow 3.4) \times 10^{21} \text{ Wcm}^{-2}$ ,  $\phi_L \sim 1.5 \text{ } \mu\text{m}$ ,  $\lambda_L \sim 800 \text{ nm}$  and  $\tau_L \sim 40 \text{ fs}$ .

Although the magnitude of the maximum proton energies are overestimated in the simulations, the observed power law scaling is similar to these experimental results, and is generally in good agreement with a range of previously published TNSA scaling results [30, 40, 41, 51]. As such, the simulation results presented in figure 4.6 provide a useful benchmark data-set to evaluate the performance of a Bayesian optimisation algorithm in identifying optimal laser and target conditions for maximising TNSA proton energies.

#### 4.6.2 Efficiency of Bayesian optimisation versus grid-search

The results detailed in the previous section represent a conventional linear grid-search of two key laser input variables that are known to influence the maximum proton energy achieved during TNSA. Optimal input conditions were identified for each of these variables, though many more laser and target parameters are known to influence source properties in the TNSA regime, such as the laser focal spot size [42, 212, 286], target foil thickness [38, 39, 43], and pre-plasma density scale length [46, 47]. As such, it is likely that many interaction parameters must be optimised simultaneously to significantly improve these novel ion sources. However, linearly varying parameters as in section 4.6.1 would result in an exponential increase in the number of simulations required as more parameters are considered. This quickly becomes unfeasible given the computational resources and time required to produce these high resolution simulations. Clearly, a more efficient method is required to enable multi-parameter (>2D) optimisation. As

discussed in section 4.1, Bayesian optimisation is an effective technique for optimising complex multi-dimensional functions. An open-source implementation [266] was therefore integrated into BISHOP and will be used to identify optimal conditions for TNSA as a function of four laser and target parameters in section 4.6.3. However, before optimising for many parameters in this way, it was useful to first benchmark the algorithm’s ability to optimise the maximum proton energy over the same parameter space covered by the conventional linear grid-search of figure 4.6 to ensure the algorithm could successfully optimise simulation data.

To investigate this, the Bayesian optimisation algorithm was restricted to sampling parameter combinations that had been simulated as part of the data-set in figure 4.6 such as  $(E_L, \tau_L)=(3.0 \text{ J}, 25.0 \text{ fs}), (3.0 \text{ J}, 33.3 \text{ fs}) \dots (3.0 \text{ J}, 100.0 \text{ fs})$ , rather than having the freedom to select any input values between the bounds of  $E_L=(3 \rightarrow 15) \text{ J}$  and  $\tau_L=(25 \rightarrow 100) \text{ fs}$ , as would typically be the case. This enabled the algorithm to be tested on already available data with a known global optimum, in a similar fashion to the testing presented in section 4.3.

Initially, the input laser energy and pulse duration were selected at random from the combinations simulated in the grid-search of figure 4.6. The resultant maximum proton energy for each laser set-up was then used to construct an initial GP model as described in section 4.5. An acquisition function then suggested a laser energy and pulse duration combination that was expected to increase the maximum proton energy compared to that which had been measured to that point, initiating the Bayesian optimisation feedback loop which ran for 100 iterations.

This process was repeated for several different configurations of the algorithm to identify a set-up that consistently identified the known global optimum from the data-set in the least iterations. At first, three Bayesian optimisation processes were initialised with 5, 10 and 20 iterations with randomly selected inputs ( $n_{\text{random}}$ ), and hedging was used to select the acquisition function between iterations in each case. A further two configurations were then initialised with  $n_{\text{random}}=20$  but were subsequently guided by a UCB acquisition function with  $\kappa=1$  and  $\kappa=1000$  respectively, rather than by an acquisition function selected

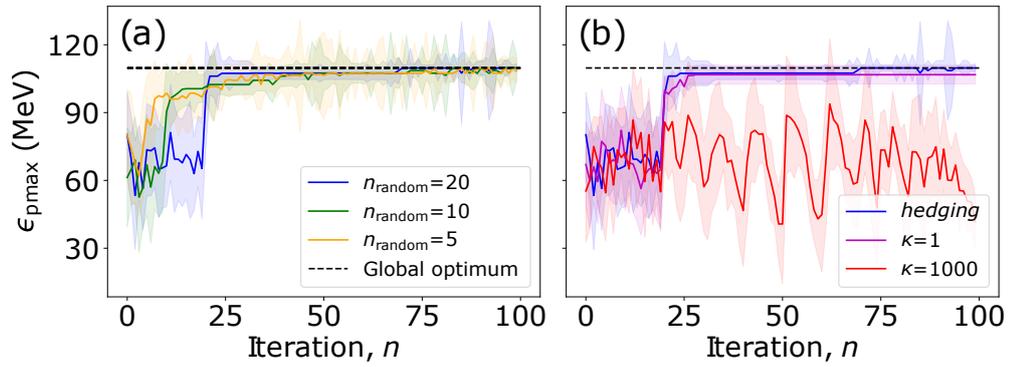


Figure 4.8: (a) Average performance of a Bayesian optimisation algorithm in three different configurations initialised by  $n_{\text{random}}=(5, 10, 20)$  simulations where the laser energy and pulse duration are randomly selected from the combinations of the grid-search. The optimisation was run 20 times for each configuration, and the mean, and standard deviation in the mean, maximum proton energy at each iteration is shown as solid coloured lines and shaded regions respectively. The global optimum maximum proton energy from the grid-search (dashed black) is shown for comparison. The configuration with  $n_{\text{random}}=20$  converges to the optimum in the least iterations and is compared to two further processes in (b) which also use  $n_{\text{random}}=20$  but have a fixed UCB acquisition function with a kappa parameter of  $\kappa=1$  and  $\kappa=1000$ .

through hedging.

For each configuration the Bayesian optimisation process was repeated 20 times to account for uncertainty in the results caused by the random nature of the GP model initialisation. To give an example of the effect this could have, imagine a singular optimisation process which samples the optimum parameters on the first and second of ten random iterations. In this case the GP model would identify the optimal parameters very quickly once initialised. Compare this to a repeat process where the parameter values sampled during the ten random iterations are very far from the optimum. In this case the algorithm would require many more iterations to identify the optimum, despite having an identical kernel, hyper-parameters, and acquisition function as in the first example.

The results of this benchmarking are shown in figure 4.8, where the maximum proton energy is averaged at each iteration over 20 repetitions of the Bayesian optimisation process under different configurations, and compared to the global optimum within the data-set. The associated uncertainty is shown for each iteration as  $\pm 1$  standard deviation. The performance of each configuration is further examined in table 4.3, which shows the median number of iterations,  $n_{\text{iters}}$ , until the algorithm first reaches 95% of the global optimum maximum proton energy.

$n_{\text{random}}$ Iterations	Acquisition Function Set-up	Median $n_{\text{iters}}$ until 95% Convergence	$\sigma$ $n_{\text{iters}}$ until 95% Convergence
5	<i>hedging</i>	26	18
10	<i>hedging</i>	36	28
20	<i>hedging</i>	21	14
20	UCB, $\kappa=1$	24	39
20	UCB, $\kappa=1000$	28	9

Table 4.3: Performance of the Bayesian optimisation algorithm under different configurations.

On average, the operation which uses hedging and is initialised with  $n_{\text{random}}=20$  converges to the global optimum identified during the grid-search of section 4.6.1 in the fewest iterations with the least uncertainty. The algorithm also performs favourably on the other operations where hedging is used, demonstrating the utility of this technique in improving the performance of Bayesian optimisation. As mentioned in section 4.5, hedging also reduces the number of considerations to be made when setting up the algorithm, mitigating the risk of hyper-parameters being misspecified, which can cause poor optimisation results. An example of this is shown in figure 4.8(b), where the best configuration of the algorithm, initialised with  $n_{\text{random}}=20$  and using hedging, is compared to two configurations which are also initialised with  $n_{\text{random}}=20$ , but are subsequently guided by the UCB acquisition function with the kappa hyper-parameter misspecified as  $\kappa=1$  and  $\kappa=1000$ , respectively.

When using  $\kappa=1$ , the algorithm does not converge to the global optimum on average within 100 iterations. This is due to the algorithm spending too many iterations in areas of the parameter space that are well known and where high maximum proton energy is expected, and not enough iterations exploring unknown regions. In this case the algorithm has favoured *exploitation* over *exploration* as mentioned in section 4.5, causing it to become stuck in local optima on most repetitions rather than identifying the global optima. This is demonstrated more clearly in table 4.3, where there is extreme uncertainty in the number of iterations required to achieve  $\geq 95\%$  convergence to the global optimum when  $\kappa=1$ , caused by repetitions where the algorithm has spent many iterations at a local optima

with low maximum proton energy.

When  $\kappa=1000$ , the opposite is true. Here, the algorithm spends too many iterations *exploring* the parameter space and not enough iterations *exploiting* input values that are predicted to result in high maximum proton energy. In this case, the input values are essentially being randomly varied, leading to significant uncertainty in the maximum proton energy identified at each iteration as shown in figure 4.8(b). This random variation results in an apparently favourable performance by the metric of table 4.3. However, the relatively low  $n_{\text{iters}}$  and uncertainty required to achieve a high maximum proton energy for this configuration is due to the objective function being used. Here, the algorithm can only sample a relatively small parameter space of 100 possible inputs. The multi-parameter optimisation exercise of section 4.6.3 presents a much less restrictive sample space where essentially unlimited input conditions are available to be sampled. In this case, it is extremely unlikely that this configuration could identify optimal parameters, as demonstrated when random sampling was applied to the branin function with no limitations to the input conditions in section 4.3.

Overall, these results show the risk of misspecifying acquisition function hyper-parameters. This risk is mitigated by hedging which, when used after initialising the algorithm with 20 randomised data-points, results in a Bayesian optimisation process which consistently identifies the global optimum from the grid-search ( $\epsilon_{\text{pmax}}=109$  MeV, at  $E_L=15$  J and  $\tau_L=33$  fs) in just 21 iterations on average. The sampling performed by the algorithm during one of the 20 processes where the algorithm first identified the global optimum in 21 iterations is visualised in figure 4.9. Here, unique combinations of laser parameter values sampled during the optimisation process are compared to the full linear search grid, highlighting the significant increase in efficiency when Bayesian optimisation is used to search the parameter space.

Each simulation in the grid-search ran for  $\sim 22$  minutes on 1520 cores for a total of 562 core-hours. As such, using Bayesian optimisation to identify the global optimum maximum proton energy in an average of 21 iterations results in an approximately  $5\times$  reduction in the computational resources required to identify

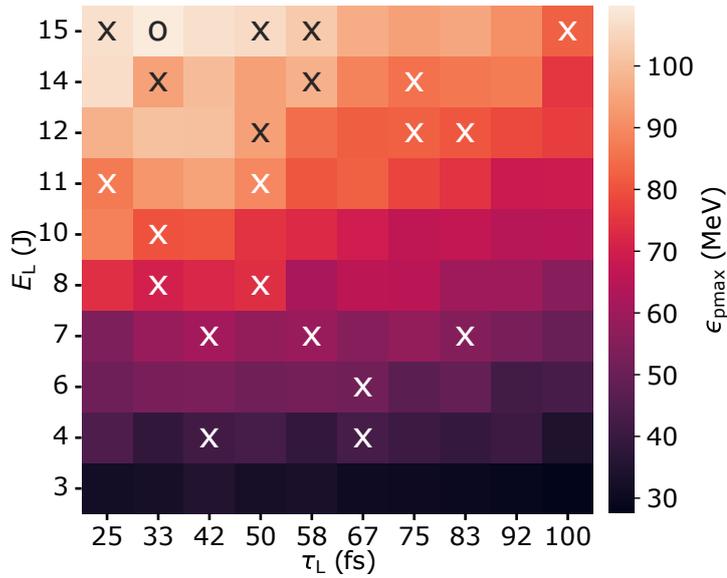


Figure 4.9: Maximum proton energy,  $\epsilon_{p\max}$ , as a function of laser pulse duration,  $\tau_L$ , and energy,  $E_L$ , from the  $10 \times 10$  grid-search of figure 4.6. Bayesian optimisation was applied to find the optimum of this data-set 20 times. Unique parameter combinations sampled during run 6/20 which converged to the optimum (marked by an ‘o’) in the fewest iterations are labelled ‘x’.

optimal conditions, compared to a grid-search that requires 100 iterations. This is equivalent to a reduction of  $\sim 29$  hours and  $\sim 45,000$  core-hours for the HPC cluster used in this investigation.

In this section, Bayesian optimisation has been shown to consistently optimise the maximum proton energy achieved during TNSA as a function of two input laser parameters. A significant increase in efficiency was demonstrated compared to optimising via conventional grid-search, though it is in multi-dimensional parameter spaces where such conventional techniques cannot be applied that Bayesian optimisation truly excels. As such, this technique will now be applied to optimise the maximum proton energy achieved during TNSA as a function of four input laser and target parameters.

### 4.6.3 Multi-dimensional optimisation of proton acceleration in the TNSA regime

Having successfully applied Bayesian optimisation in the TNSA regime as a function of two input laser parameters, the maximum proton energy will now be optimised further by considering other important input variables. The first addi-

tional parameter to be optimised is target foil thickness,  $l$ , which, combined with the laser temporal-intensity profile, plays a significant role in defining the electric sheath field formation at the target rear surface, strongly influencing TNSA maximum proton energies. For a given laser temporal-intensity profile and pulse duration, reducing the target foil thickness results in fast electrons recirculating within the target [38, 194] and absorbing more energy from the laser. This results in an increased fast electron temperature and density at the rear surface, which increases the sheath field strength and thus TNSA maximum proton energies [39]. However, there is a limit to this enhancement. When the target foil thickness is reduced too far, shock waves generated by light from amplified spontaneous emission (ASE) can propagate through the target and form a plasma on the rear surface before the main pulse arrives, reducing the sheath field strength, and thus the maximum proton energy [43, 44].

The second parameter to be added to the optimisation process is the pre-plasma density scale length,  $L$ . This parameter defines the pre-plasma density profile along the laser pulse propagation axis ( $X$ ), described by  $n_e(X) = n_{e0}\exp(X/L)$  for  $X \leq 0$ , where  $n_{e0}$  is solid-density. Studies have shown that pre-plasmas of an optimal scale length induce self-focusing of the laser pulse, resulting in increased laser-to-fast electron energy absorption, sheath field strength, and thus maximum proton energy [46]. In contrast, filamentation of the laser pulse occurs in pre-plasmas of increased scale length, reducing laser-to-fast electron energy absorption and thus the maximum proton energy [47].

For this four parameter optimisation, BISHOP was provided with bounds of  $E_L=(0.3 \rightarrow 15.0)$  J;  $\tau_L=(25 \rightarrow 100)$  fs;  $l=(500 \rightarrow 3000)$  nm; and,  $L=(0.01 \rightarrow 4.00)$   $\mu\text{m}$ . A combination of the Matérn kernel and a white noise kernel were used with their hyper-parameters tuned as detailed in section 4.5. The algorithm was initialised with 20 iterations where values of the four parameters were randomly varied, as justified by the results of table 4.3, before hedging was used to guide parameter values for a subsequent 61 iterations until the optimal parameters were deemed to have been conclusively identified.

The results of this multi-parameter optimisation are shown in figure 4.10.

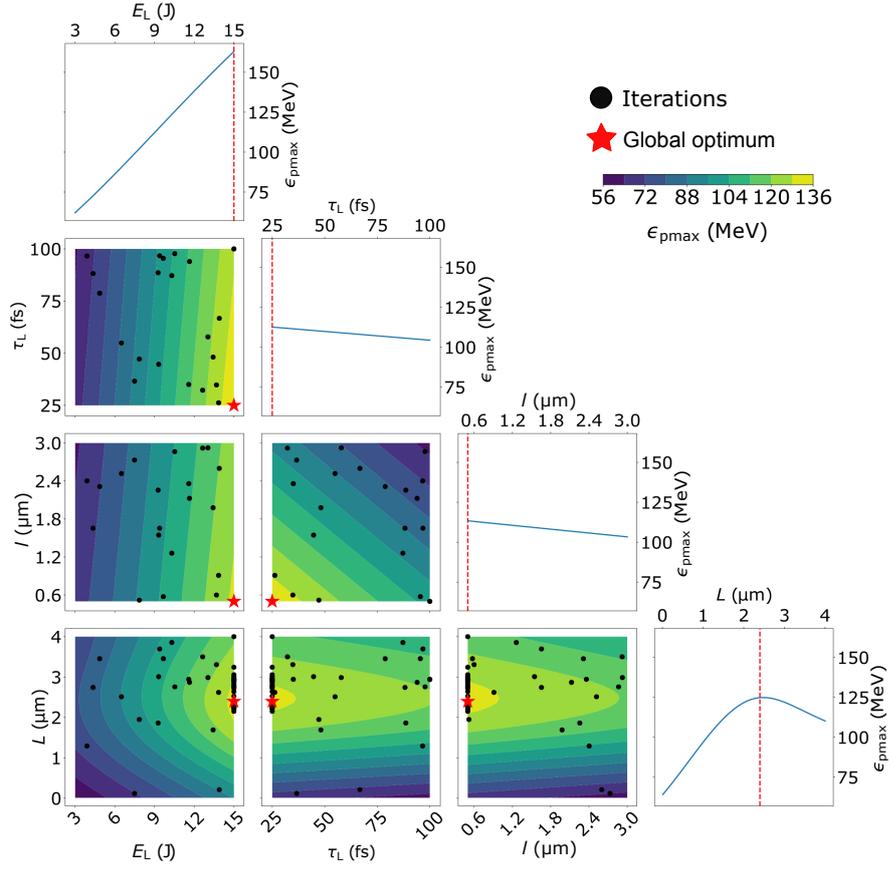


Figure 4.10: Partial dependence plots showing the influence of individual parameters (line plots) and multiple parameters (contour plots) on the maximum proton energy. Points sampled during the optimisation procedure are shown as black dots and the optimum maximum proton energy is denoted by a dashed red line and a red star in the 1D and 2D plots, respectively.

Combinations of input parameter values simulated at each iteration are shown on the contour plots and the influence of individual input parameters on the maximum proton energy is shown in the line plots. As in the grid-search results of section 4.6.1, the maximum proton energy is strongly influenced by the energy, pulse duration, and thus intensity of the driving laser, as is expected for TNSA [40, 41, 51, 282]. Figure 4.10 also shows that the maximum proton energy increases with decreasing target foil thickness, likely due to additional electron recirculation [38]. The thinnest target foil ( $l=500$  nm) available to the algorithm is shown to be optimal, in contrast to previous studies which demonstrate a thicker optimal target. In the present study the laser pulse has an idealised Gaussian temporal profile, without an intense rising edge to launch a shock through the target, as is responsible for degrading the TNSA sheath field and decreasing the maximum

proton energy for thinner targets in previous studies [43, 44]. Reducing the target foil thickness further may result in higher maximum proton energies, though the algorithm was not permitted to sample such targets to avoid complicating this first demonstration of the Bayesian optimisation approach. This is due to other ion acceleration mechanisms such as RPA and RSIT-enhanced hybrid mechanisms occurring alongside TNSA when dealing with ultrathin foils, and resulting in optimal maximum proton energies, as investigated in chapter 5.

Finally, figure 4.10 shows the maximum proton energy to increase with pre-plasma density scale length to some optimal value, before reducing. This was partially expected from previous studies, though in these investigations the maximum proton energy continually increased with scale length up to a value of  $L \sim 100 \mu\text{m}$ , reducing beyond this point due to laser pulse filamentation being induced upon propagation in the plasma [47]. This is obviously much longer than the  $L = 2.40 \mu\text{m}$  optimum shown in figure 4.10, and no laser pulse filamentation is observed for these simulations. As such, the present results demonstrate a much more subtle optimisation of maximum proton energy with pre-plasma density scale length than has previously been observed. This will be explored in section 4.6.5. However, before doing so, it is important to verify that the maximum proton energies are being driven by TNSA in these simulations, so that the comparisons to previously published TNSA results are valid. To this end, figure 4.11 shows the particle energy distribution, in  $X$ , for the simulation with the optimal parameters identified by the Bayesian optimisation algorithm. Here, the proton kinetic energy is plotted at each longitudinal position  $X$  for the final time step of this simulation.

A dominant tail of high energy protons is shown to originate at the target rear surface in figure 4.11, coinciding with the fast electrons reaching vacuum, after being accelerated by the peak laser field and propagating throughout the target. This is consistent with the TNSA mechanism [285]. A component of protons does originate from the plasma at the critical density surface at the front side of the target and propagate through the target into vacuum at the rear surface. However, the energies within this proton population never exceeds  $\sim 40\%$

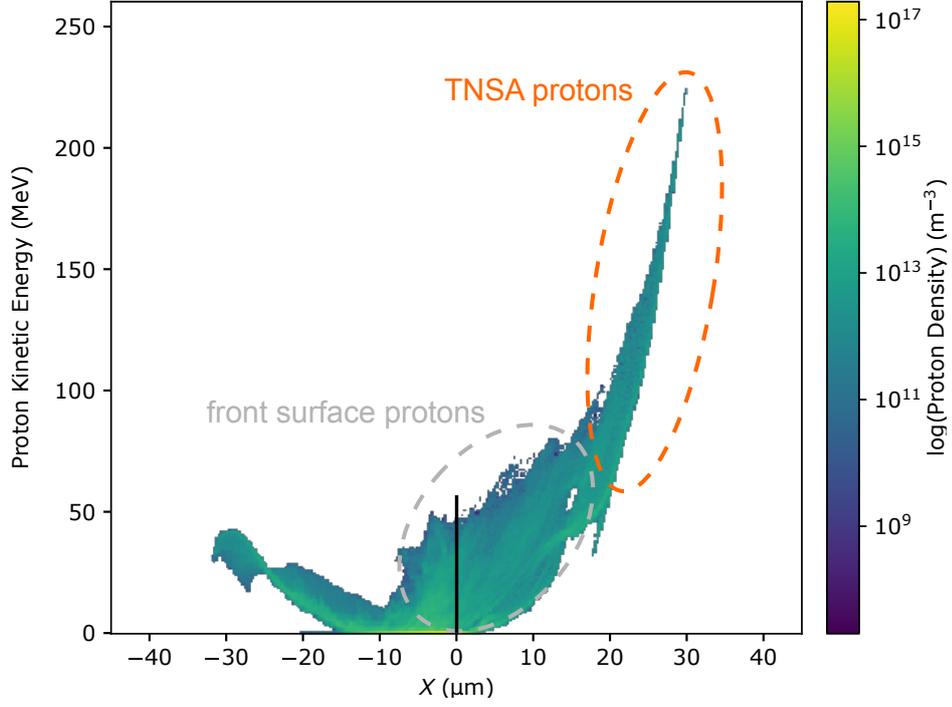


Figure 4.11: Proton energy as a function of longitudinal position, at the final timestep,  $t=230$  fs, of the simulation with optimal laser and target parameters. The target front surface is located at  $X=0$   $\mu\text{m}$ , as illustrated by a black line.

of the maximum energy reached by the protons that originate from the target rear surface. TNSA is therefore confirmed as the primary mechanism responsible for accelerating protons to the maximum energies recorded throughout the investigation presented in this chapter.

Overall, after 20 simulations with randomised inputs and a further 30 simulations guided by a GP model, the algorithm converged to optimal conditions of  $E_L=15$  J,  $\tau_L=25$  fs,  $l=500$  nm and  $L=2.40$   $\mu\text{m}$ , as shown in figure 4.10. This optimal interaction resulted in a maximum proton energy of  $\epsilon_{p\text{max}}=220$  MeV, a twofold improvement over the results of the linear grid-search shown in figure 4.6, for which the scale length was fixed. This highlights the importance of optimising for many input laser and target parameters, which is not always possible in a conventional grid-search. In this case, extending the  $10 \times 10$  grid-search of figure 4.6 to include the target foil thickness and pre-plasma density scale length would require 10,000 simulations, running for over 10,000 hours. This is over a year of run time, and is thus unfeasible in terms of time and computational resource.

This demonstrates a major advantage of using Bayesian optimisation com-

pared to conventional techniques, though it also flags up a potential issue. That is, given that an exhaustive search of the parameter space is not possible for this objective function, the global optimum is not known, making it impossible to state that it has been identified for these results. Using Bayesian optimisation does inherently mitigate this risk, since it is less susceptible to becoming stuck in local optima compared to other ML techniques. This was demonstrated in section 4.3, where Bayesian optimisation consistently identified the known global optimum of the branin function [267]. Similarly, Bayesian optimisation was shown to consistently identify the known global optimum from the linear grid-search in section 4.6.2, using the same configuration of the algorithm as used for the four parameter optimisation results in this section.

Despite these precautions, it is still important to be confident that the global optimum within this multi-dimensional parameter space has been identified. To this end, the results of the Bayesian optimisation process are analysed in detail in the following section.

#### **4.6.4 Investigating optimal laser and target conditions for TNSA of protons in a multi-dimensional regime**

In the previous section, Bayesian optimisation was used to optimise the maximum proton energy within a four dimensional parameter space where the global optimum was not known prior to the optimisation process. The optimal parameters for laser energy and pulse duration were consistent with the results of section 4.6.1 and generally published TNSA results [40, 41, 51]. The optimal target foil thickness was similarly expected from the literature [47], giving some confidence that the global optimum had been identified. However, the optimisation of maximum proton energy with pre-plasma density scale length was not expected prior to these results. As such, it is important to investigate the Bayesian optimisation process. Firstly to ensure that the global optimum has been identified, and secondly, to understand the novel optimisation of maximum proton energy with pre-plasma density scale length. To this end, figure 4.12 shows the objective func-

tion evaluation,  $\epsilon_{\text{pmax}}$ , made at each iteration,  $n$ , of the Bayesian optimisation process.

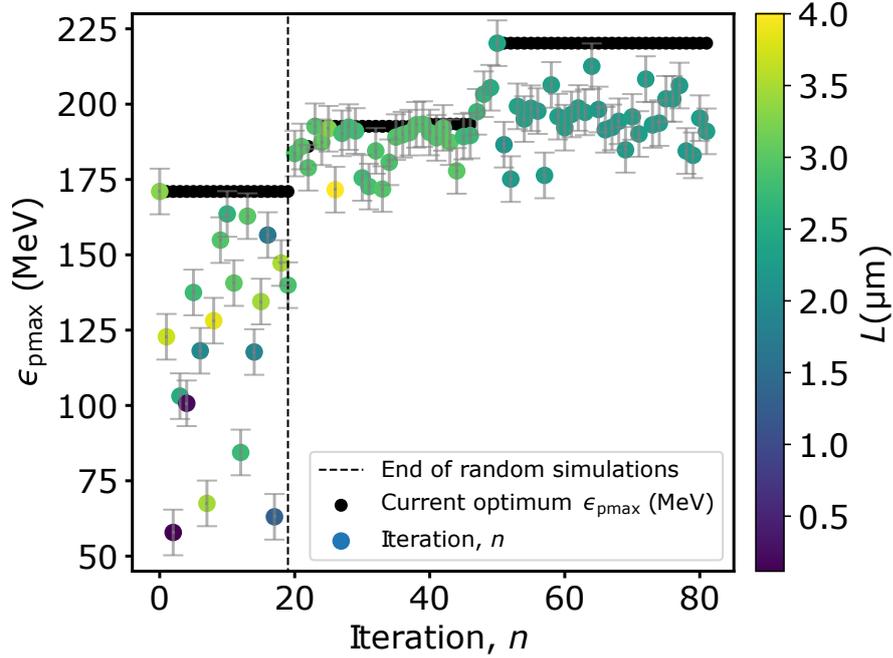


Figure 4.12: Value of the objective function, the maximum proton energy,  $\epsilon_{\text{pmax}}$ , as a function of the iteration number,  $n$ , during a Bayesian optimisation procedure in which the four input parameters  $E_L$ ,  $\tau_L$ ,  $l$  and  $L$  are varied. The highest  $\epsilon_{\text{pmax}}$  identified up to a given iteration is marked by black circles. Before the vertical dashed black line ( $n < 20$ ) inputs are randomly varied and afterwards ( $n \geq 20$ ) a GP model is constructed and used to guide input conditions.

The colour of the symbol identifies the corresponding pre-plasma density scale length.

Once initialised by 20 simulations with randomised inputs, the acquisition function immediately suggested conditions that improved the maximum proton energy compared to the values that had been achieved via random sampling. For three of the four parameters being varied, the values sampled at this iteration were those that proved to be optimal by the end of the optimisation routine. These conditions represent the thinnest target foil, shortest laser pulse duration and highest laser energy available to be sampled by the algorithm, as is consistent with the grid-search results of section 4.6.1 and the generally accepted understanding of the TNSA mechanism [40, 41, 51, 282]. Over a subsequent 60 iterations, the algorithm did not suggest varying the target foil thickness or laser energy, and only varied the pulse duration twice. This means the acquisition function had assigned a high probability to these parameters being optimal after only a few iterations. Moreover, the model identified laser and target conditions which are

expected to be optimal in this acceleration regime with a high probability, giving confidence that the global optimum maximum proton energy had been identified for these parameters.

The algorithm spent many more iterations optimising the pre-plasma density scale length, as shown in figure 4.12, where the marker colour represents the scale length sampled at each iteration. Once the model had been initialised at iteration 20, the scale length was varied over a range of  $L=(2.50\rightarrow 4.00)$   $\mu\text{m}$  until iteration 50 where an optimum of  $L=2.40$   $\mu\text{m}$  was identified. The model only varied the scale length within 5% of this value over a subsequent 30 iterations, likely due to uncertainty in the maximum proton energy evaluations, as discussed in detail in the following section.

Given that three of the four laser and target parameters had been kept constant for almost 60 iterations, and the final parameter was only varied over a very small range for the final 30 iterations, likely due to noisy evaluations, the GP model had assigned a high probability to the global optimum having been identified. The optimal values for three of the four parameters were consistent with those expected for TNSA, giving further confidence that the global optimum had been identified. As such, the Bayesian optimisation process was terminated after 81 iterations.

Now that the global optimum is considered to have been identified, it is important to understand the significant optimisation of maximum proton energy over a pre-plasma density scale length range of only a few microns, since this is much more subtle than was expected from previously published results, where optimisation was observed over tens, or hundreds of microns [46, 47].

#### **4.6.5 Novel optimisation of maximum proton energy with pre-plasma density scale length**

The results presented in section 4.6.3 demonstrate the utility of using Bayesian optimisation to improve the maximum proton energy achievable via TNSA by optimising for multiple key laser and target parameters simultaneously. Another

benefit of this technique is that it can potentially produce interesting optima which were not previously expected. This appears to be the case for these results, where an interesting optimisation of maximum proton energy with pre-plasma density scale length is demonstrated. However, on studying the optimisation process in detail, it was noticed that the algorithm spent many iterations varying the scale length over a very small range once the optimal value of  $L=2.40\ \mu\text{m}$  had been identified. This was expected to be caused by uncertainty in the maximum proton energy evaluations, resulting from subtle changes to the macroparticle distribution in EPOCH. In short, numerical noise in the simulations results in noise in the maximum proton energies. As such, it was important to determine that the optimisation of maximum proton energy with pre-plasma density scale length was physically significant, rather than being driven by statistical uncertainty.

For each simulation performed in the investigations presented in this chapter, the initial location and momentum distribution of the macroparticles is generated using a random number generator with a fixed seed. Initialising a repeat simulation at identical laser and plasma conditions with a fixed seed results in the exact same macroparticle distribution, resulting in the same laser-plasma interaction and maximum proton energy. However, if the scale length varies even by a very small amount, the initial macroparticle distribution changes, resulting in a subtly different laser-plasma interaction which causes variations in the maximum proton energy. By enabling the random seed in EPOCH this behaviour can be induced for simulations with exactly the same initial laser and plasma conditions. As such, an exact repeat simulation with the optimal laser and target conditions identified in section 4.6.3 was performed with the random seed enabled to investigate the level of statistical uncertainty in the maximum proton energy evaluations.

The maximum proton energy reduced by 13% from  $\epsilon_{\text{pmax}}=220\ \text{MeV}$  in the original simulation to  $\epsilon_{\text{pmax}}=191\ \text{MeV}$  on the repeat. This suggests that statistical fluctuations originating from the initial macroparticle distribution in EPOCH contribute to significant uncertainty in the maximum proton energy evaluations at each iteration of the Bayesian optimisation process in section 4.6.3. Given this result, it is important to ensure that statistical uncertainty didn't drive the

optimisation of maximum proton energy with pre-plasma density scale length.

To confirm this, 13 simulations from the optimisation results were investigated where the laser energy, pulse duration and target foil thickness were at their respective optimal values and the scale length was within 5% of the optimum value. The mean scale length was  $L=(2.31\pm 0.03)$   $\mu\text{m}$  for these simulations, resulting in  $\epsilon_{\text{pmax}}=(200\pm 8)$  MeV. In both cases the expressed uncertainty is the standard deviation which defines the error-bars on the maximum proton energy at each iteration in figure 4.12. Similarly, for 20 simulations with a mean scale length of  $L=(2.91\pm 0.06)$   $\mu\text{m}$  and optimal values for the other three parameters, the mean maximum proton energy was  $\epsilon_{\text{pmax}}=(187\pm 6)$  MeV. As such, the optimisation of maximum proton energy at a scale length of  $L=2.40$   $\mu\text{m}$  is statistically significant.

Now that this has been established, it is important to understand why the maximum proton energy was optimised for this pre-plasma density scale length. To investigate this, four simulations from the optimisation routine were repeated with particle tracking enabled in EPOCH so that the source and trajectory of the highest energy protons could be investigated. These repeat simulations were run for a reduced simulation time, finishing at  $t=230$  fs, compared to  $t=560$  fs in the optimisation routine, due to the increased computational resource required to include particle outputs. The optimal input conditions of  $E_L=15$  J,  $\tau_L=25$  fs and  $l=500$  nm were used for each simulation, and the scale length was varied,  $L=(2.15, 2.40, 3.25, 4.00)$   $\mu\text{m}$ .

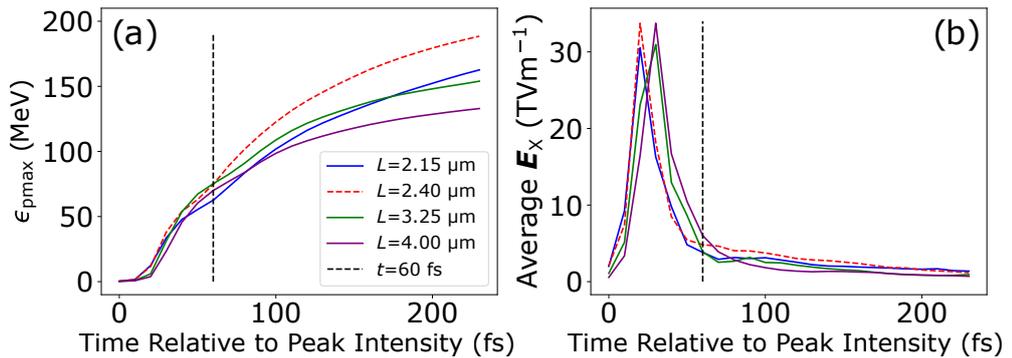


Figure 4.13: Temporal evolution of the (a) maximum proton energy and (b) mean longitudinal electric field within  $3 \mu\text{m} \times 3 \mu\text{m}$  ( $X, Y$ ) of the proton which has the highest energy by the end of four simulations with varying pre-plasma density scale length. A dashed black line illustrates time,  $t=60$  fs, when the tracked proton leaves the target rear surface in each case.

For each case, the location and energy of the proton which has the maximum energy at the final timestep ( $t=230$  fs) was tracked from the beginning of the simulation. The energy of this proton is shown for each timestep and for each scale length in figure 4.13(a). As in the Bayesian optimisation results of section 4.6.3, the maximum proton energy at the final timestep is highest for a scale length of  $L=2.40$   $\mu\text{m}$ . For these laser and target conditions the protons are accelerated by the electric sheath field at the target rear surface via the TNSA mechanism. As such, the longitudinal electric field strength ( $\mathbf{E}_X$ ) was averaged within a  $3 \mu\text{m} \times 3 \mu\text{m}$  region ( $X, Y$ ) around the tracked proton at each timestep. This is shown in figure 4.13(b) for each scale length case.

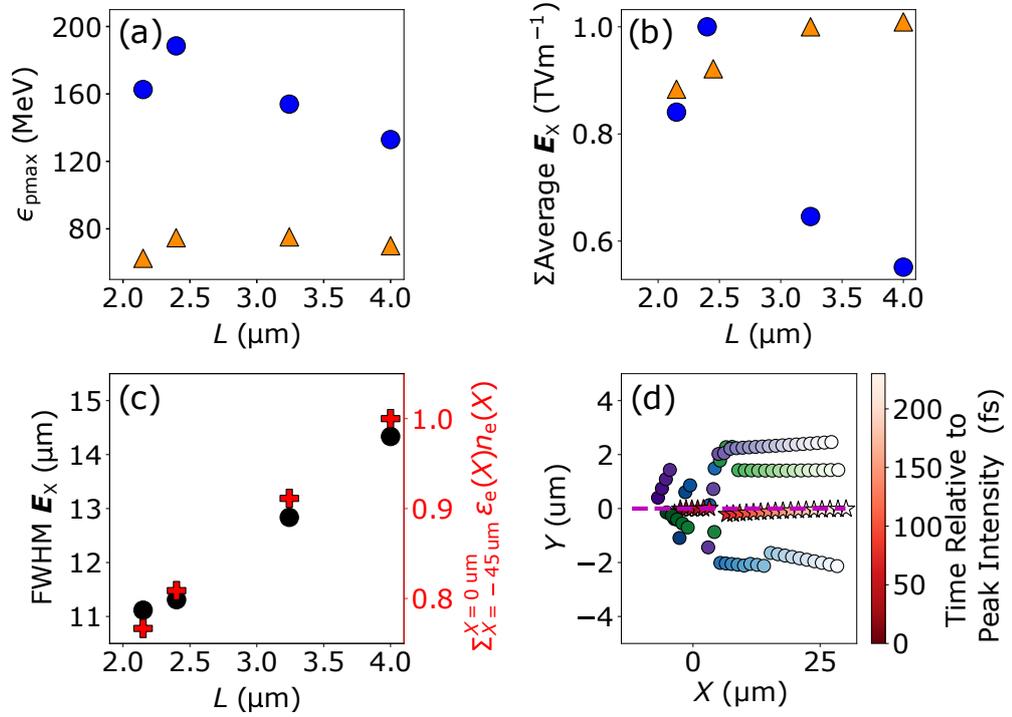


Figure 4.14: (a) Maximum energy and (b) summation of the mean  $\mathbf{E}_X$  around the proton with the highest energy before (orange) and after (blue) it leaves the target rear at  $t=60$  fs, shown as a function of pre-plasma density scale length, with all other laser and target parameters fixed. (c) Size of the sheath field (FWHM) in the dimension transverse to the laser pulse propagation ( $Y$ ) at  $t=0$  fs (black) and the laser-to-electron energy absorption (number of energetic electrons) at the target front side (red), as a function of scale length. (d) Evolution of the position of the proton which finishes the simulation with the highest energy as a function of simulation time (colourbar) for  $L=2.15$   $\mu\text{m}$  (blue),  $L=2.40$   $\mu\text{m}$  (red),  $L=3.25$   $\mu\text{m}$  (green),  $L=4.00$   $\mu\text{m}$  (purple).

In each simulation, the tracked proton leaves the target rear surface at  $t \sim 60$  fs, and it is from this timestep onward where the tracked proton energy deviates significantly for the different pre-plasma density scale lengths considered, as

shown in figure 4.13(a). This is shown more clearly in figure 4.14(a), where the maximum proton energy is shown for each scale length before and after  $t=60$  fs respectively. Measured before  $t=60$  fs, the maximum proton energy is relatively constant as the scale length increases, with a slight peak at  $L=3.25$   $\mu\text{m}$ . However, by the end of the simulation at  $t=230$  fs, the maximum proton energy is significantly higher for the optimum  $L=2.40$   $\mu\text{m}$  case. This is reflected in figure 4.14(b) where the electric field experienced by the proton before it leaves the target rear at  $t=60$  fs is higher for longer scale lengths. This is attributed to increased laser-to-electron energy absorption at the target front side for longer scale lengths, as shown in figure 4.14(c), where absorption is approximated by multiplying the fast electron density and mean kinetic energy in the region at the front side of the target,  $X=(-45\rightarrow 0)$   $\mu\text{m}$ . However, figure 4.14(b) also shows that the electric field experienced by the proton after it leaves the target rear at  $t=60$  fs is lower for longer scale lengths. This results in lower maximum proton energies by the end of the simulation, despite the increased laser-electron energy absorption which is known to increase TNSA maximum proton energies [46, 51].

When considering why increased laser-to-electron energy absorption wasn't resulting in the protons experiencing a higher electric field and thus energy for longer scale lengths, it was realised that the fast electron population will originate further from the target front surface in these simulations. This causes the fast electrons to propagate over an extended plasma region and, given the electron population is divergent, would result in a spatially broader electron distribution in the transverse ( $Y$ ) dimension. This would potentially reduce the electric sheath field strength along the target normal axis ( $Y=0$   $\mu\text{m}$ ), reducing the energy of the protons accelerated to the highest energy in each simulation, which are always accelerated along this axis, as shown in figure 4.14(d).

To confirm this hypothesis, the  $\mathbf{E}_X$  field was summed for each cell along the  $Y$ -axis within  $3$   $\mu\text{m} \times 40$   $\mu\text{m}$  ( $X, Y$ ) of the target rear. A Gaussian profile was fit to this data, and the associated FWHM is shown to increase with scale length in figure 4.14(c). Next, the position of the relativistically-corrected critical density surface,  $\gamma n_{\text{crit}}$ , where  $\gamma$  is the electron Lorentz factor, was determined at each

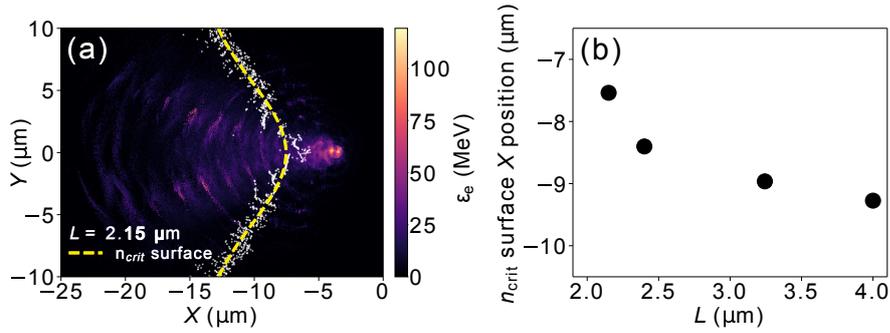


Figure 4.15: (a) 2D spatial distribution of the mean electron kinetic energy at the target front side ( $X < 0 \mu\text{m}$ ) at  $t=0$  fs, for the  $L=2.15 \mu\text{m}$  simulation. White dots represent the location of the  $\gamma n_{\text{crit}}$  surface and the yellow dashed line denotes a Gaussian profile fit to this data. (b) Position of the  $\gamma n_{\text{crit}}$  surface along the laser pulse propagation axis ( $X$ ) with respect to the target initial front surface ( $0 \mu\text{m}$ ), as a function of pre-plasma density scale length.

point along the  $Y$ -axis by averaging the electron kinetic energy over all cells from  $X=(-45 \rightarrow 0) \mu\text{m}$  (i.e. the front side of the target) at the time at which the peak laser intensity arrives ( $t=0$  fs). The position of the critical density surface is defined as the first  $X$  position where  $n_e/\gamma n_{\text{crit}}=1$ , and is shown at each point in the  $Y$  plane for  $L=2.15 \mu\text{m}$  in figure 4.15(a). The  $\gamma n_{\text{crit}}$  position along the laser pulse propagation axis is determined from the peak of a Gaussian fit to this data and is shown to be located farther from the target front surface for longer scale lengths in figure 4.15(b).

The optimisation of maximum proton energy with pre-plasma density scale length in section 4.6.3 therefore results from a trade off between increasing the laser energy absorbed by electrons when the laser pulse interacts with the critical density surface, whilst minimising the distance these electrons must travel to reach the target rear and set-up an electric sheath field. For a pre-plasma density scale length of  $L=2.40 \mu\text{m}$  this results in an optimal fast electron population of high temperature and density along the target normal axis. Both of these parameters are known to be important in defining the electric sheath field strength, and thus the maximum proton energy achieved via TNSA [51].

## 4.7 Conclusion

This chapter reports on the development of a new code called BISHOP, which automatically generates and processes data from PIC simulations of laser-plasma interactions. Initially, BISHOP was used to vary the energy and duration of a high-power laser pulse incident on a planar CH target of  $l=1$   $\mu\text{m}$  thickness, over 100 2D PIC simulations in a conventional grid-search approach. By automating this process, the maximum proton energy within the parameter space was identified in much less time than would have been required if the parameters were varied manually.

In addition to increasing the efficiency of conventional data generation methods, BISHOP also facilitates the application of ML-based optimisation approaches. Several such techniques were applied to an example 2D function to benchmark their performance in identifying optima. Bayesian optimisation was shown to perform most favourably, and was therefore incorporated into BISHOP as a tool to optimise a selected source property from a laser-solid interaction as a function of multiple input variables.

The viability and efficiency of this approach was first demonstrated by applying it to identify the same optimal conditions found during the grid-search. Using this approach, optimal conditions were identified in  $\sim 5$  times fewer simulations than in the grid-search. This corresponds to a significant reduction of  $\sim 45,000$  core-hours for the HPC cluster used to generate the simulations.

Bayesian optimisation was then used to optimise the maximum proton energy as a function of four parameters including target foil thickness and pre-plasma density scale length as variable inputs. In this case, the optimum conditions were identified in  $\sim 200$  times fewer simulations than would have been required to do so by a conventional linear grid-search. Given this is a difference of  $\sim 48$  days of simulation time, a grid-search of four parameters would not have been possible. As the maximum proton energy was doubled by optimising for four input parameters rather than two, the utility of Bayesian optimisation in enhancing desired source properties in laser-solid interactions has been clearly demonstrated.

Another beneficial feature of Bayesian optimisation is the ability to identify non-trivial optimal conditions. In the results presented, the algorithm identified pre-plasma density scale length conditions that produced an optimal balance of increasing laser energy coupling to fast electrons, whilst also increasing the electron density along the target normal axis. This combination resulted in an optimal peak electric sheath field at the target rear surface, resulting in a boost to the TNSA maximum proton energy. This optimisation was identified over a scale length range of a few microns, which was unexpected given previous studies identified optimal scale lengths over tens or hundreds of microns [46, 47], attributing de-optimisation to laser filamentation at very long scale lengths, which was not observed in the present study. This much more subtle optimisation would likely have been missed without employing ML, given that a grid-search over four input parameters would not have been possible.

To summarise, this chapter reports on the development and first implementation of the BISHOP code to automatically generate and analyse large simulation data-sets to be modelled and optimised using ML. Following this first demonstration of the approach, BISHOP has since been used to generate large data-sets of simulation results that have been used to create a model of laser-solid interactions using a neural network. Furthermore, BISHOP and Bayesian optimisation have since been used to optimise the fast electron temperature in simulations as a function of multiple laser and target parameters, and the results have been used to guide the choice of input conditions in an experiment at the Gemini laser facility. Similarly, the methods demonstrated in this chapter have been extended to optimise synchrotron radiation produced from laser-solid interactions as a function of up to six laser and target parameters in simulations. Multiple desirable source properties were optimised simultaneously in this investigation, published in *High Power Laser Science and Engineering* [263], demonstrating the flexibility of this approach.

# Chapter 5

## Stability of enhanced ion acceleration in the relativistic self-induced transparency regime

### 5.1 Introduction

Chapter 4 presents an investigation into the optimisation of maximum proton energies which can be achieved via TNSA. This is the most studied laser-driven ion acceleration mechanism, which has been experimentally demonstrated to produce proton energies of up to  $\sim 90$  MeV [62, 63]. Hybrid acceleration mechanisms enhanced by a target foil undergoing relativistic self-induced transparency (RSIT) have been demonstrated to produce increased maximum proton energies of near-100 MeV [71], and  $\sim 150$  MeV [31] for the same laser conditions. This results from efficient laser-to-electron energy transfer throughout the target volume [65, 66, 76, 78, 287], which boosts the accelerating electric field, and thus the energy of protons accelerated initially due to the laser radiation pressure (RPA) [73, 74, 221, 288, 289], and resultant collisionless shocks (CSA) [290], and by the sheath field at the target rear (TNSA) [28, 29, 33]. Further improvement of these RSIT-enhanced acceleration mechanisms could result in proton beams of sufficient energy for applications such as radiotherapy [155], driven by relatively affordable and compact laser systems with peak intensities on the order

of  $I_L \sim 10^{21}$  Wcm<sup>-2</sup>, rather than requiring systems of ever increasing intensity, footprint and cost.

However, enhanced maximum proton energies have been shown to depend strongly on the RSIT onset time in this regime. This is sensitive to multiple laser and target parameters, and can change significantly if the values of these parameters change, such that the RSIT onset time is no longer optimal, diminishing proton energy enhancement compared to TNSA. This is demonstrated in Higginson *et al.* [71], where the maximum proton energy varies significantly with target foil thickness, from a peak of  $\epsilon_{pmax} \sim 96$  MeV at  $l=90$  nm, to  $\epsilon_{pmax} \sim 55$  MeV at  $l=40$  nm. Simulations demonstrated that this was due to RSIT being induced by laser light preceding the peak intensity in target foils thinner than the optimal, rather than by this peak intensity light in targets of the optimal thickness [79]. A similar reduction was demonstrated in the maximum energy of protons accelerated from target foils thicker than the optimal, in which RSIT is induced by light in the falling edge of the laser temporal intensity-profile, rather than that at the peak intensity. Another experimental investigation [31] demonstrated a peak maximum proton energy of  $\epsilon_{pmax} \sim 150$  MeV which similarly reduced when the RSIT onset time varied, though in this case the target foil thickness was not varied. The discussed simulations also demonstrated the RSIT onset time, and thus maximum proton energy, to vary by  $\sim 40\%$  when the laser had a pronounced rising edge intensity profile preceding the main pulse, rather than a simple Gaussian temporal-intensity profile [79].

Clearly, variations in target foil thickness and the laser temporal-intensity profile induce large variation in the maximum proton energy achieved within an RSIT-enhanced acceleration regime. However, in Higginson *et al.* [71], after the target foil thickness was initially varied to identify an optimum, there were several repeat interactions with this optimal target where the maximum proton energy varied significantly, as shown in figure 5.2 of this chapter. The target foil thickness was kept constant in reference [31] and, though the laser temporal-intensity profile could be fluctuating throughout these experiments, since it can't be measured on every shot, it will not do so to the extent that it was deliberately varied in the

discussed simulations [79]. As such, other factors must be contributing to the observed maximum proton energy variation in these experimental results.

This is investigated in the present chapter, which begins with an investigation into fluctuations in Vulcan-PW laser parameter values on the experiment reported in Higginson *et al.* [71], henceforth referred to as experiment A. The variation in maximum proton energy resulting from laser energy, pulse duration, and intensity fluctuations is then analysed for a series of repeat interactions with the same target foil thickness,  $l=110$  nm, and approximate laser temporal-intensity profile on a following experiment, performed using the same laser, in 2021, henceforth referred to as experiment B. The maximum proton energy is shown to vary by up to 25% when the laser energy and pulse duration vary by just 8%. This is supported by 2D PIC simulations, which show a similar variation in maximum proton energy of up to 21% for the same fluctuation in laser energy and pulse duration, and of up to 44% when these parameters vary more significantly, by 35%. Furthermore, the maximum proton energy is shown to vary by up to 18% across exact repeat simulations where every laser and target parameter is kept constant, demonstrating an inherent sensitivity of proton acceleration in the RSIT-enhanced regime.

This sensitivity is shown to be more significant compared to TNSA dominant proton acceleration from  $l=6$   $\mu\text{m}$  target foils that do not undergo RSIT, and RSIT-enhanced acceleration from  $l=30$  nm target foils where RSIT is induced by light in the rising edge of the laser temporal-intensity profile, rather than by light of peak intensity. As expected, the maximum energy of protons accelerated from these  $l=30$  nm target foils is lower compared to those accelerated from the  $l=110$  nm target foils, which undergo RSIT at an optimal time, though they are not significantly lower, and are still enhanced over those of TNSA protons produced from  $l=6$   $\mu\text{m}$  target foils.

This represents a potential best of both worlds, where the use of  $l=30$  nm target foils results in increased maximum proton energies compared to a predominantly TNSA regime without RSIT, but with a similar robustness to changing interaction dynamics as for this well-established mechanism. Particularly inter-

esting is that this regime, of relatively long pulse duration ( $\tau_L \sim 900$  fs) interactions with  $l \sim 30$  nm target foils, has not been studied in detail. As such, further understanding and optimisation of proton acceleration under these conditions will likely increase maximum proton energies beyond this initial demonstration, as has been the case for TNSA and ‘optimised’ RSIT-enhanced acceleration from  $l \sim 110$  nm target foils over a number of years.

## 5.2 Methodology

### 5.2.1 Experimental set-up

Experimental results presented in this chapter were obtained using the Vulcan-PW laser system, which delivers hundreds of Joules of energy, in a pulse duration of  $\tau_L \sim 900$  fs, as detailed in chapter 3. For these experiments, an off-axis parabola (OAP) of  $f$ -number,  $F_\# = 3$ , was used to focus the beam to a spot size of,  $\phi_L \sim 5$   $\mu\text{m}$ , as illustrated in figure 5.1(a). This results in a laser intensity of  $I_L \sim 10^{20}$   $\text{Wcm}^{-2}$ , used to irradiate target foils ranging from  $l = 30$  nm to  $l = 6$   $\mu\text{m}$  in thickness.

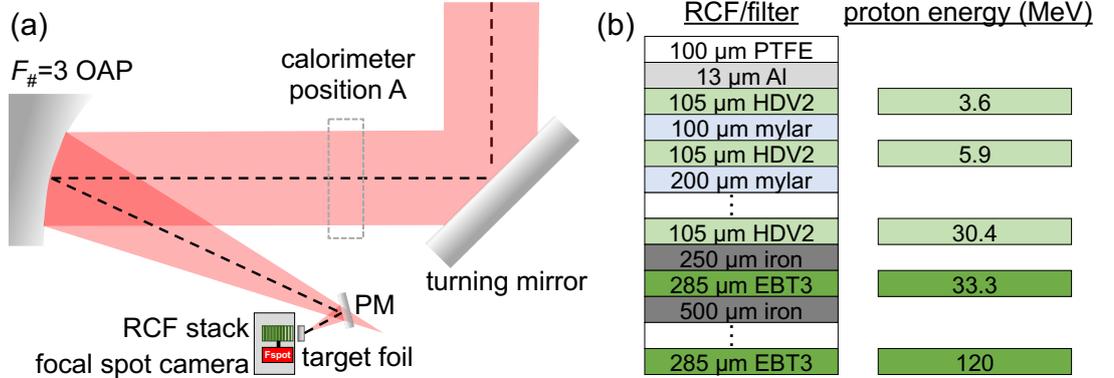


Figure 5.1: (a) Illustration of the Vulcan-PW laser beam focusing onto a planar plasma mirror (PM), and target foil. A calorimeter can be moved into position A, and then into a position after the PM to measure the laser energy at these points in the beamline. A camera is moved into the focused beam before each shot to optimise the focal spot spatial intensity profile. (b) Sample of an RCF stack used to diagnose the energy of protons accelerated from a target foil.

A calorimeter was placed in the collimated beam for two shots on both experiments discussed in this chapter, to calibrate the fraction of laser energy transmitted through the compressor. Dividing the average energy measured on the calorimeter in this position by that measured on a separate calorimeter, placed

behind a partially transmissive optic before the compressor, the throughput was determined as  $\text{comp}_{\% \text{exptA}} = (66 \pm 4)\%$  and  $\text{comp}_{\% \text{exptB}} = (67 \pm 2)\%$ , for the respective experiments. The position of the calorimeter before the compressor does not block the beam, and so the laser energy ‘on target’ was measured throughout the respective experiments as the energy before the compressor multiplied by the throughput,  $E_L = E_{L\text{precomp}} \times \text{comp}_{\%}$ .

A plasma mirror (PM) was used to improve the laser temporal-intensity contrast by up to two orders of magnitude from that shown in figure 2.2. The PM energy throughput was similarly calibrated by moving the calorimeter to a position after the PM, and comparing the laser energy on target to that measured in the new position. Over two shots, the mean throughput was  $\text{PM}_{\% \text{exptA}} = (55 \pm 5)\%$  and  $\text{PM}_{\% \text{exptB}} = (62 \pm 4)\%$ , for the respective experiments.

A second-order autocorrelator [56, 237] placed behind a partially transmissive optic within the compressor was used to characterise the laser pulse duration on each shot, as detailed in chapter 3. The focal spot quality was optimised before each shot using a camera with objective lens focusing to magnify the beam by  $20\times$ , and an image was saved before each shot to determine the focal spot size, and the fraction of laser energy encircled within the spot, as detailed in chapter 3. The camera was then replaced exactly by a target foil before the shot was delivered, ensuring the target is placed at the laser best focus. Combined, the laser encircled energy, pulse duration, and focal spot size determine the peak optical intensity which irradiates the target foil in each interaction, according to equation 3.4.

This chapter reports on an investigation of the level of shot-to-shot fluctuation in each of these parameters, and the influence this has on the maximum energy of protons accelerated in a TNSA dominant, and RSIT-enhanced mechanism. The maximum proton energy was determined for each interaction from the final layer where protons deposited dose in a radiochromic film (RCF) stack, placed behind the target. A partial illustration of this stack is shown in figure 5.1(b), where RCF and filters are layered to measure the number, and spatial profile, of protons accelerated to various energies [291, 292], as detailed in chapter 3.

## 5.3 Results

As described in section 5.1, RSIT-enhanced acceleration has produced the highest proton energies from any laser-driven source [31, 71]. This is promising for improving proton energies towards that required for applications such as radiotherapy [155], though proton energies produced in this regime have been demonstrated to be susceptible to variations in target foil thickness and the laser temporal-intensity profile [45, 79]. There are also indications that RSIT-enhanced proton acceleration is susceptible to variation in other interaction conditions which are yet to be explored. Investigating these dependencies will provide valuable insight into the underpinning physics of the mechanism, allowing it to be controlled and optimised, so that high maximum proton energies can be produced consistently and reliably, as required for applications. This is the focus of the present chapter, which begins by re-examining results reported in Higginson *et al.* [71], to investigate significant variation in maximum proton energy that was observed despite the target foil thickness and laser temporal-intensity profile being nominally constant.

### 5.3.1 Fluctuations in Vulcan-PW laser properties

In experiment A, the Vulcan-PW laser was fired at CH polymer target foils of varying thickness, and the maximum energy of protons accelerated was recorded from the dose deposited on the last layer of an RCF stack. The absolute maximum proton energy from several repeat interactions with each target foil is shown as a function of their thickness in figure 5.2(a). Error bars represent the energy difference between protons stopped by the last layer of the RCF stack, the layer before this (negative), and the layer after this (positive).

This data was published in Higginson *et al.* [71], and shows the maximum proton energy to be optimised for a CH target foil of  $l=90$  nm thickness, resulting in maximum proton energies of near  $\epsilon_{\text{pmax}}\sim 100$  MeV. Numerical simulations were performed to investigate the interaction conditions for each CH target foil, and confirmed the relationship between maximum proton energy and target thickness.

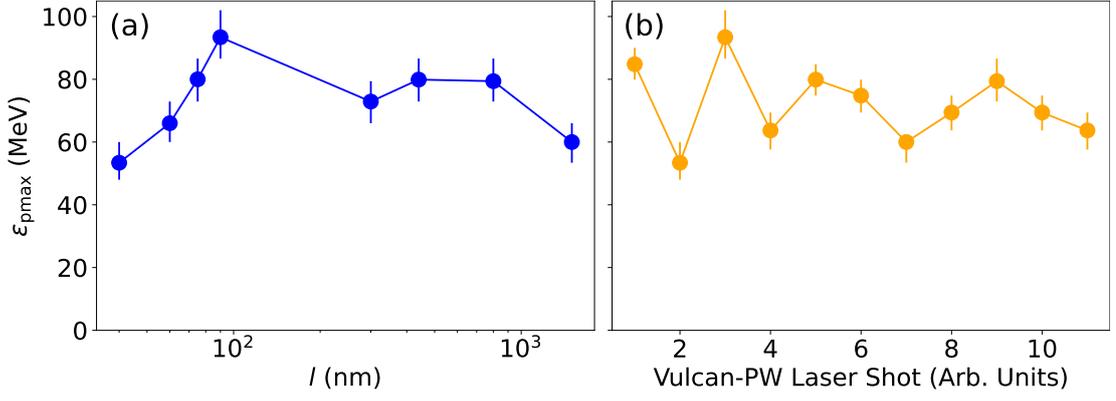


Figure 5.2: (a) Absolute maximum proton energy from a range of Vulcan-PW laser shots onto CH polymer target foils of varying thickness,  $l$ . (b) Maximum proton energy from 11 repeat Vulcan-PW laser shots onto CH target foils of the ‘optimal’,  $l=90$  nm, thickness.

Furthermore, in the simulations, the target foil of optimal thickness was demonstrated to undergo RSIT just as it was irradiated by the peak laser intensity. In contrast, RSIT was induced by the rising edge of the laser temporal-intensity profile in target foils thinner than the optimum, due to their reduced areal electron density, diminishing the maximum proton energy. For target foils thicker than the optimum the opposite is true. Here, the increased areal electron density delayed RSIT onset until after the laser intensity had peaked, reducing the maximum proton energy compared to that produced from the target foil of optimum thickness. This was the main result discussed in the publication, which reported the highest measured proton energy for each target foil thickness. There was a significant spread in the measured maximum proton energy over 11 repeat shots onto foils of the ‘optimal’,  $l=90$  nm, thickness, as shown in figure 5.2(b). This corresponds to a maximum proton energy range of 44%, from the optimum  $\epsilon_{pmax} \sim 96$  MeV, to a low of  $\epsilon_{pmax} \sim 53$  MeV, despite the fixed target foil thickness.

Clearly these results indicate that the maximum proton energy achieved in an RSIT-enhanced regime is sensitive to more than just the thickness of the target foil. Other publications have shown variations in the laser temporal-intensity profile to influence the maximum proton energy in this regime [45, 79], and this is a possible explanation for the variation shown here. However, this parameter can not be measured before every interaction, as this requires a third-order scanning autocorrelator to analyse a series of laser shots, which is not feasible given the low

repetition rate of the Vulcan-PW laser. As such, it is prudent to consider other laser variables that are known to significantly influence interaction dynamics. Among the most influential are the laser energy, pulse duration, and focal spot size, which were measured for this experiment as outlined in section 5.2.1. Figure 5.3 shows the value of each parameter for every interaction where they were successfully measured, as well as the laser intensity calculated from these values using equation 3.4.

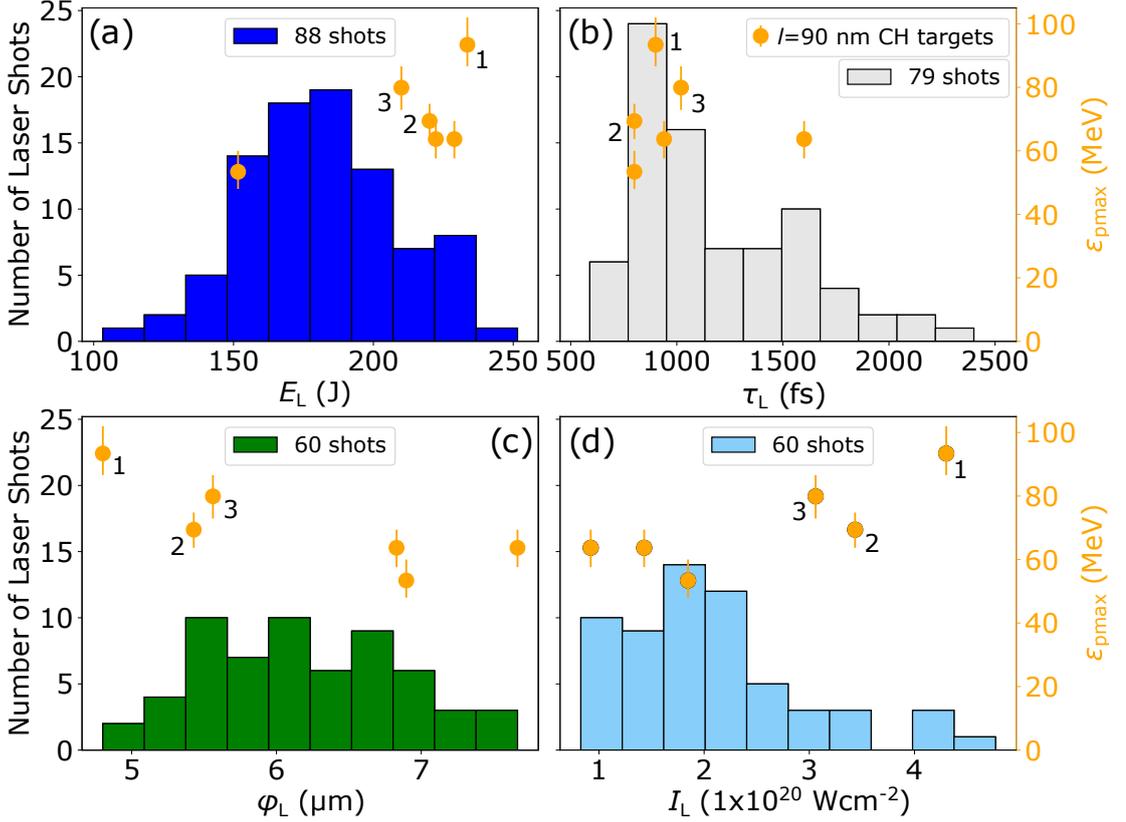


Figure 5.3: Histograms of the (a) energy, (b) pulse duration, (c) focal spot size and (d) intensity delivered by the Vulcan-PW laser during experiment A. The number of Vulcan-PW shots where an individual parameter was measured successfully is shown in each case alongside 6 shots onto  $l=90$  nm CH target foils, where every parameter was successfully measured.

Despite there being no deliberate efforts to change these parameters throughout the experiment, there is significant fluctuation in each, as shown in figure 5.3. This is demonstrated further in table 5.1, where the range of laser intensity values divided by the mean is 190%, highlighting significant fluctuation in this key parameter from one laser shot to another.

Much of this fluctuation can be attributed to the high energy nature of the

Laser Parameter	Min	Max	Mean $\pm$ S.Dev	S.Dev/Mean (%)	Range/Mean (%)
$E_L$ (J)	103	251	$180 \pm 30$	15	82
$\tau_L$ (fs)	590	2400	$1200 \pm 400$	33	152
$\phi_L$ ( $\mu\text{m}$ )	4.8	7.7	$6 \pm 0.7$	11	46
$I_L$ ( $10^{20}$ $\text{Wcm}^{-2}$ )	0.8	4.8	$2 \pm 0.9$	42	191

Table 5.1: Statistics describing Vulcan-PW laser parameter fluctuations during experiment A

system, which requires amplifying crystals in the beam-line to be heated to very high temperatures, introducing thermal gradients which can cause fluctuation in the output energy. Thermal gradients can also cause the beam to be misaligned, particularly through the compressor, which can influence the laser pulse duration and induce spatial distortions which affect the beam wavefront, and thus the focal spot size and encircled energy. To mitigate these effects, optics in the system are left to cool for a minimum of 30 minutes between shots. This combination of a low repetition rate, and significant fluctuation in key laser parameters, makes it difficult to obtain large data-sets with consistent interaction conditions, so that individual parameters can be intentionally varied to investigate their influence.

This lack of statistics is made worse if laser diagnostics are not filtered adequately on a given shot, so that data is saturated and the parameters cannot be measured. This is demonstrated in figure 5.3(b-c), where the laser pulse duration and focal spot size are not measured on every shot, due to saturation of the second-order autocorrelator, and of the Shack-Hartmann wavefront sensor used to measure the wavefront defocus. The result is that the data-set of 11 shots onto  $l=90$  nm CH target foils, shown in figure 5.2(b), effectively halves to 6 shots where all parameters were successfully measured, as shown in figure 5.3.

Analysing this data, the highest maximum proton energy,  $\epsilon_{\text{pmax}} \sim 96$  MeV, resulted from a shot with the highest laser energy, smallest focal spot, close to the shortest pulse duration, and thus the highest laser intensity, marked by a ‘1’ on each panel of figure 5.3. However, consider the shot marked ‘2’, where the laser energy, pulse duration, and focal spot size are close to the highest, shortest, and smallest in the dataset, resulting in the second highest laser intensity. The maximum proton energy is significantly lower on this shot, not only when compared to

the peak value, but also when compared to that resulting from shot ‘3’, where every laser parameter contributes to a lower intensity. Understanding this requires many more data points to build an average across the domain. To this end, a further experiment, experiment B, was carried out. Target foils of a fixed ‘optimal’ thickness were used throughout, so that the sensitivity of proton acceleration to fluctuations in key Vulcan-PW laser parameters could be investigated.

### 5.3.2 Investigating the stability of proton acceleration in the TNSA and RSIT-enhanced regimes

The beam-line was set-up for experiment B in the same way as for experiment A. Initially, a series of Vulcan-PW laser shots were directed onto CH target foils of varying thickness in the range,  $l=(30 \rightarrow 200)$  nm. The maximum proton energy from each interaction was recorded using an RCF stack, and is shown as a function of target thickness in figure 5.4.

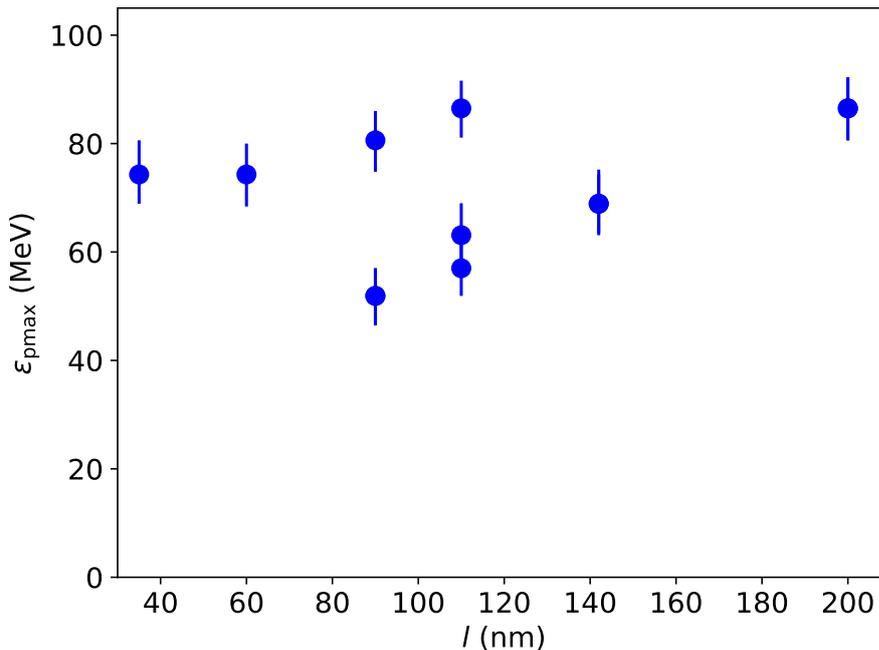


Figure 5.4: Maximum proton energy recorded for Vulcan-PW laser shots onto CH target foils of varying thickness,  $l$ , during experiment B.

Similar to the results of experiment A, optimum proton acceleration occurs from a  $l=110$  nm CH target foil. Protons of a similar energy were recorded from an interaction with a  $l=200$  nm CH target foil, however  $l=110$  nm is closer to the

$l=90$  nm CH target foil from which the highest energy protons were accelerated in experiment A.  $l=110$  nm CH target foils were therefore expected to facilitate ‘optimal’ RSIT onset time relative to the arrival of the peak laser intensity again in this experiment. As a result, of the 56 shots taken over the remaining time at the facility, 27 were fired onto  $l=110$  nm CH target foils. Added to the initial  $l=110$  nm shots, this makes for a data-set of 30 shots, enabling a detailed investigation of the stability of RSIT-enhanced proton acceleration for the first time. To compliment this data-set, there were also 9 shots onto  $l=6$   $\mu\text{m}$  aluminium target foils which do not undergo RSIT, facilitating pure TNSA proton acceleration. There were also 10 shots onto  $l=30$  nm CH target foils that undergo RSIT very early relative to the arrival time of the peak pulse intensity, rather than at approximately the same time, as has been demonstrated to optimise maximum proton energies. This enables comparison between the maximum proton energies which can be consistently produced from each target type.

As done for experiment A in figure 5.3, the laser energy, pulse duration, focal spot size, and intensity are shown for each interaction throughout experiment B in figure 5.5. The maximum proton energy is shown for 30 shots onto  $l=110$  nm CH target foils. Compared to experiment A, the laser energy during experiment B was higher on average, and the pulse duration and focal spot size are shorter and smaller respectively, resulting in a higher laser intensity. As before, the level of fluctuation in each of these variables is shown in more detail in table 5.2.

Laser Parameter	Min	Max	Mean $\pm$ S.Dev	S.Dev/Mean (%)	Range/Mean (%)
$E_L$ (J)	130	283	$220 \pm 30$	16	70
$\tau_L$ (fs)	450	1866	$900 \pm 300$	32	151
$\phi_L$ ( $\mu\text{m}$ )	4.5	6.8	$5 \pm 0.5$	10.7	44.5
$I_L$ ( $10^{20}$ $\text{Wcm}^{-2}$ )	1.4	7.2	$4 \pm 2$	34.4	137.2

Table 5.2: Statistics describing Vulcan-PW laser parameter fluctuations during experiment B.

The large data-set accumulated here enables the susceptibility of RSIT-enhanced proton acceleration to fluctuations in laser energy, pulse duration, focal spot size, and intensity, to be clearly demonstrated for the first time. This exacerbates known sensitivity of proton acceleration to target foil thickness and the laser

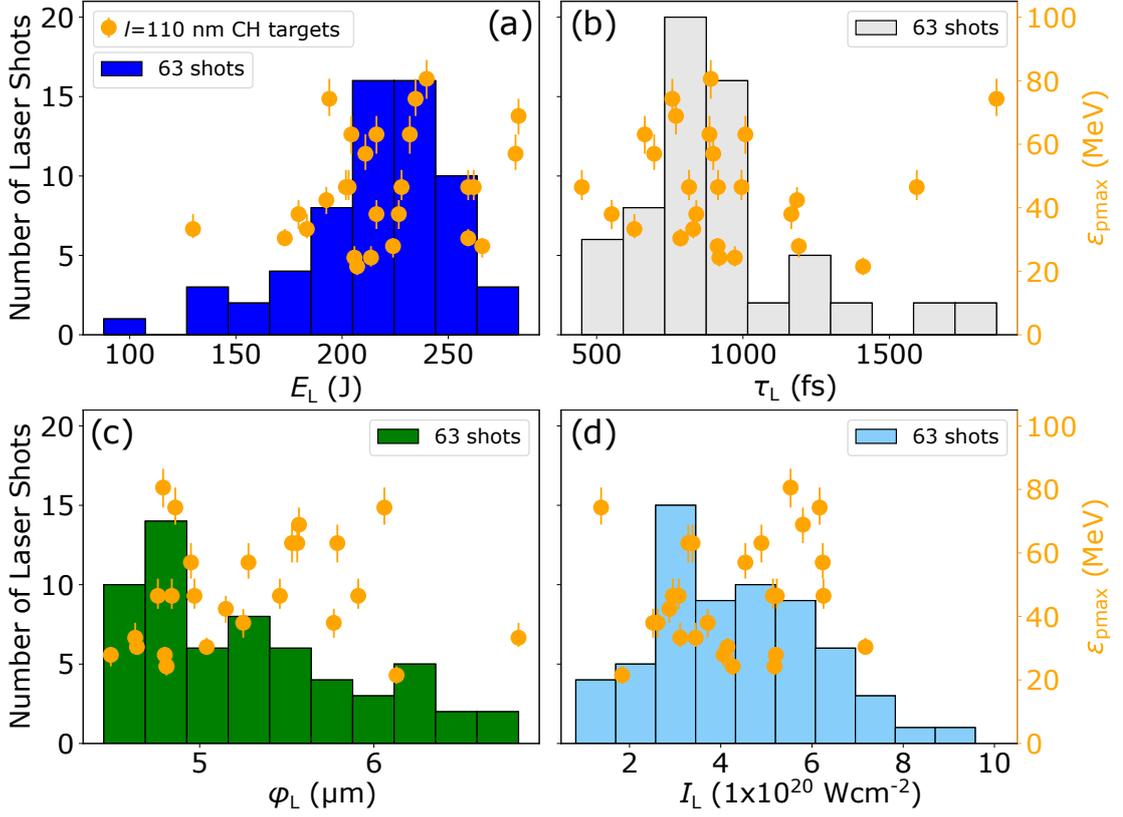


Figure 5.5: Histograms of the (a) energy, (b) pulse duration, (c) focal spot size and (d) intensity delivered by the Vulcan-PW laser during experiment B. The number of Vulcan-PW shots where an individual parameter was measured successfully is shown in each case alongside 30 shots onto  $l=110$  nm CH target foils where every parameter was successfully measured.

temporal-intensity profile within this regime, and likely explains the fluctuations in maximum proton energy observed for repeat interactions with target foils of fixed thickness on experiment A [71].

Initially, there is no obvious correlation between laser parameters and maximum proton energies in figure 5.5. However, there are clear outliers in the data-set, which will now be identified and removed. Maximum proton energies will then be compared for shots with similar laser conditions, so that the susceptibility of RSIT-enhanced proton acceleration to fluctuating laser parameters can be investigated.

### 5.3.3 Influence of laser energy and pulse duration fluctuations on maximum proton energies

The first outliers to be removed from the data-set of 30 shots onto  $l=110$  nm CH target foils shown in figure 5.5 are 10 shots where laser alignment through the compressor gratings was changed to increase the pulse duration. This change was made because there were indications that the maximum proton energy increased with pulse duration to some optimum ( $\tau_L \sim 900$  fs), before decreasing, as observed in figure 5.5(b). The measured pulse duration is shown in figure 5.6 for every shot with the compressor in its initial position, and for those where the compressor alignment was varied.

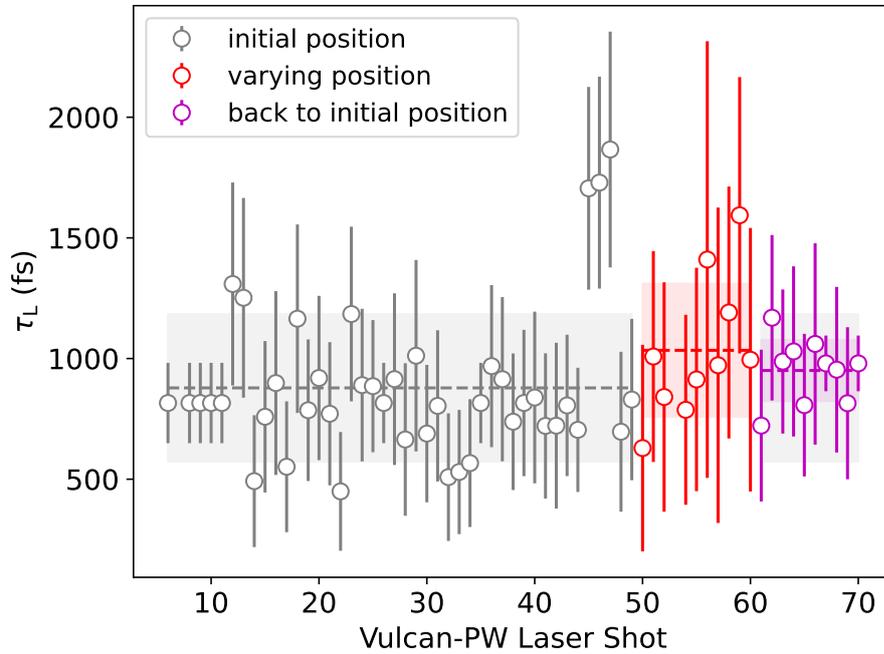


Figure 5.6: Pulse duration of the Vulcan-PW laser measured throughout experiment B. Dashed grey, red, and pink lines indicate the mean pulse duration for shots where optics in the compressor were at an initial position, moved to different positions, and then moved back to the initial position. Shaded regions represent the standard deviation in the mean measurements for each region and error bars indicate the standard deviation for individual measurements

The mean pulse duration was  $\tau_L = (900 \pm 300)$  fs with the compressor in the initial position, and  $\tau_L = (1000 \pm 300)$  fs on the shots where it was varied. For the latter, the compressor alignment was varied four times, by moving optics in the compressor to four different positions. The uncertainty in individual pulse duration measurements is extremely high for these shots, and so they are removed from

the data-set of shots onto  $l=110$  nm CH target foils. Compressor alignment was changed back to the initial position for the last 9 shots of the experiment, where the laser had a mean recorded pulse duration of  $\tau_L=(1000\pm 100)$  fs and was fired at  $l=30$  nm CH target foils.

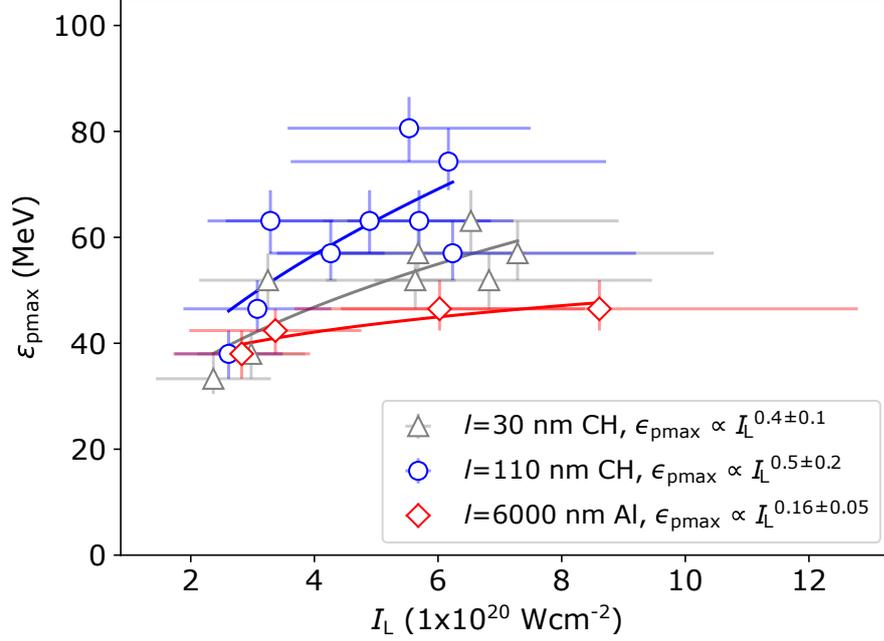


Figure 5.7: Maximum proton energy as a function of Vulcan-PW laser intensity for a range of shots onto  $l=30$  nm CH,  $l=110$  nm CH, and  $l=6$   $\mu$ m aluminium target foils. Error bars express the resolution of the RCF stack as the difference in energy between layers before and after the final RCF layer where protons deposit their dose (maximum proton energy). A power scaling of the form  $\epsilon_{pmax} \propto I_L^b$  is fit to each data-set

Removing this data leaves 20 shots onto  $l=110$  nm CH target foils. Analysing the remaining data, it was useful to look for outliers where the measured laser parameters were far from their respective average within the data-set. Such outliers are clearly observed in figure 5.6, where the measured laser pulse duration is greater than  $\pm 1$  standard deviation from the mean on four shots. Similarly, there are four shots where the laser energy is greater than  $\pm 1$  standard deviation from the mean, and two shots where the focal spot size is similarly far from the average. Removing these outliers leaves a final data-set of 10 shots onto  $l=110$  nm CH target foils. The maximum proton energy for each is shown as a function of laser intensity in figure 5.7. After handling the data in this way there is now a clear relationship between the maximum proton energy and the Vulcan-PW laser intensity, in contrast to the lack of correlation seen in figure 5.5. The maximum

proton energy is shown to scale with intensity as,  $\epsilon_{\text{pmax}} \propto I_{\text{L}}^{0.5 \pm 0.2}$ , similar to that reported for TNSA [30, 40, 41, 51, 214]. Also shown on this figure are two data-sets of shots onto  $l=30$  nm CH, and  $l=6$   $\mu\text{m}$  aluminium target foils. Outliers were removed from these data-sets in the same way as described for the  $l=110$  nm CH target foil data-set.

RSIT is not induced in the  $l=6$   $\mu\text{m}$  aluminium target foils, and so TNSA is expected to be the dominant proton acceleration mechanism responsible for maximum energies which scale with laser intensity as  $\epsilon_{\text{pmax}} \propto I_{\text{L}}^{0.16 \pm 0.05}$ . As stated, TNSA is expected to produce a much stronger energy scaling with intensity, of  $\epsilon_{\text{pmax}} \propto I_{\text{L}}^{0.5}$  [30, 40, 41, 51, 214]. The weaker scaling demonstrated may be due to the restricted laser intensity range, or the relatively low number of data points. However, it may also be true that increases to TNSA maximum proton energies saturate at these high laser intensities. This is the case when the intensity is increased by reducing the focal spot size [42].

Regardless of scaling with laser intensity, RSIT-enhanced acceleration from  $l=110$  nm target foils produces protons of higher maximum energy compared to TNSA from  $l=6$   $\mu\text{m}$  target foils which don't undergo RSIT. Higher maximum proton energies have been achieved via TNSA using thinner target foils than those used here [62, 63], though not to the levels achieved from ultra-thin foils undergoing RSIT [31, 71]. Despite this, RSIT-enhanced proton acceleration is less stable compared to TNSA, as demonstrated by the uncertainty in the respective intensity scalings of figure 5.7. To investigate this, pairs of shots with similar laser conditions are compared in table 5.3.

Shot	$E_{\text{E}}$ (J)	$\tau_{\text{L}}$ (fs)	$\phi_{\text{L}}$ ( $\mu\text{m}$ )	$I_{\text{L}}$ ( $10^{20}$ $\text{Wcm}^{-2}$ )	$\epsilon_{\text{pmax}}$ (MeV)
A	87	759	4.9	6.2	74.3
B	84	697	5.0	6.2	57.0
C	89	890	4.8	5.5	80.6
D	84	816	4.8	5.7	63.1
E	71	886	5.6	3.3	63.1
F	66	915	5.5	3.1	46.5

Table 5.3: Pairs of comparable Vulcan-PW laser shots onto  $l=110$  nm CH target foils.

For shots A and B, the laser focal spot size and intensity are the same, whilst the encircled energy ( $E_E$ ) and pulse duration vary. This leads to a higher maximum proton energy for the shot with a higher encircled energy and longer pulse duration. This is also the case when comparing shots C and D. For shots E and F, all parameters are very similar, and the higher encircled energy for shot E appears to increase the maximum proton energy compared to shot F.

A similar dynamic is observed when comparing shots of different laser intensity. Shots A and C have comparable encircled energy and focal spot size, though the pulse duration is longer for shot C, resulting in a higher maximum proton energy. The same is true when comparing shots B and D, suggesting that shots with a higher laser energy and longer pulse duration produce higher maximum proton energies in the RSIT-enhanced regime. This provides more evidence to support the trend indicated in figure 5.5, that maximum proton energies increase with laser pulse duration to an optimum, before decreasing. Furthermore, these laser energy and pulse duration fluctuations between pairs of shots in table 5.3 are relatively minor ( $\sim 8\%$ ), yet they result in maximum proton energy variation of  $\sim (22-26)\%$ . Given this sensitivity to relatively small fluctuations in laser parameters, the observed variation in maximum proton energy when parameters vary much more significantly over the entire  $l=110$  nm data-set is hardly surprising.

It should be noted that factors such as the laser temporal-intensity profile cannot be measured shot-to-shot, and may be changing between interactions, causing maximum proton energy variation. Despite this, the data presented in table 5.3 demonstrates large variability in RSIT-enhanced proton energies for small fluctuations in key laser parameters. This builds on evidence from experiment A, presented in figure 5.2 and figure 5.3. To investigate this further, 2D PIC simulations were run using EPOCH [171], with comparable laser energy and pulse duration fluctuations as the experimental data presented in table 5.3. The results are compared to these experimental data points in the next section.

### 5.3.4 Comparing the maximum proton energy sensitivity between experiment and simulations

Most of the experimental data discussed in the previous sections result from measurements made using the Vulcan-PW laser to irradiate ultra-thin CH target foils. These targets are now replicated using the fully relativistic 2D EPOCH PIC code [171] by defining an equal mixture of  $H^+$  and  $C^{6+}$  ions, and a neutralising electron population with an initial density of  $n_e=349n_{\text{crit}}$ , where  $n_{\text{crit}}$  is the plasma critical density, defined in section 2.4.2. Both ion populations had an initial temperature of 0.01 keV, and the electrons had a higher initial temperature of 1 keV due to their relatively lower mass. There were also experimental shots onto aluminium target foils. These targets are replicated similarly in EPOCH, by defining an equal mixture of  $Al^{11+}$  ions, and a neutralising electron population with an initial density of  $n_e=660n_{\text{crit}}$ . For clarity,  $Al^{11+}$  is used because this is the highest charge state of aluminium ions detected in the experiments, not  $Al^{13+}$ , suggesting aluminium atoms are not fully ionised at present laser intensities. Experiments have also shown that, for metal target foils, the highest energy protons are accelerated from rear surface hydrocarbon contaminants. To replicate this, a thin,  $l=12$  nm, layer of  $H^+$  and  $C^{6+}$  ions was added to the rear surface of the aluminium target in EPOCH.

The grid dimensions for each simulation were  $18,500 \times 3,456$  cells, corresponding to a domain of  $(-30 \rightarrow 160)$   $\mu\text{m}$  and  $(-25.5 \rightarrow 55.1)$   $\mu\text{m}$ , and a resolution of  $\sim 10$  nm/cell and  $\sim 23$  nm/cell in the  $X$  and  $Y$  dimensions, respectively. Targets were angled at  $30^\circ$  with respect to the incident laser, as they were in the experiment. The laser wavelength was  $\lambda_L=1053$  nm, and the focal spot size was fixed as  $\phi_L=5.2$   $\mu\text{m}$ , the average measurement across experiment B. Each simulation took  $\sim 2$  hours to run using the EPOCH 2D PIC code on 2304 of the CPUs provided by the ARCHER2 HPC cluster [293].

The laser pulse duration fluctuated from  $\tau_L \sim 900$  fs to  $\tau_L \sim 600$  fs on experiment B, however it is not feasible to simulate interactions with such a long pulse duration due to the substantial computational resources required to do so. Instead a

range of  $\tau_L=(253-400)$  fs was used in the simulations to represent an experimental range of  $\tau_L\sim(600-900)$  fs, matching the magnitude of fluctuation, and facilitating comparison between simulation and experimental data. As discussed in chapter 4, the entire laser energy defined in these simulations is contained within the  $\phi_L=5.2$   $\mu\text{m}$  focal spot, in contrast to the experiments where only  $\sim 40\%$  of the laser energy delivered to the OAP is focused into a spot of equivalent size. The laser energy range for the simulations was therefore effectively the equivalent of the experimental encircled energy, selected as  $E_L=(24-52)$  J by rearranging equation 2.2 so that the laser intensity range covered experimentally,  $I_L=(2-9)\times 10^{20}$   $\text{Wcm}^{-2}$ , was matched for the shorter simulated pulse durations,  $\tau_L=(253-400)$  fs.

Experimental CH target foils of  $l=30$  nm,  $l=110$  nm, and  $l=6$   $\mu\text{m}$  were replicated by target slabs of solid-density plasma, of thickness  $l=25$  nm,  $l=95$  nm, and  $l=1$   $\mu\text{m}$ , in respective simulations. The ultra-thin targets are slightly thinner than the experimental equivalents, so that the RSIT onset time is close to the same in simulation and experiment, since targets undergo RSIT later in the simulations due to the comparatively shorter laser pulse duration. For the thickest target there is a considerable difference, since significant computational resources would be required to simulate a  $l=6$   $\mu\text{m}$  target. Importantly, simulations at  $l=1$   $\mu\text{m}$  still result in a TNSA dominant interaction where no RSIT occurs, accurately replicating the experimental conditions for the purposes of this investigation.

For the first set of simulations, the laser parameters were fixed as  $E_L=43$  J,  $\tau_L=400$  fs,  $\phi_L=5.2$   $\mu\text{m}$ , and  $I_L=5.1\times 10^{20}$   $\text{Wcm}^{-2}$ , and the target thickness was varied over the range  $l=(50-110)$  nm to find the ‘optimal’ target thickness, as was done experimentally. The *Vector* mode of BISHOP was used to generate these simulations as described in chapter 4, and the maximum proton energy was determined for each as the mean of the most energetic 1% of protons, to avoid outliers caused by only considering 1 proton with the highest energy. Three repeat simulations were performed at each target thickness, and the mean maximum proton energy is shown as a function of target thickness in figure 5.8, with error bars representing the associated standard deviation. The maximum proton energy peaks at a target thickness of  $l=95$  nm, replicating the trend from experiments A

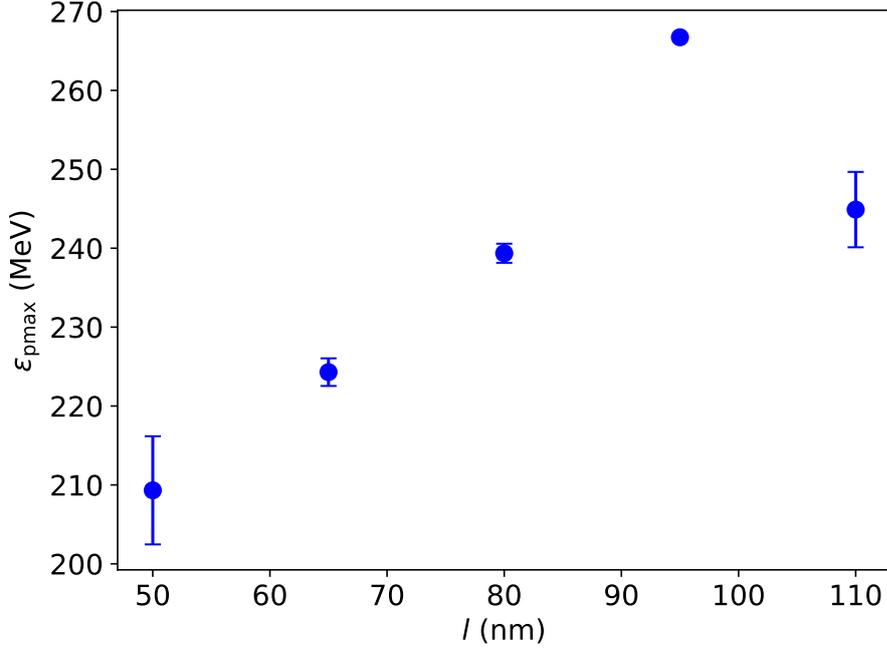


Figure 5.8: Maximum proton energy as a function of target thickness in 2D PIC simulations.

and B. This is thinner than the optimal target foil thickness of  $l=110$  nm in the data-set from experiment B, which is expected, for the reasons already discussed.

Now that the optimal target thickness is defined, a second series of simulations were performed where the laser pulse duration and focal spot size were kept fixed as  $\tau_L=400$  fs and  $\phi_L=5.2$   $\mu\text{m}$  respectively, and the laser energy was varied over the range  $E_L=(24-52)$  J, to replicate the laser intensity range explored experimentally via the inherent shot-to-shot fluctuation of the Vulcan-PW laser,  $I_L=(2.9-6.2)\times 10^{20}$   $\text{Wcm}^{-2}$ , when fired onto  $l=110$  nm CH target foils. Three simulations were performed for each encircled energy (intensity) point, to obtain a mean maximum proton energy, and associated standard deviation measurement for each. These measurements are shown as a function of laser intensity in figure 5.9(a). Maximum proton energy values are known to be overestimated by  $\sim 3$  times in 2D PIC simulations compared to experimental results [283–285], and so values from these simulations are divided by 3.9 to enable comparison with the experimental results in this figure.

Also shown on figure 5.9(a) is a third simulation dataset which is the the same as the second but for the laser pulse duration being  $\sim 8\%$  shorter at  $\tau_L=367$  fs. This replicates the fluctuation between pairs of Vulcan-PW shots detailed

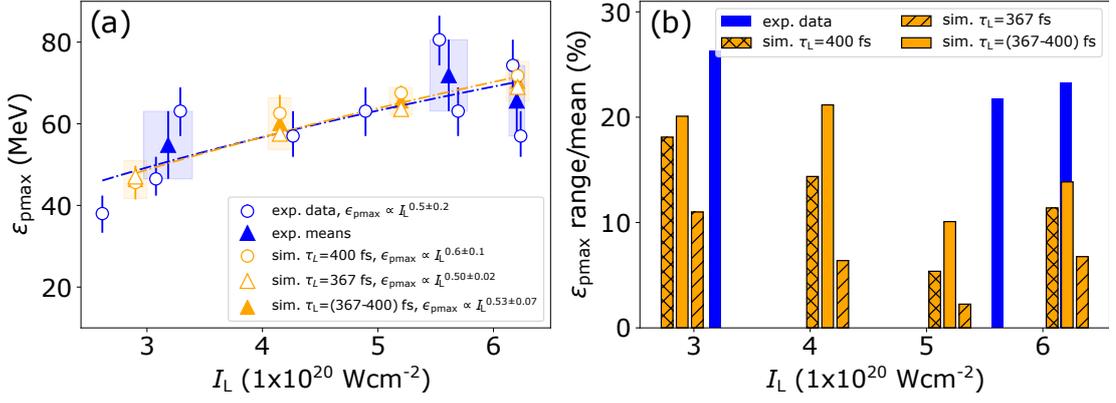


Figure 5.9: (a) Mean maximum proton energy across repeat simulations with a  $l=95$  nm CH target and a laser pulse duration of  $\tau_L=400$  fs,  $\tau_L=367$  fs, and  $\tau_L=(367-400)$  fs, as a function of laser intensity. An  $\epsilon_{pmax} \propto I_L^b$  scaling is shown for each pulse duration case. The range of maximum proton energy values across the repeat simulations is expressed as orange error bars on the mean for  $\tau_L=400$  fs and  $\tau_L=367$  fs, and as a shaded orange region for the combined  $\tau_L=(367-400)$  fs data-set. Values from the simulations are divided by 3.9 for comparison to the maximum energy of protons accelerated from  $l=110$  nm CH target foils irradiated by the Vulcan-PW laser. The mean and range of energies across three pairs of experimental points with comparable laser conditions is shown as a blue triangle and shaded region respectively. (b) shows the range of maximum proton energy values as a percentage of the mean for each of the simulation data-sets and the comparable experimental points

in table 5.3. The maximum proton energies produced on these pairs of shots is averaged and shown in figure 5.9(a). The laser energy for this simulation data-set is correspondingly  $\sim 8\%$  lower, in the range  $E_L=(22-48)$  J rather than  $E_L=(24-52)$  J, so that the laser intensity range is matched in the  $\tau_L=367$  fs and  $\tau_L=400$  fs simulations. Finally, the average maximum proton energy is also shown across the combined six simulations,  $\tau_L=(367,400)$  fs, at each laser intensity point. A power scaling of  $\epsilon_{pmax} \propto I_L^{0.53 \pm 0.07}$  is found for the combined data-set, in good agreement with the experimental scaling of  $\epsilon_{pmax} \propto I_L^{0.5 \pm 0.2}$ . In addition, the maximum proton energy is generally higher for the  $\tau_L=400$  fs simulations compared to the  $\tau_L=367$  fs simulations, replicating the experimental trend presented in table 5.3.

Figure 5.9(a) also shows the range of maximum proton energies across repeat simulations at each laser intensity point, on the individual and combined simulation data-sets. These values are shown as a percentage of the mean maximum proton energy at each laser intensity point in figure 5.9(b), as is the range of maximum proton energies across comparable experimental data points from table 5.3. The laser intensity has been shifted by  $-1 \times 10^{19}$  Wcm $^{-2}$  and  $+1 \times 10^{19}$  Wcm $^{-2}$  from the true simulated intensity  $\epsilon_{pmax}$  on this plot for the  $\tau_L=400$  fs and  $\tau_L=367$  fs

data respectively, so that the bars are visible without overlapping those of the combined data-set.

The maximum proton energy varies by (5-18)% for  $\tau_L=400$  fs and (2-11)% for  $\tau_L=367$  fs over the intensity range for repeat simulations with identical conditions, the only difference being that the initial particle distribution is randomised by enabling the random seed in EPOCH, as discussed in section 4.6.5. This demonstrates high inherent uncertainty in simulations of RSIT-enhanced proton acceleration, even when the interaction conditions are identical. This is explored in detail in section 5.3.5.

Furthermore, these simulations show that small (8%) variations in the laser pulse duration and energy result in more significant maximum proton energy variation of (10-21)%. As such, it is little surprise that the maximum proton energy varies by (22-26)% for experimental data points where the measured interaction parameters fluctuate by a similar 8%, but other variables are not measured and could be fluctuating. As an example, the experimental data points compared in table 5.3 and figure 5.9 were taken on different days, between which the laser was shut down and restarted many times, potentially changing the temperature, alignment, beam profile, and temporal-intensity profile. In addition, though the target foil thickness was nominally constant throughout the experiment, there is  $\sim 10\%$  uncertainty in the thickness measurement. Suffice to say that one, or a combination of these variables, could increase the maximum proton energy variation from up to 21% in controlled simulations where the laser energy and pulse duration were varied by 8%, to (22-26)% as recorded for experimental interactions with comparable fluctuations in these parameters.

Another potential explanation for the increased variation in experiment compared to simulations is uncertainty in the measured experimental laser conditions. In particular, there is significant uncertainty in the measured laser pulse duration, as shown in figure 5.6. For the experimental data presented in figure 5.9, the mean uncertainty (standard deviation) is  $\pm 292$  fs, corresponding to a difference of 35% from the mean value,  $\tau_L=845$  fs. With such uncertainty in these measurements, it is prudent to investigate the effect of varying laser pulse duration over

a larger range than in figure 5.9.

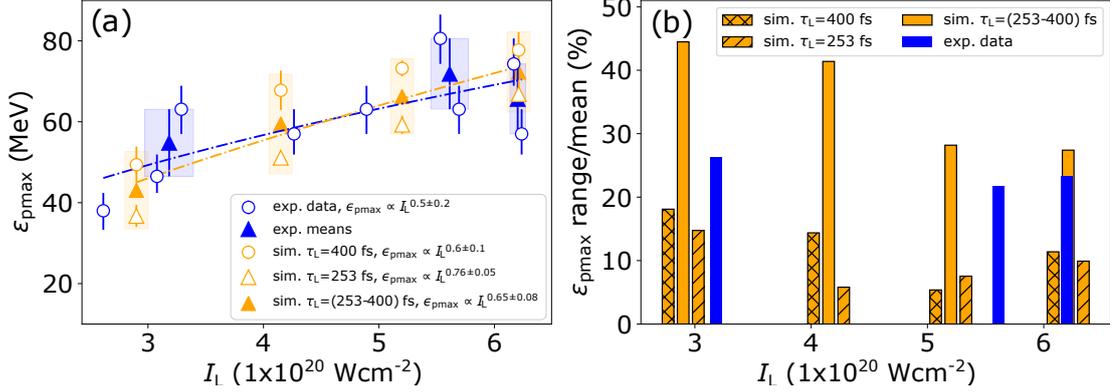


Figure 5.10: Same as figure 5.9 but now with simulations at  $\tau_L=253$  fs rather than  $\tau_L=367$  fs. Values from the simulations are divided by 3.6 for comparison to the maximum energy of protons accelerated from  $l=110$  nm CH target foils irradiated by the Vulcan-PW laser.

To this end, a fourth set of simulations were run that were identical to those already presented, except that the laser pulse duration was now set as  $\tau_L=253$  fs, 35% shorter than the initial  $\tau_L=400$  fs simulations. As before, the laser energy was reduced to a range of  $E_L=(15-33)$  J, to maintain the laser intensity between the simulation data-sets. Again, three simulations were performed at each laser intensity point, and the mean and range of maximum proton energy values are shown as a function of laser intensity in figure 5.10.

As before, the mean and range of maximum proton energy values was also calculated over the combined simulations,  $\tau_L=(253-400)$  fs, and a power scaling of  $\epsilon_{pmax} \propto I_L^{0.65 \pm 0.08}$  was observed, within the error of that observed experimentally,  $\epsilon_{pmax} \propto I_L^{0.5 \pm 0.2}$ . Maximum proton energy values from the simulations were divided by 3.6 to be comparable in magnitude to the experimental data in figure 5.10(a), rather than being divided by 3.9 to be comparable in the same way in figure 5.9(a), given the lower energies at  $\tau_L=253$  fs, compared to  $\tau_L=367$  fs. The most notable result from these simulations is that the maximum proton energy varies much more significantly due to the increased 35% variation in laser pulse duration (and encircled energy), now varying by (27-44)%, dwarfing the variation observed for experimental shots where the pulse duration varies by  $\sim 8\%$ , as shown in figure 5.10(b). Notably, at the highest laser intensities,  $I_L=(5.2, 6.2) \times 10^{20} \text{ Wcm}^{-2}$ , the maximum proton energy range is 27% and 28% in simulations, similar to that

observed for pairs of experimental shots at these intensities (22% and 23%).

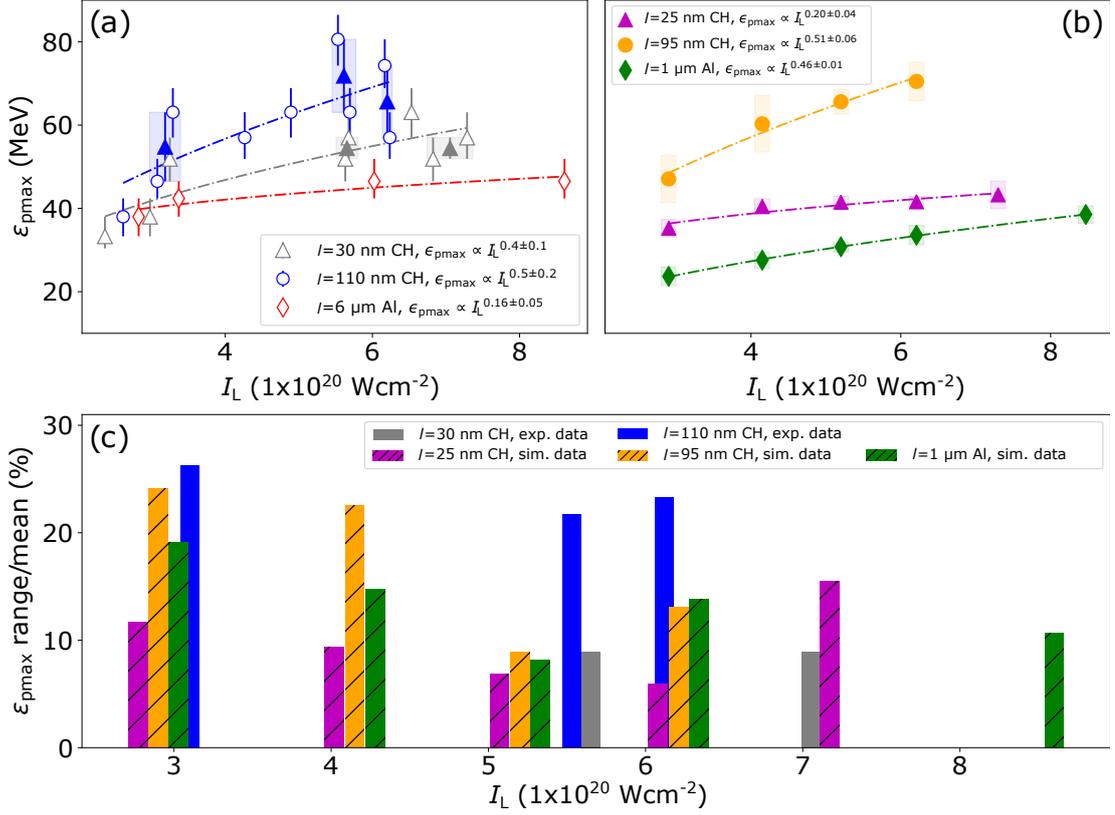


Figure 5.11: (a) Maximum proton energy as a function of the Vulcan-PW laser intensity focused onto  $l=30$  nm CH,  $l=110$  nm CH, and  $l=6$   $\mu$ m aluminium target foils. The mean and range of maximum proton energies across pairs of Vulcan-PW shots with comparable laser properties is illustrated by filled triangles and shaded regions respectively for  $l=30$  nm (grey), and  $l=110$  nm (blue) CH target foils. (b) shows the mean maximum proton energy (divided by 3.9) across repeat simulations with  $\tau_L=(367-400)$  fs as a function of intensity for  $l=25$  nm and  $l=95$  nm CH targets, and  $l=1$   $\mu$ m aluminium targets. The range of maximum proton energies for these repeat simulations is illustrated by shaded regions for each intensity point, and is shown as a percentage of the mean maximum proton energy in (c) alongside the range of maximum proton energies for the pairs of experimental data points shown in (a), expressed as a percentage of their respective means.

Now that the susceptibility of maximum proton energies to laser parameter fluctuations has been explored in a regime of RSIT-enhanced acceleration from target foils of an ‘optimal’ thickness, it will now be investigated for acceleration from  $l=30$  nm CH and  $l=6$   $\mu$ m aluminium target foils, for which there was less maximum proton energy variation on experiment, as shown in figure 5.7, and highlighted again in figure 5.11(a). Here, the maximum proton energy produced from two pairs of comparable Vulcan-PW laser shots onto  $l=30$  nm CH target foils is shown, adding to that compared previously over three pairs of shots onto

$l=110$  nm CH target foils. For the lower laser intensity shots onto  $l=30$  nm CH target foils, the laser parameter values were  $E_E=(86, 87)$  J,  $\tau_L=(980, 988)$  fs, and  $\phi_L=(4.5, 4.5)$   $\mu\text{m}$ , representing a fluctuation of  $\sim 1\%$  in each parameter. The maximum proton energy produced by these shots varies by 10%, from  $\epsilon_{\text{pmax}}=52$  MeV to  $\epsilon_{\text{pmax}}=57$  MeV, as a result. For the pair of higher laser intensity shots onto  $l=30$  nm CH target foils, the parameter values were  $E_E=(91, 92)$  J,  $\tau_L=(722, 815)$  fs, and  $\phi_L=(4.7, 4.6)$   $\mu\text{m}$ , representing a fluctuation of  $\sim 1\%$  in the encircled energy and focal spot size, and of  $\sim 11\%$  in the pulse duration. Again, this results in the maximum proton energy varying by 10%, from  $\epsilon_{\text{pmax}}=57$  MeV to  $\epsilon_{\text{pmax}}=52$  MeV. Unfortunately, no Vulcan-PW laser shots with comparable properties onto  $l=6$   $\mu\text{m}$  target foils were recorded during experiment B, though the susceptibility of TNSA maximum proton energies produced from these target foils will now be investigated in simulations, alongside that of RSIT-enhanced acceleration from the ultra-thin target foils.

To do so, the same simulations were run as before, this time with a plasma region of  $l=25$  nm, and  $l=1$   $\mu\text{m}$ , to replicate the  $l=30$  nm CH target foils and  $l=6$   $\mu\text{m}$  aluminium target foils used in the experiment. As before, the mean and range of maximum proton energies over the combined simulations,  $\tau_L=(367-400)$  fs, is shown over the laser intensity range for the respective targets in figure 5.11(b), alongside the same data for the  $l=95$  nm CH targets.

As in the experimental data, maximum proton energies are highest across the laser intensity range for  $l=95$  nm CH targets. This confirms that proton energies are optimised when RSIT is induced just as the peak laser intensity irradiates the target. The maximum energy of protons accelerated from  $l=25$  nm CH targets is lower, and scales with laser intensity as  $\epsilon_{\text{pmax}} \propto I_L^{0.25 \pm 0.07}$ , within the error of the experimental equivalent,  $\epsilon_{\text{pmax}} \propto I_L^{0.4 \pm 0.1}$ . The maximum energy of protons accelerated from  $l=1$   $\mu\text{m}$  CH targets is lower still. In this case the observed scaling in simulations is  $\epsilon_{\text{pmax}} \propto I_L^{0.70 \pm 0.02}$ , quite different to the  $\epsilon_{\text{pmax}} \propto I_L^{0.16 \pm 0.05}$  scaling demonstrated in experiment B. This is likely due to limited experimental data points, or the large difference in target thickness between simulation and experiment. As before, the range of maximum proton energies produced when the

laser pulse duration varies from  $\tau_L=400$  fs to  $\tau_L=367$  fs in the simulations is shown in figure 5.11(c), for each intensity point, and now for each target thickness. The maximum energy variation of protons accelerated from three pairs of Vulcan-PW laser shots with comparable properties onto  $l=110$  nm CH target foils is shown as before, and is now shown for both pairs of comparable Vulcan-PW laser shots onto  $l=30$  nm CH target foils. Individual maximum proton energies for each of these shots onto  $l=30$  nm CH target foils is shown as a function of laser intensity in figure 5.11(a).

As in the data from experiment B, maximum proton energies vary most significantly across the laser intensity range in simulations of RSIT-enhanced acceleration from  $l=95$  nm CH targets, where the RSIT time is optimised to produce the highest proton energies. Significantly, variations in maximum energy are smallest for RSIT-enhanced proton acceleration from  $l=25$  nm CH targets in simulations, and are comparable to those observed for this regime in experiment B. Also interesting is that this variation generally decreases with increasing laser intensity in the simulations. These results indicate that earlier onset of RSIT in thinner target foils, irradiated at higher laser intensities, results in improved shot-to-shot stability in the RSIT-enhanced proton acceleration mechanism. This will be investigated as part of a programme of further work, outlined in section 5.4.

Having demonstrated RSIT-enhanced proton energies to vary in response to fluctuating laser encircled energy and pulse duration, this will now be explored in more detail in the following section. Also explored are significant inherent fluctuations of maximum proton energies in this acceleration regime, indicated on repeat simulations which are identical in all interaction variables, with only the initial random distribution of plasma changing between simulations.

### **5.3.5 Diagnosing the sensitivity of RSIT-enhanced proton acceleration to fluctuating laser conditions**

The results presented in section 5.3.4 demonstrate significant maximum energy variation of protons accelerated from  $l=95$  nm CH targets undergoing RSIT,

between simulations where the laser pulse duration changes by  $\sim 8\%$ , from  $\tau_L=400$  fs to  $\tau_L=367$  fs, and by  $\sim 35\%$ , from  $\tau_L=400$  fs to  $\tau_L=253$  fs. To investigate why this occurs, the temporal evolution of the mean and range of maximum proton energy values over three simulations at  $\tau_L=400$  fs, and three simulations at  $\tau_L=253$  fs, is shown in  $t=100$  fs increments, in figure 5.12. The laser energy is  $E_L=24$  J, and  $E_L=15$  J, in the respective simulation sets, so that the peak laser intensity is conserved at  $I_L=3\times 10^{20}$  Wcm $^{-2}$ . The acceleration of protons is also shown as the difference in maximum energy between the simulation timesteps. Protons are accelerated earlier in the longer pulse duration simulations, due to the prolonged laser intensity rising edge profile. This results in higher maximum proton energies which, despite a spike in proton acceleration at  $t\sim(100\rightarrow 300)$  fs in the shorter,  $\tau_L=253$  fs, simulations, remains higher for the remainder of the  $\tau_L=400$  fs simulations. In particular, the prolonged falling edge intensity profile of the longer duration pulse results in a considerable difference in acceleration from  $t\sim(400\rightarrow 800)$  fs, and thus maximum proton energy by the end of the respective simulations.

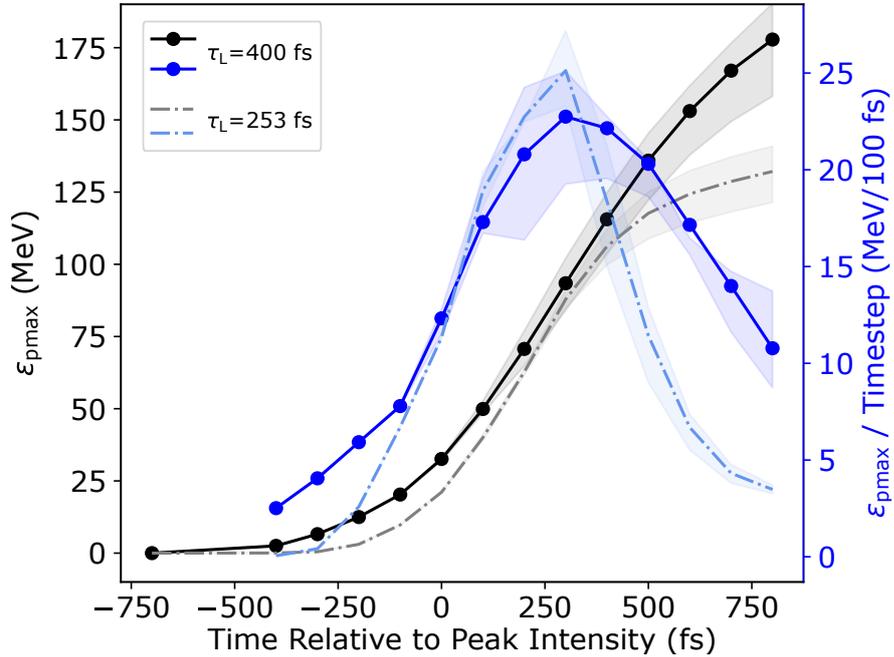


Figure 5.12: Maximum proton energy at each timestep, and difference in maximum proton energy between timesteps (acceleration), both averaged across three repeat simulations with  $\tau_L=400$  fs, and three repeat simulations with  $\tau_L=253$  fs. The range of acceleration and maximum energy is shown as a shaded region for the respective repeat simulations at each timestep.

Also apparent from figure 5.12 is that proton acceleration, and resultant maximum energies, vary significantly over three exact repeat simulations at the same pulse duration. This is most noticeable from  $t \sim (100 \rightarrow 400)$  fs, and from  $t \sim (700 \rightarrow 800)$  fs, in the repeat  $\tau_L = 400$  fs simulations. To understand this, figure 5.13 shows the maximum proton energy, proton acceleration, and mean longitudinal electric field experienced by the highest energy protons for times,  $t = (100 \rightarrow 800)$  fs, in each of these repeat simulations. Protons initially experience a stronger longitudinal electric field in repeat simulation B, compared to simulations A and C, thus being accelerated more readily according to equation 2.16, and reaching higher maximum energies over times  $t = (100 \rightarrow 300)$  fs. However, as the simulations progress, protons experience a stronger longitudinal electric field in simulation C, resulting in more significant acceleration, and higher maximum energies by the end of this simulation, compared to simulations A and B.

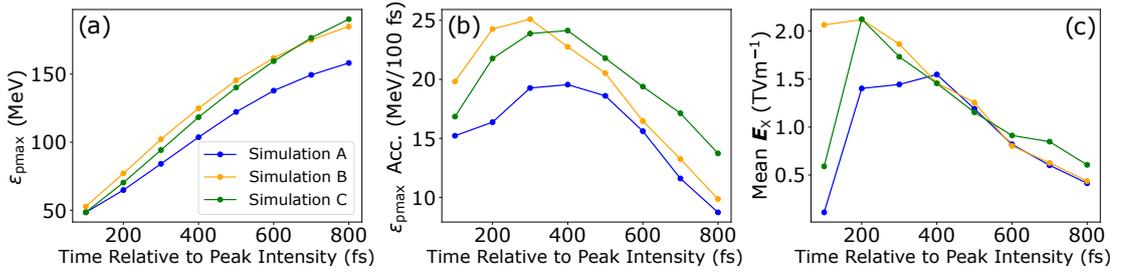


Figure 5.13: Proton (a) maximum energy and (b) acceleration resulting from an (c) average longitudinal electric field at each timestep of three repeat simulations with identical conditions, but for the initial plasma macroparticle distribution which is randomly varied.

To understand why these differences are arising for nominally identical simulations, the longitudinal electric field in a region of  $X = (-1 \rightarrow 1)$   $\mu\text{m}$  by  $Y = (-1 \rightarrow 1)$   $\mu\text{m}$  around these protons is illustrated in figure 5.14. The fields look very similar for most of the simulations, though at the final times of  $t = (700, 800)$  fs, the field is almost entirely positive in simulation C, compared to the other simulations where there are some neutral regions. This results in a significant difference in the mean accelerating field, slightly increasing the maximum energy of protons by the end of simulation C, compared to the other simulations, as shown in figure 5.13.

As no laser or target parameters are changing across these simulations, dif-

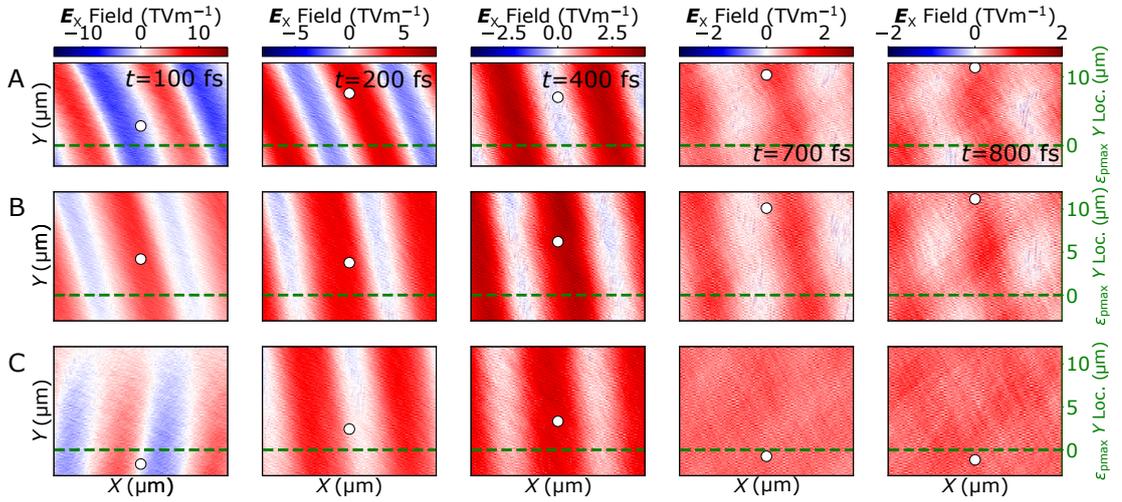


Figure 5.14: Illustrations of the longitudinal electric field in a  $(X, Y) = (-1 \rightarrow 1)$   $\mu\text{m}$  region around the proton with the highest energy (shown relative to the laser axis by a white circle) at times,  $t = (100, 200, 400, 700, 800)$  fs in three repeat simulations - A (top), B (middle), and C (bottom).

ferences between them are attributed to changes in the macroparticle distribution within EPOCH, induced by varying the random seed as discussed in section 4.6.5. These subtle changes result in the laser interacting with the plasma volume slightly differently between the simulations, changing the longitudinal electric field experienced by the highest energy protons, which are at different positions in each simulation, as shown on figure 5.14. These differences are particularly accentuated between simulations A and C, resulting in a difference of 18% in the maximum proton energy at the end of the respective simulations, as shown in figure 5.13(a).

Compared to the maximum proton energy variation of up to 21%, and up to 44%, observed in simulations where the laser energy and pulse duration were varied by 8%, and 35% respectively in figure 5.9 and figure 5.10, this indicates inherent instability in the RSIT-enhanced proton acceleration regime, in simulations. This behaviour is also noted in the simulations with  $l = 25$  nm CH and  $l = 1$   $\mu\text{m}$  aluminium targets, though the maximum proton energy variation is not as severe. It is proposed that this is because RSIT is induced much earlier in  $l = 25$  nm CH targets, meaning the laser pulse propagates through relativistically underdense plasma within the target well before the main proton acceleration takes place, at which point the longitudinal electric fields have relatively stabilised,

compared to simulations where RSIT is induced as the peak laser intensity irradiates  $l=95$  nm CH targets, enhancing proton acceleration at this time. This could also explain why variation of maximum energy generally reduces in simulations where protons are accelerated from  $l=25$  nm CH targets by a laser of increasing intensity, as shown in figure 5.11, since RSIT onset time is induced earlier by lasers of increased intensity.

In the case of simulations with  $l=1$   $\mu\text{m}$  aluminium targets, the laser pulse doesn't propagate through the bulk plasma at all, and there is less variation in the maximum proton energy on repeat simulations compared to those of  $l=95$  nm CH targets, as shown in figure 5.11. There is, however, still relatively high variation in these simulations, which is similarly attributed to changes in the macroparticle distribution across repeat simulations, in this case within a pre-plasma density profile of scale length,  $L=1$   $\mu\text{m}$ . This pre-plasma was included at the  $l=1$   $\mu\text{m}$  target front surface in these simulations to enable fair comparison with the simulations of  $l=25$  nm and  $l=95$  nm CH targets, ensuring the maximum proton energy variation observed in these simulations is not due solely to the laser pulse propagating through a region of underdense plasma, but rather propagating in relativistically underdense plasma throughout the target volume.

## 5.4 Conclusion

This chapter reports on the sensitivity of proton acceleration within the laser-driven, RSIT-enhanced regime. An experiment carried out at the Vulcan-PW laser facility demonstrated that the maximum energy of protons accelerated from CH target foils of  $l=110$  nm thickness can vary by  $\sim 25\%$  from one shot to another due to small fluctuations of  $\sim 8\%$  in the laser energy and pulse duration. This result is supported by simulations which showed that the maximum proton energy varies by up to 21% for the same level of laser energy and pulse duration variation. Additional variation in the experiment is attributed to fluctuations in other key laser parameters which cannot be accounted for on every shot, and to high uncertainty in the variables that are measured. To investigate this, further

simulations were performed in which the laser pulse duration was varied by 35%, corresponding to fluctuation of  $\pm 1$  standard deviation from the mean value measured in the experiment. The maximum proton energy was shown to vary by up to 44% in these simulations.

In addition to highlighting the sensitivity of proton acceleration to laser pulse duration within this regime, maximum proton energies were demonstrated to vary by up to 18% variation across repeat simulations with identical conditions, but for the initial macroparticle distribution which was randomly varied. Overall, these results demonstrate that the maximum energy of protons accelerated in the RSIT-enhanced regime is extremely sensitive, varying significantly from one identical interaction to another, and even more significantly when the value of key laser parameters fluctuate.

The chapter also reports on experimental interactions with target foils of  $l=30$  nm thickness, that undergo RSIT well before the laser intensity peaks, rather than at approximately the same time, for  $l=110$  nm target foils. Maximum proton energies are consistently lower for these interactions compared to those with  $l=110$  nm target foils, though they are more robust to fluctuations in key laser parameters, and to fluctuations in simulation macroparticle distributions. This stabilisation of RSIT-enhanced proton energies is proposed to result from accelerating fields stabilising by the time the peak laser intensity arrives, when RSIT is induced earlier using thinner,  $l=30$  nm target foils. This requires further investigation beyond the scope of this thesis.

On that point, the results presented in this chapter open up several interesting avenues for future research. Firstly, the sensitivity of RSIT-enhanced proton acceleration to fluctuations in the laser energy and pulse duration has been carefully quantified for the first time. Now that these sensitivities are understood, future research will explore ways in which they can be mitigated, so that the RSIT-enhanced regime can be utilised not only to produce the highest maximum proton energies from a laser-driven source, but to do so with the stability required for applications. Work in this direction is already underway in the form of a simulation campaign investigating whether defocused laser beams can control

the plasma macroparticle distribution, and thus inherent fluctuation in maximum proton energy. Other interaction parameters will also be investigated towards this goal and, depending on the results, may prompt future experimental campaigns.

In conjunction with this research programme, it would be interesting to investigate optimising RSIT-enhanced proton acceleration from the thinner,  $l=30$  nm CH target foils used in this chapter. Interactions with such targets induce RSIT much earlier than is optimal for achieving the highest possible maximum proton energies, but still result in improved energies compared to protons produced in TNSA, importantly with significantly less variation compared to these, and protons accelerated from  $l=110$  nm CH target foils. Numerous experimental and simulation campaigns have optimised maximum proton energies in the latter regime by varying target foil thickness and the laser temporal-intensity profile [31, 45, 71, 79], though no optimisation has been carried out for potentially more stable RSIT-enhanced proton acceleration from  $l=30$  nm target foils. Future research will therefore investigate optimising this regime, so that high maximum proton energies can be produced stably, in a way that is robust to fluctuations in key interaction conditions.

# Chapter 6

## Conclusions and outlooks for future research

The investigations presented in this thesis contribute to our understanding of interactions between short duration, high-power laser pulses and solid target foils. More specifically, new understanding of laser-driven proton acceleration mechanisms has been described, and new methods to optimise and stabilise these mechanisms developed, accelerating progress towards laser-driven proton beams being utilised in important applications. These contributions are now summarised, and avenues for future research are proposed.

### 6.1 Investigation into multi-parameter Bayesian optimisation of laser-driven ion acceleration in particle-in-cell simulations

In chapter 4, techniques were developed and implemented to optimise the maximum energy of protons accelerated in 2D particle-in-cell (PIC) simulations of laser-solid interactions more quickly, with less data, and for more parameters than was previously possible. This was achieved using a new code called BISHOP, which automatically generates input parameter values for simulations according to a researcher's instructions, submits simulations with these inputs to a high-

performance computing (HPC) cluster, and analyses the resultant data. This facilitates the generation of large simulation data-sets, required to understand and optimise radiation production in laser-solid interactions, which depends on many input laser and target parameters. Analysing these data-sets presents another challenge, since humans have a limited capacity to understand data-sets with more than three dimensions. In contrast, machine learning (ML) techniques are well suited to analysing high dimensional data-sets. Several ML algorithms were tested to see which could most effectively optimise in a regime of limited and noisy data, as often produced in numerical and experimental laser-solid interaction studies. Bayesian optimisation was selected, and applied to optimise the maximum energy of protons accelerated via TNSA in  $\sim 5$  times fewer simulations than by using a conventional grid search technique, corresponding to a reduction of  $\sim 45,000$  hours in simulation time. This was for two interaction parameters, the laser energy and pulse duration, and proved the utility of the method. Bayesian optimisation was then applied to optimise the maximum proton energy as a function of four interaction parameters, now also including the target thickness and pre-plasma density scale length. In this case optimisation was achieved in  $\sim 200$  times fewer simulations compared to a conventional grid search, equating to a reduction of  $\sim 48$  days of simulation time. The maximum proton energy also doubled compared to only optimising for two interaction parameters, demonstrating the need to efficiently optimise many interaction parameters to increase maximum proton energies towards that required for applications. The Bayesian optimisation algorithm also identified a dependency of maximum proton energy on scale length which had not been shown previously, over the scale length range,  $L=(0 \rightarrow 4) \mu\text{m}$ . This was attributed to a subtle optimal balance between increasing laser-to-electron energy coupling whilst mitigating electron divergence, to maintain a high number of electrons along the target normal axis, optimising the sheath field strength along this axis, and thus the maximum energy of TNSA protons.

Overall, the BISHOP code, integrated Bayesian optimisation algorithm, and the new physical insights they facilitated, represent a significant contribution

to optimising maximum proton energies in laser-solid interactions, as recognised through publication in *New Journal of Physics* [100]. Beyond the scope of this thesis, the code was subsequently used to optimise multiple features of synchrotron radiation produced in simulated laser-solid interactions simultaneously, in results published in *High Power Laser Science and Engineering* [263]. This demonstrates the flexibility of the approach, which has been extended to guide an experimental investigation using the Gemini laser system in 2022. Here, laser and target conditions were guided by the Bayesian algorithm’s optimisation of fast electron temperatures in simulated, TNSA dominant interactions. This work is being prepared for submission to a peer reviewed journal. Further ML algorithms in the form of deep neural networks (DNN) have now been integrated into BISHOP, and will be directly applied in experiments investigating proton acceleration in the TNSA and RSIT-enhanced regime.

Alongside these further investigations, there is scope to extend the optimisation presented in chapter 4, by including further interaction parameters that have been demonstrated to influence the maximum energy of TNSA protons. In particular, given subtle optimisation of laser-to-electron energy coupling and transport significantly increases the maximum proton energy, a further investigation will allow the BISHOP and Bayesian optimisation algorithm to vary the laser incident angle, polarisation and wavelength. Combined with pre-plasma density scale length and laser intensity, these parameters determine the extent to which laser-to-electron energy coupling occurs through resonance absorption, vacuum heating, and  $\mathbf{j} \times \mathbf{B}$  heating. Including these parameters could therefore result in a very interesting optimisation in seven dimensions, far beyond that possible using conventional grid search techniques. Similarly, a further investigation will extend the range of parameter values available to the algorithm to explore different physical regimes compared to the optimisation of chapter 5, where the target thickness was deliberately restricted to  $l=(0.5 \rightarrow 3) \mu\text{m}$  to optimise proton energies in the well understood TNSA regime. Reducing the lower target thickness limit to  $l=10 \text{ nm}$  would allow the algorithm to explore whether higher maximum proton energies can be accelerated from thinner targets through radiation pressure

acceleration (RPA), or a hybrid acceleration mechanism enhanced by relativistic self-induced transparency (RSIT). Increased maximum proton energies beyond that achievable in TNSA are likely in the latter regime, as explored in chapter 5, and summarised in the following section.

## 6.2 Investigation into the stability of enhanced ion acceleration in the relativistic self-induced transparency regime

In chapter 5, the stability of enhanced proton acceleration from ultra-thin foils undergoing RSIT was investigated towards the goal of stably producing energetic protons sufficient for applications. This began by expanding on experimental results presented in Higginson *et al.* [71], that demonstrated maximum proton energies of near 100 MeV using this mechanism, but only for a narrow range of target foil thicknesses. Proton energy enhancement was optimised when the target foil underwent RSIT just as it was irradiated by the peak pulse intensity, making this process very sensitive to target foil thickness [79] and laser temporal-intensity contrast [45], unlike relatively robust proton acceleration in the TNSA regime [62, 63]. Further analysis of these results presented in chapter 5 demonstrate that the maximum proton energy varies significantly even between interactions where the target foil thickness was kept constant. This also appears to be the case in more recent results, that demonstrated even higher proton maximum energies of up to 150 MeV [31]. The underlying reasons for this instability must be understood and controlled if proton energies are to be stably enhanced in the RSIT regime.

To this end, results from an experiment at the Vulcan-PW laser facility are presented in chapter 5. Maximum proton energies are shown to vary with target foil thickness as expected, but also by up to 25% when the laser energy and pulse duration fluctuate by just 8% between interactions. This is supported by 2D PIC simulations that show up to 21% variation in maximum proton energy

for the same laser parameter fluctuations, with additional experimental variation attributed to fluctuation in other interaction parameters which can't be monitored for every interaction. When the laser energy and pulse duration fluctuated more significantly, by up to 35% between interactions, the maximum proton energy varied by up to 44%. Further investigation demonstrated maximum proton energy fluctuations of up to 18% on repeat simulations with identical parameters, highlighting inherent instabilities in this regime. Overall, these fluctuations, in interactions with  $l=110$  nm targets, were more significant compared to those demonstrated in a TNSA regime, using  $l=6$   $\mu\text{m}$  targets, in the experiment and in simulations, though maximum proton energies were generally higher in the RSIT-enhanced regime. This highlights RSIT-enhanced acceleration as a promising route to producing proton energies sufficient for certain applications, but one that must be better understood if it is to produce these protons with the same stability as TNSA.

Experimental and simulated interactions using thinner,  $l=30$  nm, targets showed promise in realising this goal, despite RSIT occurring much earlier than the peak pulse intensity irradiates the target for optimal proton energy enhancement. Even deoptimising this regime, proton energies were enhanced compared to those produced in the TNSA regime. Importantly, they were also less susceptible to laser parameter fluctuations than the energy of protons accelerated in the optimum RSIT-enhanced and TNSA regimes, varying by a maximum of just 15% when the laser energy and pulse duration fluctuated by 8%. Proton acceleration in this early RSIT onset regime therefore represents a potential best of both worlds, and has not previously been investigated to optimise proton energies, unlike the other regimes.

Overall, chapter 5 presents new insight into the stability of RSIT-enhanced proton acceleration, significantly contributing to a major goal of laser-solid interaction research - to stably produce energetic protons suitable for applications. As such, the findings will be prepared for submission to a peer reviewed journal. Further investigations will explore the findings in more detail to understand why the RSIT-enhanced regime is so sensitive, particularly even in simulations

when all measurable interaction parameters are kept constant. This investigation is already underway, with some indication that maximum proton energies are correlated to changes in the laser beam propagation within relativistically underdense plasma. This work will be continued, and supplemented by investigating another set of simulations where the laser focal spot has been iteratively defocused to try to control the laser beam divergence within the relativistically underdense plasma, and thus the inherent maximum proton energy fluctuations. Alongside this work, enhanced proton acceleration in interactions where RSIT occurs earlier than is optimal for the highest maximum proton energies will be investigated and optimised, so that proton energies can be consistently enhanced above those produced via TNSA, in a way that is robust to fluctuating interaction conditions.

In addition, building on the work of chapter 4, Bayesian optimisation will be applied to optimise not only proton energies, but simultaneously the stability with which they can be delivered in the RSIT regime. This will be investigated in simulations, and potentially in experiments planned at the Gemini laser facility later this year. This laser fires every 20 seconds, rather than Vulcan-PW which fires every 30 minutes, meaning more statistically significant measurements can be made and investigated. Proton acceleration in the RSIT and TNSA regime could then not only be optimised, but rigorously modelled on these experiments, using more data-intensive DNN machine learning algorithms, integrated into BISHOP following the work presented in chapter 4. This will provide further understanding of maximum proton energies which can stably be produced in each regime, as a function of many important laser and target parameters.

### **6.3 Final remarks and outlook for future research**

Having summarised the investigations presented in this thesis, and the contributions made towards realising a stable laser-driven proton source of improved energy for applications such as radiotherapy, it is fitting to broadly reflect on pathways to achieving this goal, and the goal of improving other proton source

properties. Three pathways in particular provide optimism that these goals can be achieved.

Firstly, a new generation of laser systems are currently being commissioned to deliver increased intensity at shorter pulse duration. Proton energy and other important properties are known to scale with laser intensity, and will therefore be improved in experiments at these new facilities. Pulse properties are also expected to be delivered more stably, at higher repetition rates, compared to the previous generation of lasers. Systems are therefore being established specifically with ‘user end stations’, as are commonplace in conventional accelerator facilities, so that researchers can utilise fixed laser beamlines, targetry and diagnostics, making only minor adjustments over a few days, before gathering thousands of data-points. This is in stark contrast to experiments of the past and present, such as that presented in chapter 5, where researchers typically set up beamlines, targetry and diagnostics from scratch, to collect tens or hundreds of data-points using relatively low repetition rate lasers. This transition requires development not just of laser facilities, but also diagnostics and analysis software. But once these challenges are met, large data-sets of the most intense laser-solid interactions will be available, facilitating a new frontier of statistical and ML-based analysis.

As discussed in the introductory chapter 1, ML has made significant contributions to scientific research, such as the first demonstration of ignition in inertial fusion reactions. A review paper by Hatfield *et al.* [103] describes a suite of ML techniques that contributed towards this goal. Individual applications of ML to laser-plasma interactions, such as in chapter 4, show promise, and a combined approach could similarly contribute to breakthroughs in this field. An experiment can then be imagined where laser intensity is maximised by using ML to optimise the thermal load on amplifying media, beamline optics, and compressor gratings, whilst monitoring and mitigating damage to avoid degradation of performance. Sources of unwanted pre-pulses and amplified spontaneous emission could be similarly monitored to improve laser-temporal intensity contrast. So too could the spatial-intensity profile, to optimise the size and encircled energy within a Gaussian focal spot, and perhaps to manipulate light into flat-top, and

other beam profiles on demand. Combined with ML control and optimisation of data from active diagnostics, laser-driven accelerators could be predominantly automated to produce radiation sources sufficient for applications. By changing the objective function of an algorithm being utilised to optimise experimental measurements, researchers could switch between optimising properties of protons, electrons, x-rays,  $\gamma$ -rays and THz radiation, exploiting the multi-modality of laser-solid interactions. Furthermore, ML contributed new experimental designs of benefit to the ignition result, making novel insights beyond a human guide. This leads to the third stand out pathway to improved laser-driven proton acceleration, the discovery and optimisation of new physics.

Experiments at the new generation of laser facilities delivering laser intensities beyond  $I_L=10^{23}$  Wcm<sup>-2</sup> in the immediate future will investigate previously unexplored physics in statistically significant data-sets. These intensities may be sufficient to stably accelerate laser-driven protons with improved properties required for applications using TNSA. Alternatively, this may be achieved using RSIT-enhanced acceleration, particularly given understanding of this regime is relatively recent, and has been developed from hundreds of data-points over several experiments spanning a decade or so. Such data-sets are produced in minutes on new systems, and much larger data-sets will facilitate improved understanding of the influence of many laser and target parameters in this regime, potentially meaning proton beam properties can be enhanced with sufficient stability for applications. Furthermore, radiation pressure acceleration has been proposed to produce improved proton beam properties compared to TNSA. This has been difficult to demonstrate at petawatt scale facilities achieving  $I_L=10^{21}$  Wcm<sup>-2</sup>, but is expected to be observed in interactions beyond  $I_L=10^{23}$  Wcm<sup>-2</sup>.

# Bibliography

- [1] M.N Martins and T.F Silva. Electron accelerators: History, applications, and perspectives. *Radiation Physics and Chemistry*, 95:78–85, 2014.
- [2] E.O Lawrence and M.S Livingston. The production of high speed light ions without the use of high voltages. *Physical Review*, 40(1):19, 1932.
- [3] E.O Lawrence. The evolution of the cyclotron. *Nobel Lecture*, 1951.
- [4] W.K Panofsky. Evolution of particle accelerators. *SLAC Beam Line*, 27(1):36–44, 1997.
- [5] S Hofmann, V Ninov, F.P Heßberger, P Armbruster, H Folger, G Münzenberg, et al. The new element 112. *Zeitschrift für Physik A Hadrons and Nuclei*, 354:229–230, 1996.
- [6] J.M Gates, C.E Düllmann, M Schädel, A Yakushev, A Türler, K Eberhardt, et al. First superheavy element experiments at the gsi recoil separator tasca: The production and decay of element 114 in the pu 244 (ca 48, 3-4 n) reaction. *Physical Review C*, 83(5):054618, 2011.
- [7] P.J Karol, R.C Barber, B.M Sherrill, E Vardaci, and T Yamazaki. Discovery of the elements with atomic numbers  $z= 113, 115$  and  $117$  (iupac technical report). *Pure and Applied Chemistry*, 88(1-2):139–153, 2016.
- [8] O Barbalat. Applications of particle accelerators. Technical report, CERN, 1990.
- [9] T Feder. Accelerator school travels university circuit. *Physics Today*, 63(2):20–22, 2010.

- [10] Manchester cancer research centre, proton beam therapy [online, accessed 1 march 2024]. available from: <https://www.mrc.manchester.ac.uk/research/research-themes/radiotherapy/proton-beam-therapy>.
- [11] University college london hospitals, proton beam therapy [online, accessed 1 march 2024]. available from: <https://www.uclh.nhs.uk/our-services/find-service/cancer-services/proton-beam-therapy-pbt>.
- [12] The christie nhs foundation trust, proton beam therapy [online, accessed 1 march 2024]. available from: <https://www.christie.nhs.uk/patients-and-visitors/services/proton-beam-therapy>.
- [13] T.H Maiman. Stimulated optical radiation in ruby. *Nature*, 187(4736):493–494, 1960.
- [14] W.I Linlor. Plasmas produced by laser bursts. *Bull. Am. Phys. Soc*, 7(440):1, 1962.
- [15] J.J Muray. Photoelectric effect induced by high-intensity laser light beam from quartz and borosilicate glass. Technical report, Stanford Univ., Calif. SLAC, 1962.
- [16] C.M Verber and A.H Adelman. Laser-induced thermionic emission. *Applied Physics Letters*, 2(11):220–222, 1963.
- [17] R.E Honig and J.R Woolston. Laser-induced emission of electrons, ions, and neutral atoms from solid surfaces. *Applied Physics Letters*, 2(7):138–139, 1963.
- [18] F Giori, L.A MacKenzie, and E.J McKinney. Laser-induced thermionic emission. *Applied Physics Letters*, 3(2):25–27, 1963.
- [19] D Lichtman and J.F Ready. Laser beam induced electron emission. *Physical Review Letters*, 10(8):342, 1963.

- [20] F.J McClung and R.W Hellwarth. Giant optical pulsations from ruby. *Applied Optics*, 1(101):103–105, 1962.
- [21] W.I Linlor. Ion energies produced by laser giant pulse. *Applied Physics Letters*, 3(11):210–211, 1963.
- [22] J.F Ready. Development of plume of material vaporized by giant-pulse laser. *Applied Physics Letters*, 3(1):11–13, 1963.
- [23] W.I Linlor. Some properties of plasma produced by laser giant pulse. *Physical Review Letters*, 12(14):383, 1964.
- [24] L.E Hargrove, R.L Fork, and M.A Pollack. Locking of he–ne laser modes induced by synchronous intracavity modulation. *Applied Physics Letters*, 5(1):4–5, 1964.
- [25] S.J Gitomer, R.D Jones, F Begay, A.W Ehler, J.F Kephart, and R Kristal. Fast ions and hot electrons in the laser–plasma interaction. *The Physics of Fluids*, 29(8):2679–2688, 1986.
- [26] D Strickland and G Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 55(6):447–449, 1985.
- [27] S.P Hatchett, C.G Brown, T.E Cowan, E.A Henry, J.S Johnson, M.H Key, et al. Electron, photon, and ion beams from the relativistic interaction of petawatt laser pulses with solid targets. *Physics of Plasmas*, 7(5):2076–2082, 2000.
- [28] E.L Clark, K Krushelnick, M Zepf, F.N Beg, M Tatarakis, A Machacek, et al. Energetic heavy-ion and proton generation from ultraintense laser-plasma interactions with solids. *Physical Review Letters*, 85(8):1654, 2000.
- [29] R.A Snavely, M.H Key, S.P Hatchett, T.E Cowan, M Roth, T.W Phillips, et al. Intense high-energy proton beams from petawatt-laser irradiation of solids. *Physical Review Letters*, 85(14):2945, 2000.

- [30] M Zimmer, S Scheuren, T Ebert, G Schaumann, B Schmitz, J Hornung, et al. Analysis of laser-proton acceleration experiments for development of empirical scaling laws and others. *Physical Review E*, 104(4):045210, 2021.
- [31] T Ziegler, I Göthel, S Assenbaum, C Bernert, F.E Brack, T.E Cowan, et al. Laser-driven high-energy proton beams from cascaded acceleration regimes. *Nature Physics*, 2024.
- [32] T.E Cowan, J Fuchs, H Ruhl, A Kemp, P Audebert, M Roth, et al. Ultralow emittance, multi-mev proton beams from a laser virtual-cathode plasma accelerator. *Physical Review Letters*, 92(20):204801, 2004.
- [33] S.C Wilks, A.B Langdon, T.E Cowan, M Roth, M Singh, S.P Hatchett, et al. Energetic proton generation in ultra-intense laser–solid interactions. *Physics of Plasmas*, 8(2):542–549, 2001.
- [34] M Omet, H Hayano, A Kuramoto, T Matsumoto, S Michizono, T Miura, et al. High-gradient near-quench-limit operation of superconducting tesla-type cavities in scope of the international linear collider. *Physical Review Special Topics-Accelerators and Beams*, 17(7):072003, 2014.
- [35] H.S Padamsee. Superconducting radio-frequency cavities. *Annual Review of Nuclear and Particle Science*, 64:175–196, 2014.
- [36] H Daido, M Nishiuchi, and A.S Pirozhkov. Review of laser-driven ion sources and their applications. *Reports on progress in physics*, 75(5):056401, 2012.
- [37] F Kroll, F.E Brack, C Bernert, S Bock, E Bodenstein, K Brüchner, et al. Tumour irradiation in mice with a laser-accelerated proton beam. *Nature Physics*, 18(3):316–322, 2022.
- [38] A.J Mackinnon, Y Sentoku, P.K Patel, D.W Price, S.P Hatchett, M.H Key, et al. Enhancement of proton acceleration by hot-electron recirculation in thin foils irradiated by ultraintense laser pulses. *Physical Review Letters*, 88(21):215006, 2002.

- [39] D Neely, P Foster, A.P Robinson, F Lindau, O Lundh, A Persson, et al. Enhanced proton beams from ultrathin targets driven by high contrast laser pulses. *Applied Physics Letters*, 89(2), 2006.
- [40] J Fuchs, P Antici, E d’Humières, E Lefebvre, M Borghesi, E Brambrink, et al. Laser-driven proton scaling laws and new paths towards energy increase. *Nature Physics*, 2(1):48–54, 2006.
- [41] L Robson, P.T Simpson, R.J Clarke, K.W Ledingham, F Lindau, O Lundh, et al. Scaling of proton acceleration driven by petawatt-laser–plasma interactions. *Nature Physics*, 3(1):58–62, 2007.
- [42] N.P Dover, M Nishiuchi, H Sakaki, K Kondo, M.A Alkhimova, A.Y Faenov, et al. Effect of small focus on electron heating and proton acceleration in ultrarelativistic laser-solid interactions. *Physical Review Letters*, 124(8):084802, 2020.
- [43] M Kaluza, J Schreiber, M.I Santala, G.D Tsakiris, K Eidmann, J Meyer-ter Vehn, et al. Influence of the laser prepulse on proton acceleration in thin-foil experiments. *Physical Review Letters*, 93(4):045003, 2004.
- [44] P McKenna, F Lindau, O Lundh, D Neely, A Persson, C.G Wahlström, et al. High-intensity laser-driven proton acceleration: influence of pulse contrast. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 364(1840):711–723, 2006.
- [45] N.P Dover, T Ziegler, S Assenbaum, C Bernert, S Bock, F.E Brack, et al. Enhanced ion acceleration from transparency-driven foils demonstrated at two ultraintense laser facilities. *Light: Science and Applications*, 12(1):71, 2023.
- [46] R.J Gray, D.C Carroll, X.H Yuan, C.M Brenner, M Burza, M Coury, et al. Laser pulse propagation and enhanced energy coupling to fast electrons in dense plasma gradients. *New Journal of Physics*, 16(11):113075, 2014.

- [47] P McKenna, D.C Carroll, O Lundh, F Nürnberg, K Markey, S Bandyopadhyay, et al. Effects of front surface plasma expansion on proton acceleration in ultraintense laser irradiation of foil targets. *Laser and Particle Beams*, 26(4):591–596, 2008.
- [48] J.S Green, V.M Ovchinnikov, R.G Evans, K.U Akli, H Azechi, F.N Beg, et al. Effect of laser intensity on fast-electron-beam divergence in solid-density plasmas. *Physical Review Letters*, 100(1):015003, 2008.
- [49] J.S Green, C.D Murphy, N Booth, R.J Dance, R.J Gray, D.A MacLellan, et al. Single shot, temporally and spatially resolved measurements of fast electron dynamics using a chirped optical probe. *Journal of Instrumentation*, 9(03):03003, 2014.
- [50] R.J Gray, R Wilson, M King, S.D Williamson, R.J Dance, C Armstrong, et al. Enhanced laser-energy coupling to dense plasmas driven by recirculating electron currents. *New Journal of Physics*, 20(3):033021, 2018.
- [51] P Mora. Plasma expansion into a vacuum. *Physical Review Letters*, 90(18):185002, 2003.
- [52] J.M Beckers. Adaptive optics for astronomy: principles, performance, and applications. *Annual Review of Astronomy and Astrophysics*, 31(1):13–62, 1993.
- [53] B Dromey, S Kar, M Zepf, and P Foster. The plasma mirror—a subpicosecond optical switch for ultrahigh power lasers. *Review of Scientific Instruments*, 75(3):645–649, 2004.
- [54] R Wilson. On the role of focal spot size in ultra-intense laser-solid interaction physics. PhD thesis, University of Strathclyde, 2018.
- [55] T.P Frazer. Investigation of laser-solid interaction physics with tightly focused, ultra-intense laser pulses. PhD thesis, University of Strathclyde, 2021.

- [56] W Shaikh, I Musgrave, C Hernandez-Gomez, B Parry, and D Johnson. A high contrast dual opcpa pre-amplifier system using both picosecond and nanosecond pump pulses for the vulcan petawatt facility. In *Advanced Solid-State Photonics*, page ATuA18. Optica Publishing Group, 2010.
- [57] C Hooker, Y Tang, O Chekhlov, J.L Collier, E Divall, K Ertel, et al. Improving coherent contrast of petawatt laser pulses. *Optics Express*, 19(3):2193–2203, 2011.
- [58] I Musgrave, A Boyle, D.C Carroll, R.J Clarke, R Heathcote, M Galimberti, et al. Recent developments on the vulcan high power laser facility. *High-Power, High-Energy, and High-Intensity Laser Technology; and Research Using Extreme Light: Entering New Frontiers with Petawatt-Class Lasers*, 8780:17–30, 2013.
- [59] K Ogura, M Nishiuchi, A.S Pirozhkov, T Tanimoto, A Sagisaka, T.Z Esirkepov, et al. Proton acceleration to 40 mev using a high intensity, high contrast optical parametric chirped-pulse amplification/ti: sapphire hybrid laser system. *Optics Letters*, 37(14):2868–2870, 2012.
- [60] H Kiriya, Y Miyasaka, A Kon, M Nishiuchi, A Sagisaka, H Sasao, et al. Enhancement of pre-pulse and picosecond pedestal contrast of the petawatt j-karen-p laser. *High Power Laser Science and Engineering*, 9:e62, 2021.
- [61] X Lu, H Zhang, J Li, and Y Leng. Reducing temporal pedestal in a ti: sapphire chirped-pulse amplification system by using a stretcher based on two concave mirrors. *Optics Letters*, 46(21):5320–5323, 2021.
- [62] F Wagner, O Deppert, C Brabetz, P Fiala, A Kleinschmidt, P Poth, et al. Maximum proton energy above 85 mev from the relativistic interaction of laser pulses with micrometer thick ch<sub>2</sub> targets. *Physical Review Letters*, 116(20):205002, 2016.
- [63] J Hornung, Y Zobus, P Boller, C Brabetz, U Eisenbarth, T Kühl, et al.

Enhancement of the laser-driven proton source at phelix. *High Power Laser Science and Engineering*, 8:e24, 2020.

- [64] B.M Hegelich, I Pomerantz, L Yin, H.C Wu, D Jung, B.J Albright, et al. Laser-driven ion acceleration from relativistically transparent nanotargets. *New Journal of Physics*, 15(8):085015, 2013.
- [65] H.W Powell, M King, R.J Gray, D.A MacLellan, B Gonzalez-Izquierdo, L.C Stockhausen, et al. Proton acceleration enhanced by a plasma jet in expanding foils undergoing relativistic transparency. *New Journal of Physics*, 17(10):103033, 2015.
- [66] M King, R.J Gray, H.W Powell, D.A MacLellan, B Gonzalez-Izquierdo, L.C Stockhausen, et al. Ion acceleration and plasma jet formation in ultra-thin foils undergoing expansion and relativistic transparency. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 829:163–166, 2016.
- [67] S Palaniyappan, C Huang, D.C Gautier, C.E Hamilton, M.A Santiago, C Kreuzer, et al. Efficient quasi-monoenergetic ion beams from laser-driven relativistic plasmas. *Nature Communications*, 6(1):10170, 2015.
- [68] M.J Duff, R Wilson, M King, B Gonzalez-Izquierdo, A Higginson, S.D Williamson, et al. High order mode structure of intense light fields generated via a laser-driven relativistic plasma aperture. *Scientific Reports*, 10(1):105, 2020.
- [69] B Gonzalez-Izquierdo, R.J Gray, M King, R.J Dance, R Wilson, J McCreddie, et al. Optically controlled dense current structures driven by relativistic plasma aperture-induced diffraction. *Nature Physics*, 12(5):505–512, 2016.
- [70] B Gonzalez-Izquierdo, M King, R.J Gray, R Wilson, R.J Dance, H Powell, et al. Towards optical polarization control of laser-driven proton accelera-

- tion in foils undergoing relativistic transparency. *Nature Communications*, 7(1):12891, 2016.
- [71] A Higginson, R.J Gray, M King, R.J Dance, S.D Williamson, N.M Butler, et al. Near-100 mev protons via a laser-driven transparency-enhanced hybrid acceleration scheme. *Nature Communications*, 9(1):724, 2018.
- [72] S.C Wilks, W.L Kruer, M Tabak, and A.B Langdon. Absorption of ultra-intense laser pulses. *Physical Review Letters*, 69(9):1383, 1992.
- [73] T.Z Esirkepov, M Borghesi, S.V Bulanov, G Mourou, and T Tajima. Highly efficient relativistic-ion generation in the laser-piston regime. *Physical Review Letters*, 92(17):175003, 2004.
- [74] A.P Robinson, P Gibbon, M Zepf, S Kar, R.G Evans, and C Bellei. Relativistically correct hole-boring and ion acceleration by circularly polarized laser pulses. *Plasma Physics and Controlled Fusion*, 51(2):024004, 2009.
- [75] L Yin, B.J Albright, B.M Hegelich, and J.C Fernández. GeV laser ion acceleration from ultrathin targets: The laser break-out afterburner. *Laser and Particle Beams*, 24(2):291–298, 2006.
- [76] L Yin, B.J Albright, B.M Hegelich, K.J Bowers, K.A Flippo, T.J Kwan, et al. Monoenergetic and GeV ion acceleration from the laser breakout afterburner using ultrathin targets. *Physics of Plasmas*, 14(5), 2007.
- [77] B.J Albright, L Yin, K.J Bowers, B.M Hegelich, K.A Flippo, T.J Kwan, et al. Relativistic buneman instability in the laser breakout afterburner. *Physics of Plasmas*, 14(9), 2007.
- [78] M King, R Wilson, E.F Bacon, E.J Dolier, T.P Frazer, J Goodman, et al. Perspectives on laser-plasma physics in the relativistic transparency regime. *The European Physical Journal A*, 59(6):132, 2023.
- [79] J Goodman, M King, R Wilson, R.J Gray, and P McKenna. Optimisation of multi-petawatt laser-driven proton acceleration in the relativistic transparency regime. *New Journal of Physics*, 24(5):053016, 2022.

- [80] J.P Zou, C Le Blanc, D.N Papadopoulos, G Chériaux, P Georges, G Menerat, et al. Design and current progress of the apollon 10 pw project. *High Power Laser Science and Engineering*, 3:e2, 2015.
- [81] S Kühn, M Dumergue, S Kahaly, S Mondal, M Füle, T Csizmadia, et al. The eli-alps facility: the next generation of attosecond sources. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 50(13):132002, 2017.
- [82] C.N Danson, C Haefner, J Bromage, T Butcher, J.C Chanteloup, E.A Chowdhury, et al. Petawatt and exawatt class lasers worldwide. *High Power Laser Science and Engineering*, 7:e54, 2019.
- [83] C.N Danson, M White, J.R Barr, T Bett, P Blyth, D Bowley, et al. A history of high-power laser research and development in the united kingdom. *High Power Laser Science and Engineering*, 9:e18, 2021.
- [84] P Mason, N Stuart, J Phillips, R Heathcote, S Buck, A Wojtusiak, et al. Progress on laser development at the extreme photonics applications centre. In *The European Conference on Lasers and Electro-Optics*, page ca.8.2. Optica Publishing Group, 2023.
- [85] P McKenna, K.W Ledingham, I Spencer, T McCany, R.P Singhal, C Ziener, et al. Characterization of multiterawatt laser-solid interactions for proton acceleration. *Review of Scientific Instruments*, 73(12):4176–4184, 2002.
- [86] S Ter-Avetisyan, M Schnürer, P.V Nickles, M Kalashnikov, E Risse, T Sokollik, et al. Quasimonoenergetic deuteron bursts produced by ultraintense laser pulses. *Physical Review Letters*, 96(14):145006, 2006.
- [87] J Hah, G.M Petrov, J.A Nees, Z.H He, M.D Hammig, K Krushelnick, et al. High repetition-rate neutron generation by several-mj, 35 fs pulses interacting with free-flowing d2o. *Applied Physics Letters*, 109(14), 2016.
- [88] M Noaman-ul Haq, H Ahmed, T Sokollik, L Yu, Z Liu, X Yuan, et al. Statistical analysis of laser driven protons using a high-repetition-rate tape

- drive target system. *Physical Review Accelerators and Beams*, 20(4):041301, 2017.
- [89] C Spindloe, G Arthur, F Hall, S Tomlinson, R Potter, S Kar, et al. High volume fabrication of laser targets using mems techniques. In *Journal of Physics: Conference Series*, volume 713, page 012002. IOP Publishing, 2016.
- [90] C.A Palmer. Paving the way for a revolution in high repetition rate laser-driven ion acceleration. *New Journal of Physics*, 20(6):061001, 2018.
- [91] N Xu, M.J Streeter, O.C Ettliger, H Ahmed, S Astbury, M Borghesi, et al. Versatile tape-drive target for high-repetition-rate laser-driven proton acceleration. *High Power Laser Science and Engineering*, 11:e23, 2023.
- [92] A Leblanc, S Monchocé, C Bourassin-Bouchet, S Kahaly, and F Quéré. Ptychographic measurements of ultrahigh-intensity laser-plasma interactions. *Nature Physics*, 12(4):301–305, 2016.
- [93] N.P Dover, M Nishiuchi, H Sakaki, M.A Alkhimova, A.Y Faenov, Y Fukuda, et al. Scintillator-based transverse proton beam profiler for laser-plasma ion sources. *Review of Scientific Instruments*, 88(7), 2017.
- [94] D.A Mariscal, B.Z Djordjević, E.S Grace, R Hollinger, T Ma, G.G Scott, et al. Design of flexible proton beam imaging energy spectrometers (probies). *Plasma Physics and Controlled Fusion*, 63(11):114003, 2021.
- [95] D.A Mariscal, B.Z Djordjević, R Anirudh, T Bremer, P.C Campbell, S Feister, et al. A flexible proton beam imaging energy spectrometer (probies) for high repetition rate or single-shot high energy density (hed) experiments. *Review of Scientific Instruments*, 94(2), 2023.
- [96] V Istokskaia, B Lefebvre, G Petringa, P Cirrone, M Guarrera, L Giuffrida, et al. Proton bragg curve and energy reconstruction using an online scintillator stack detector. *Review of Scientific Instruments*, 94(7), 2023.

- [97] C.D Armstrong, G.G Scott, S Richards, J.K Patel, K Fedorov, R.J Gray, et al. X-ray detector requirements for laser-plasma accelerators. *Frontiers in Physics*, 11:1286442, 2023.
- [98] K.M George, J.T Morrison, S Feister, G.K Ngirmang, J.R Smith, et al. High-repetition-rate (khz) targets and optics from liquid microjets for high-intensity laser-plasma interactions. *High Power Laser Science and Engineering*, 7:e50, 2019.
- [99] C.I Underwood, G Gan, Z.H He, C.D Murphy, A.G Thomas, K Krushelnick, et al. Characterization of flowing liquid films as a regenerating plasma mirror for high repetition-rate laser contrast enhancement. *Laser and Particle Beams*, 38(2):128–134, 2020.
- [100] E.J Dolier, M King, R Wilson, R.J Gray, and P McKenna. Multi-parameter bayesian optimisation of laser-driven ion acceleration in particle-in-cell simulations. *New Journal of Physics*, 24(7):073025, 2022.
- [101] R.J Gray. Lpi-py [online, accessed 1 march 2024]. available from: <https://gitlab.cis.strath.ac.uk/lpi-py>, 2022.
- [102] S Feister, K Cassou, S Dann, A Döpp, P Gauron, A.J Gonsalves, et al. Control systems and data management for high-power laser facilities. *High Power Laser Science and Engineering*, pages 1–31, 2023.
- [103] P.W Hatfield, J.A Gaffney, G.J Anderson, S Ali, L Antonelli, S Başığmez du Pree, et al. The data-driven future of high-energy-density physics. *Nature*, 593(7859):351–361, 2021.
- [104] T Ma, D Mariscal, R Anirudh, T Bremer, B.Z Djordjevic, T Galvin, et al. Accelerating the rate of discovery: Toward high-repetition-rate hed science. *Plasma Physics and Controlled Fusion*, 63(10):104003, 2021.
- [105] A.M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950.

- [106] S.J Russell and P Norvig. *Artificial Intelligence a Modern Approach*. London, 2010.
- [107] M Campbell, A.J Hoane Jr, and F.H Hsu. Deep blue. *Artificial Intelligence*, 134(1-2):57–83, 2002.
- [108] G Kasparov. Don't fear intelligent machines. work with them. *TED talk*: <https://www.youtube.com/watch?v=NP8xt8o45Q>., 2017.
- [109] G Kohns, D Hassabis, et al. Alphago - the movie. *Documentary*: <https://www.youtube.com/watch?v=WXuK6gekU1Yt=9s>, 2017.
- [110] R Toewes. Alphafold is the most important achievement in ai—ever. *Forbes*, 2021.
- [111] J Jumper, R Evans, A Pritzel, T Green, M Figurnov, O Ronneberger, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
- [112] F Ren, X Ding, M Zheng, M Korzinkin, X Cai, W Zhu, et al. Alphafold accelerates artificial intelligence powered drug discovery: efficient discovery of a novel cdk20 small molecule inhibitor. *Chemical Science*, 14(6):1443–1452, 2023.
- [113] K.T Ko, F Lennartz, D Mekhaie, B Guloglu, A Marini, D.J Deuker, et al. Structure of the malaria vaccine candidate pfs48/45 and its recognition by transmission blocking antibodies. *Nature Communications*, 13(1):5603, 2022.
- [114] M.K Higgins. Can we alphafold our way out of the next pandemic? *Journal of Molecular Biology*, 433(20):167093, 2021.
- [115] W.M Kincannon, M Zahn, R Clare, J Lusty Beech, A Romberg, J Larson, et al. Biochemical and structural characterization of an aromatic ring–hydroxylating dioxygenase for terephthalic acid catabolism. *Proceedings of the National Academy of Sciences*, 119(13):e2121426119, 2022.

- [116] J Degraeve, F Felici, J Buchli, M Neunert, B Tracey, F Carpanese, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.
- [117] A.B Zylstra, O.A Hurricane, D.A Callahan, A.L Kritcher, J.E Ralph, H.F Robey, et al. Burning plasma achieved in inertial fusion. *Nature*, 601(7894):542–548, 2022.
- [118] H Abu-Shawareb, R Acree, P Adams, J Adams, B Addis, R Aden, et al. Lawson criterion for ignition exceeded in an inertial fusion experiment. *Physical Review Letters*, 129(7):075001, 2022.
- [119] H Abu-Shawareb, R Acree, P Adams, J Adams, B Addis, R Aden, et al. Achievement of target gain larger than unity in an inertial fusion experiment. *Physical Review Letters*, 132(6):065102, 2024.
- [120] J.L Peterson, K.D Humbird, J.E Field, S.T Brandon, S.H Langer, R.C Nora, et al. Zonal flow generation in inertial confinement fusion implosions. *Physics of Plasmas*, 24(3), 2017.
- [121] V Gopaldaswamy, R Betti, J.P Knauer, N Luciani, D Patel, K.M Woo, et al. Tripled yield in direct-drive laser fusion through statistical modelling. *Nature*, 565(7741):581–586, 2019.
- [122] P.W Hatfield, S.J. Rose, and R.H Scott. The blind implosion-maker: Automated inertial confinement fusion experiment design. *Physics of Plasmas*, 26(6), 2019.
- [123] A Hsu, B Cheng, and P.A Bradley. Analysis of nif scaling using physics informed machine learning. *Physics of Plasmas*, 27(1), 2020.
- [124] D Osthus, S.A Vander Wiel, N.M. Hoffman, and F.J Wysocki. Prediction uncertainties beyond the range of experience: A case study in inertial confinement fusion implosion experiments. *SIAM/ASA Journal on Uncertainty Quantification*, 7(2):604–633, 2019.

- [125] C Amorin, L.M. Kegelmeyer, and W.P. Kegelmeyer. A hybrid deep learning architecture for classification of microscopic damage on national ignition facility laser optics. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 12(6):505–513, 2019.
- [126] R Anirudh, J.J. Thiagarajan, P.T. Bremer, and B.K. Spears. Improved surrogates in inertial confinement fusion with manifold and cycle consistencies. *Proceedings of the National Academy of Sciences*, 117(18):9741–9746, 2020.
- [127] M.E. Glinsky, T.W. Moore, W.E. Lewis, M.R. Weis, C.A. Jennings, D.J. Ampleford, et al. Quantification of maglif morphology using the mallat scattering transformation. *Physics of Plasmas*, 27(11), 2020.
- [128] S. Palaniyappan, J.P. Sauppe, B.J. Tobias, C.F. Kawaguchi, K.A. Flippo, A.B. Zylstra, et al. Hydro-scaling of direct-drive cylindrical implosions at the omega and the national ignition facility. *Physics of Plasmas*, 27(4), 2020.
- [129] A. Döpp, C. Eberle, S. Howard, F. Irshad, J. Lin, and M.J. Streeter. Data-driven science and machine learning methods in laser–plasma physics. *High Power Laser Science and Engineering*, 11:e55, 2023.
- [130] T. Pascu. A comparison between the performances of machine learning algorithms in the case of laser profile images classification. In *LPA Online Workshop on Control Systems and Machine Learning*, volume 1, 2022.
- [131] J. Lin, F. Haberstroh, S. Karsch, and A. Döpp. Applications of object detection networks in high-power laser systems and experiments. *High Power Laser Science and Engineering*, 11:e7, 2023.
- [132] Z.H. He, B. Hou, V. Lebailly, J.A. Nees, K. Krushelnick, and A.G. Thomas. Coherent control of plasma dynamics. *Nature Communications*, 6(1):7156, 2015.
- [133] J.R. Smith, C. Orban, J.T. Morrison, K.M. George, G.K. Ngirmang, E.A. Chowdhury, et al. Optimizing laser–plasma interactions for ion acceleration

- using particle-in-cell simulations and evolutionary algorithms. *New Journal of Physics*, 22(10):103067, 2020.
- [134] R.J Shalloo, S.J Dann, J.N Gruse, C.I Underwood, A.F Antoine, C Arran, et al. Automation and control of laser wakefield accelerators using bayesian optimization. *Nature Communications*, 11(1):6355, 2020.
- [135] B Loughran, M.J Streeter, H Ahmed, S Astbury, M Balcazar, M Borghesi, et al. Automated control and optimization of laser-driven ion acceleration. *High Power Laser Science and Engineering*, 11:e35, 2023.
- [136] M Borghesi, J Fuchs, S.V Bulanov, A.J Mackinnon, P.K. Patel, and M Roth. Fast ion generation by high-intensity laser irradiation of solid targets and applications. *Fusion Science and Technology*, 49(3):412–439, 2006.
- [137] V Malka, J Faure, Y.A Gauduel, E Lefebvre, A. Rousse, and K.T Phuoc. Principles and applications of compact laser–plasma accelerators. *Nature Physics*, 4(6):447–453, 2008.
- [138] R Capdessus, M King, D Del Sorbo, M Duff, C.P Ridgers, and P McKenna. Relativistic doppler-boosted  $\gamma$ -rays in high fields. *Scientific Reports*, 8(1):9155, 2018.
- [139] H.X Chang, B Qiao, T.W Huang, Z Xu, C.T Zhou, and Y.Q Gu. Brilliant petawatt gamma-ray pulse generation in quantum electrodynamic laser-plasma interaction. *Scientific Reports*, 7(1):45031, 2017.
- [140] M.J Duff. Radiation generation and high-field physics phenomena in ultra-intense laser-solid interactions. PhD thesis, University of Strathclyde, 2020.
- [141] J Nuckolls, L Wood, A Thiessen, and G Zimmerman. Laser compression of matter to super-high densities: Thermonuclear (ctr) applications. *Nature*, 239(5368):139–142, 1972.

- [142] M Tabak, J Hammer, M.E Glinsky, W.L Kruer, S.C Wilks, J Woodworth, et al. Ignition and high gain with ultrapowerful lasers. *Physics of Plasmas*, 1(5):1626–1634, 1994.
- [143] T Ruth. Accelerating production of medical isotopes. *Nature*, 457(7229):536–537, 2009.
- [144] S Fritzler, V Malka, G Grillon, J.P Rousseau, F Burgy, E Lefebvre, et al. Proton beams generated with high-intensity lasers: Applications to medical isotope production. *Applied Physics Letters*, 83(15):3039–3041, 2003.
- [145] A.J Mackinnon, P.K Patel, R.P Town, M.J Edwards, T Phillips, S.C Lerner, et al. Proton radiography as an electromagnetic field and density perturbation diagnostic. *Review of Scientific Instruments*, 75(10):3531–3536, 2004.
- [146] M Barberio, S Veltri, M. Scisciò, and P Antici. Laser-accelerated proton beams as diagnostics for cultural heritage. *Scientific Reports*, 7(1):40415, 2017.
- [147] M Barberio, M Scisciò, S Vallières, F Cardelli, S.N Chen, G Famulari, et al. Laser-accelerated particle beams for stress testing of materials. *Nature Communications*, 9(1):372, 2018.
- [148] R.R Wilson. Radiological use of fast protons. *Radiology*, 47(5):487–491, 1946.
- [149] C.A Tobias, J.H Lawrence, J.L Born, R.K McCombs, J.E Roberts, H.O Anger, et al. Pituitary irradiation with high-energy proton beams a preliminary report. *Cancer Research*, 18(2):121–134, 1958.
- [150] V Favaudon, L Caplier, V Monceau, F Pouzoulet, M Sayarath, C Fouillade, et al. Ultrahigh dose-rate flash irradiation increases the differential response between normal and tumor tissue in mice. *Science Translational Medicine*, 6(245):24593, 2014.

- [151] J Bourhis, P Montay-Gruel, P.G Jorge, C Bailat, B Petit, J Ollivier, et al. Clinical translation of flash radiotherapy: Why and how? *Radiotherapy and Oncology*, 139:11–17, 2019.
- [152] M.C Vozenin, J.H Hendry, and C.L Limoli. Biological benefits of ultra-high dose rate flash radiotherapy: sleeping beauty awoken. *Clinical Oncology*, 31(7):407–415, 2019.
- [153] J.D Wilson, E.M Hammond, G.S Higgins, and K Petersson. Ultra-high dose rate (flash) radiotherapy: silver bullet or fool’s gold? *Frontiers in Oncology*, 9:1563, 2020.
- [154] P Montay-Gruel, M.M Acharya, P Gonçalves Jorge, B Petit, I.G Petridis, P Fuchs, et al. Hypofractionated flash-rt as an effective treatment against glioblastoma that reduces neurocognitive side effects in mice. *Clinical Cancer Research*, 27(3):775–784, 2021.
- [155] S.V Bulanov, T.Z Esirkepov, V.S. Khoroshkov, A.V Kuznetsov, and F Pegoraro. Oncological hadrontherapy with laser ion accelerators. *Physics Letters A*, 299(2-3):240–247, 2002.
- [156] G Aymar, T Becker, S Boogert, M Borghesi, R Bingham, C Brenner, et al. Lhara: the laser-hybrid accelerator for radiobiological applications. *Frontiers in Physics*, 8:567738, 2020.
- [157] J Alda. Laser and gaussian beam propagation and transformation. *Encyclopedia of Optical Engineering*, 999:1013–1013, 2003.
- [158] G.S Voronov and N.B Delone. Ionization of the xenon atom by the electric field of ruby laser emission. *Soviet Physics JETP*, 1:66, 1965.
- [159] P Agostini, G Barjot, J Bonnal, G Mainfray, C Manus, and J Morellec. Multiphoton ionization of hydrogen and rare gases. *IEEE Journal of Quantum Electronics*, 4(10):667–669, 1968.

- [160] P Agostini, F Fabre, G Mainfray, G Petite, and N.K Rahman. Free-free transitions following six-photon ionization of xenon atoms. *Physical Review Letters*, 42(17):1127, 1979.
- [161] Y Gontier, M Poirier, and M Trahin. Multiphoton absorptions above the ionisation threshold. *Journal of Physics B: Atomic and Molecular Physics*, 13(7):1381, 1980.
- [162] L.V Keldysh. Ionization in the field of a strong electromagnetic wave. *Soviet Physics JETP*, 20(5):1307–1314, 1965.
- [163] V.S Popov. Tunnel and multiphoton ionization of atoms and ions in a strong laser field (keldysh theory). *Physics-Uspekhi*, 47(9):855, 2004.
- [164] A.M Perelomov. Ionization of atoms in an alternating electric field. *Soviet Physics JETP*, 23(5):924–934, 1966.
- [165] S Augst, D Strickland, D.D Meyerhofer, S.L Chin, and J.H Eberly. Tunneling ionization of noble gases in a high-intensity laser field. *Physical Review Letters*, 63(20):2212, 1989.
- [166] T Auguste, P Monot, L.A Lompré, G Mainfray, and C Manus. Multiply charged ions produced in noble gases by a 1 ps laser pulse at  $\lambda = 1053$  nm. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 25(20):4181, 1992.
- [167] C. J Evans and Y Gamal. Laser-induced breakdown of helium. *Journal of Physics D: Applied Physics*, 13(8):1447, 1980.
- [168] P Gibbon. *Short pulse laser interactions with matter: an introduction*. World Scientific, 2005.
- [169] W Kruer. *The physics of laser plasma interactions*. crc Press, 2019.
- [170] L Yang, L Huang, S Assenbaum, T.E Cowan, I Goethel, S Göde, et al. Time-resolved optical shadowgraphy of solid hydrogen jets as a testbed to benchmark particle-in-cell simulations. *Preprint*, 2023.

- [171] T.D Arber, K Bennett, C.S Brady, A Lawrence-Douglas, M.G Ramsay, N.J Sircombe, et al. Contemporary particle-in-cell approach to laser-plasma modelling. *Plasma Physics and Controlled Fusion*, 57(11):113001, 2015.
- [172] L Schlessinger and J Wright. Inverse-bremsstrahlung absorption rate in an intense laser field. *Physical Review A*, 20(5):1934, 1979.
- [173] A.B Langdon. Nonlinear inverse bremsstrahlung and heated-electron distributions. *Physical Review Letters*, 44(9):575, 1980.
- [174] H.W Koch and J.W Motz. Bremsstrahlung cross-section formulas and related data. *Reviews of Modern Physics*, 31(4):920, 1959.
- [175] M.D Perry, J.A Sefcik, T.E Cowan, S.P Hatchett, A Hunt, M Moran, et al. Hard x-ray production from high intensity laser solid interactions. *Review of Scientific Instruments*, 70(1):265–269, 1999.
- [176] K.G Estabrook, E.J Valeo, and W.L Kruer. Two-dimensional relativistic simulations of resonance absorption. *The Physics of Fluids*, 18(9):1151–1159, 1975.
- [177] D.D Meyerhofer, H Chen, J.A Delettrez, B Soom, S Uchida, and B Yaakobi. Resonance absorption in high-intensity contrast, picosecond laser–plasma interactions. *Physics of Fluids B: Plasma Physics*, 5(7):2584–2588, 1993.
- [178] J.P Freidberg, R.W Mitchell, R.L Morse, and L Rudsinski. Resonant absorption of laser light by plasma targets. *Physical Review Letters*, 28(13):795, 1972.
- [179] N.G Denisov. On a singularity of the field of an electromagnetic wave propagated in an inhomogeneous plasma. *Soviet Physics JETP*, 4(4):544–553, 1957.
- [180] F Brunel. Not-so-resonant, resonant absorption. *Physical Review Letters*, 59(1):52, 1987.

- [181] P Gibbon and A.R Bell. Collisionless absorption in sharp-edged plasmas. *Physical Review Letters*, 68(10):1535, 1992.
- [182] W.L Kruer and K Estabrook. J $\times$  b heating by very intense laser light. *The Physics of Fluids*, 28(1):430–432, 1985.
- [183] Y Ping, R Shepherd, B.F Lasinski, M Tabak, H Chen, H.K Chung, et al. Absorption of short laser pulses on solid targets in the ultrarelativistic regime. *Physical Review Letters*, 100(8):085004, 2008.
- [184] X.L Ge, X.X Lin, X.H Yuan, D.C Carroll, R.J Gray, T.P Yu, et al. Directed fast electron beams in ultraintense picosecond laser irradiated solid targets. *Applied Physics Letters*, 107(9), 2015.
- [185] B Bezzerides, S.J Gitomer, and D.W Forslund. Randomness, maxwellian distributions, and resonance absorption. *Physical Review Letters*, 44(10):651, 1980.
- [186] A Bret, L Gremillet, and M.E Dieckmann. Multidimensional electron beam-plasma instabilities in the relativistic regime. *Physics of Plasmas*, 17(12), 2010.
- [187] R.J Gray. On mechanisms of laser-coupling to fast electrons in ultraintense laser-solid interactions. PhD thesis, University of Strathclyde, 2013.
- [188] F.N Beg, A.R Bell, A.E Dangor, C.N Danson, A.P Fews, M.E Glinsky, et al. A study of picosecond laser–solid interactions up to 10<sup>19</sup> w cm<sup>-2</sup>. *Physics of Plasmas*, 4(2):447–457, 1997.
- [189] M.G Haines, M.S Wei, F.N Beg, and R.B Stephens. Hot-electron temperature and laser-light absorption in fast ignition. *Physical Review Letters*, 102(4):045008, 2009.
- [190] C.D Chen, P.K Patel, D.S Hey, A.J Mackinnon, M.H Key, K.U Akli, et al. Bremsstrahlung and  $k\alpha$  fluorescence measurements for inferring conversion efficiencies into fast ignition relevant hot electrons. *Physics of Plasmas*, 16(8), 2009.

- [191] J.R Davies. Laser absorption by overdense plasmas in the relativistic regime. *Plasma Physics and Controlled Fusion*, 51(1):014006, 2008.
- [192] M.C Levy, S.C Wilks, M Tabak, S.B Libby, and M.G Baring. Petawatt laser absorption bounded. *Nature Communications*, 5(1):4149, 2014.
- [193] R Wilson et al. Influence of spatial-intensity contrast in ultraintense laser–plasma interactions. *Scientific Reports*, 12(1):1910, 2022.
- [194] M.N Quinn, X.H Yuan, X.X Lin, D.C Carroll, O Tresca, R.J Gray, et al. Refluxing of fast electrons in solid targets irradiated by intense, picosecond laser pulses. *Plasma Physics and Controlled Fusion*, 53(2):025007, 2011.
- [195] A.R Bell and R.J Kingham. Resistive collimation of electron beams in laser-produced plasmas. *Physical Review Letters*, 91(3):035003, 2003.
- [196] W.H Bennett. Magnetically self-focussing streams. *Physical Review*, 45(12):890, 1934.
- [197] H Alfvén. On the motion of cosmic rays in interstellar space. *Physical Review*, 55(5):425, 1939.
- [198] E.S Weibel. Spontaneously growing transverse waves in a plasma due to an anisotropic velocity distribution. *Physical Review Letters*, 2(3):83, 1959.
- [199] L.O Silva, R.A Fonseca, J.W Tonge, W.B Mori, and J.M Dawson. On the role of the purely transverse weibel instability in fast ignitor scenarios. *Physics of Plasmas*, 9(6):2458–2461, 2002.
- [200] R Jung, J Osterholz, K Löwenbrück, S Kiselev, G Pretzler, A Pukhov, et al. Study of electron-beam propagation through preionized dense foam plasmas. *Physical Review Letters*, 94(19):195001, 2005.
- [201] R.G Evans. Modelling electron transport for fast ignition. *Plasma Physics and Controlled Fusion*, 49(12B):B87, 2007.
- [202] J.R Davies, J.S Green, and P.A Norreys. Electron beam hollowing in laser–solid interactions. *Plasma Physics and Controlled Fusion*, 48(8):1181, 2006.

- [203] P.A Norreys, J.S Green, J.R Davies, M Tatarakis, E.L Clark, F.N Beg, et al. Observation of annular electron beam transport in multi-terawatt laser-solid interactions. *Plasma Physics and Controlled Fusion*, 48(2):L11, 2006.
- [204] M Borghesi, A.J Mackinnon, D.H Campbell, D.G Hicks, S Kar, P.K Patel, et al. Multi-mev proton source investigations in ultraintense laser-foil interactions. *Physical Review Letters*, 92(5):055003, 2004.
- [205] K Markey, P McKenna, C.M Brenner, D.C Carroll, M.M Günther, K Harres, et al. Spectral enhancement in the double pulse regime of laser proton acceleration. *Physical Review Letters*, 105(19):195008, 2010.
- [206] A Macchi, M Borghesi, and M Passoni. Ion acceleration by superintense laser-plasma interaction. *Reviews of Modern Physics*, 85(2):751, 2013.
- [207] A Yogo, K Mima, N Iwata, S Tosaki, A Morace, Y Arikawa, et al. Boosting laser-ion acceleration with multi-picosecond pulses. *Scientific Reports*, 7(1):42451, 2017.
- [208] M Hegelich, S Karsch, G Pretzler, D Habs, K Witte, W Guenther, et al. Mev ion jets from short-pulse-laser interaction with thin foils. *Physical Review Letters*, 89(8):085002, 2002.
- [209] P Mora. Thin-foil expansion into a vacuum. *Physical Review E*, 72(5):056401, 2005.
- [210] Y Oishi, T Nayuki, T Fujii, Y Takizawa, X Wang, T Yamazaki, et al. Dependence on laser intensity and pulse duration in proton acceleration by irradiation of ultrashort laser pulses on a cu foil target. *Physics of Plasmas*, 12(7), 2005.
- [211] J Schreiber, F Bell, F Grüner, U Schramm, M Geissler, M Schnürer, et al. Analytical model for ion acceleration by high-intensity laser pulses. *Physical Review Letters*, 97(4):045005, 2006.

- [212] C.M Brenner, P McKenna, and D Neely. Modelling the effect of laser focal spot size on sheath-accelerated protons in intense laser-foil interactions. *Plasma Physics and Controlled Fusion*, 56(8):084003, 2014.
- [213] R.A Simpson, G.G Scott, D Mariscal, D Rusby, P.M King, E Grace, et al. Scaling of laser-driven electron and proton acceleration as a function of laser pulse duration, energy, and intensity in the multi-picosecond regime. *Physics of Plasmas*, 28(1), 2021.
- [214] Y Takagi, N Iwata, E d’Humieres, and Y Sentoku. Multivariate scaling of maximum proton energy in intense laser driven ion acceleration. *Physical Review Research*, 3(4):043140, 2021.
- [215] B.Z Djordjević, A.J Kemp, J Kim, R.A Simpson, S.C Wilks, T Ma, et al. Modeling laser-driven ion acceleration with deep learning. *Physics of Plasmas*, 28(4), 2021.
- [216] Y Sentoku, T.E Cowan, A Kemp, and H Ruhl. High energy proton acceleration in interaction of short laser pulse with dense plasma target. *Physics of Plasmas*, 10(5):2009–2015, 2003.
- [217] M Roth, A Blazevic, M Geissel, T Schlegel, T.E Cowan, M Allen, et al. Energetic ions generated by laser pulses: A detailed study on target properties. *Physical Review Special Topics-Accelerators and Beams*, 5(6):061301, 2002.
- [218] M Allen, P.K Patel, A Mackinnon, D Price, S Wilks, and E Morse. Direct experimental evidence of back-surface ion acceleration from laser-irradiated gold foils. *Physical Review Letters*, 93(26):265004, 2004.
- [219] S.M Pfothenhauer, O Jäckel, and A Sachtleben. Spectral shaping of laser generated proton beams. *New Journal of Physics*, 10(3):033034, 2008.
- [220] I.J Kim, K.H Pae, I.W Choi, C.L Lee, H.T Kim, H Singhal, et al. Radiation pressure acceleration of protons to 93 mev with circularly polarized petawatt laser pulses. *Physics of Plasmas*, 23(7), 2016.

- [221] A Macchi, S Veghini, and F Pegoraro. “light sail” acceleration reexamined. *Physical Review Letters*, 103(8):085003, 2009.
- [222] T.V Liseykina, M Borghesi, A Macchi, and S Tuveri. Radiation pressure acceleration by ultraintense laser pulses. *Plasma Physics and Controlled Fusion*, 50(12):124033, 2008.
- [223] A Macchi, S Veghini, T.V Liseykina, and F Pegoraro. Radiation pressure acceleration of ultrathin foils. *New Journal of Physics*, 12(4):045013, 2010.
- [224] S Steinke, P Hinz, M Schnürer, G Priebe, J Bränzel, F Abicht, et al. Stable laser-ion acceleration in the light sail regime. *Physical Review Special Topics-Accelerators and Beams*, 16(1):011303, 2013.
- [225] O Klimo, J Psikal, J Limpouch, and V.T Tikhonchuk. Monoenergetic ion beams from ultrathin foils irradiated by ultrahigh-contrast circularly polarized laser pulses. *Physical Review Special Topics-Accelerators and Beams*, 11(3):031301, 2008.
- [226] S Palaniyappan, B.M Hegelich, H.C Wu, D Jung, D.C Gautier, L Yin, et al. Dynamics of relativistic transparency and optical shuttering in expanding overdense plasmas. *Nature Physics*, 8(10):763–769, 2012.
- [227] B Qiao, S Kar, M Geissler, P Gibbon, M Zepf, and M Borghesi. Dominance of radiation pressure in ion acceleration with linearly polarized pulses at intensities of  $10^{21}$  W cm<sup>-2</sup>. *Physical Review Letters*, 108(11):115002, 2012.
- [228] X.Q Yan, T Tajima, M Hegelich, L Yin, and D Habs. Theory of laser ion acceleration from a foil target of nanometer thickness. *Applied Physics B*, 98:711–721, 2010.
- [229] I.N Ross, P Matousek, M Towrie, A.J Langley, and J.L Collier. The prospects for ultrashort pulse duration and ultrahigh intensity using optical parametric chirped pulse amplifiers. *Optics Communications*, 144(1-3):125–133, 1997.

- [230] J.L Collier, C Hernandez-Gomez, I.N Ross, P Matousek, C.N Danson, and J Walczak. Evaluation of an ultrabroadband high-gain amplification technique for chirped pulse amplification facilities. *Applied Optics*, 38(36):7486–7493, 1999.
- [231] C.N Danson, P.A Brummitt, R.J Clarke, J.L Collier, B Fell, A.J Frackiewicz, et al. Vulcan petawatt: Design, operation and interactions at  $5 \times 10^{20}$  wcm<sup>-2</sup>. *Laser and Particle Beams*, 23(1):87–93, 2005.
- [232] R. M Wood. *Laser-induced damage of optical materials*. CRC Press, 2003.
- [233] R Adair, L.L Chase, and S.A Payne. Nonlinear refractive index of optical crystals. *Physical Review B*, 39(5):3337, 1989.
- [234] Y Senatsky, A Shirakawa, Y Sato, J Hagiwara, J Lu, K Ueda, et al. Nonlinear refractive index of ceramic laser media and perspectives of their usage in a high-power laser-driver. *Laser Physics Letters*, 1(10):500, 2004.
- [235] J Kerr. XI. a new relation between electricity and light: Dielectrified media birefringent. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(332):337–348, 1875.
- [236] M.D Perry, T Ditmire, and B.C Stuart. Self-phase modulation in chirped-pulse amplification. *Optics Letters*, 19(24):2149–2151, 1994.
- [237] M Raghuramaiah, A.K Sharma, P.A Naik, P.D Gupta, and R.A Ganeev. A second-order autocorrelator for single-shot measurement of femtosecond laser pulse durations. *Sadhana*, 26:603–611, 2001.
- [238] C.J Hooker, J.L Collier, O Chekhlov, R.J Clarke, E Divall, K Ertel, et al. The astra gemini project—a dual-beam petawatt ti: Sapphire laser system. In *Journal de Physique IV (Proceedings)*, volume 133, pages 673–677. EDP sciences, 2006.
- [239] B.C Platt and R Shack. History and principles of shack-hartmann wavefront sensing, 2001.

- [240] S.D Williamson, R Wilson, M King, M Duff, B Gonzalez-Izquierdo, Z.E Davidson, et al. Self-referencing spectral interferometric probing of the onset time of relativistic transparency in intense laser-foil interactions. *Physical Review Applied*, 14(3):034018, 2020.
- [241] L Hanssen. Integrating-sphere system and method for absolute measurement of transmittance, reflectance, and absorptance of specular samples. *Applied Optics*, 40(19):3196–3204, 2001.
- [242] R Benattar, C Popovics, and R Sigel. Polarized light interferometer for laser fusion studies. *Review of Scientific Instruments*, 50(12):1583–1586, 1979.
- [243] O Lundh. Laser-driven beams of fast ions, relativistic electrons and coherent x-ray photons. PhD thesis, Lund University, 2008.
- [244] Z.E Davidson, B Gonzalez-Izquierdo, A Higginson, K.L Lancaster, S.D Williamson, M King, et al. An optically multiplexed single-shot time-resolved probe of laser–plasma dynamics. *Optics Express*, 27(4):4416–4423, 2019.
- [245] D.C Carroll, P Brummitt, D Neely, F Lindau, O Lundh, C.G Wahlström, et al. A modified thomson parabola spectrometer for high resolution multi-mev ion measurements—application to laser-driven ion acceleration. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 620(1):23–27, 2010.
- [246] W.H Bragg and R Kleeman. On the  $\alpha$  particles of radium, and their loss of range in passing through various atoms and molecules. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 10(57):318–340, 1905.
- [247] A Brown and H Suit. The centenary of the discovery of the bragg peak. *Radiotherapy and Oncology*, 73(3):265–268, 2004.

- [248] M Borghesi, D.H Campbell, A Schiavi, M.G Haines, O Willi, A.J MacKinnon, et al. Electric field detection in laser-plasma interaction experiments via the proton imaging technique. *Physics of Plasmas*, 9(5):2214–2220, 2002.
- [249] C.D Armstrong, D Neely, D Kumar, P McKenna, R.J Gray, and A.S Pirozhkov. Deconvolution of multi-boltzmann x-ray distribution from linear absorption spectrometer via analytical parameter reduction. *Review of Scientific Instruments*, 92(11), 2021.
- [250] H Padda, M King, R.J Gray, H.W Powell, B Gonzalez-Izquierdo, L.C Stockhausen, et al. Intra-pulse transition between ion acceleration mechanisms in intense laser-foil interactions. *Physics of Plasmas*, 23(6), 2016.
- [251] B Fryxell, K Olson, P Ricker, F.X Timmes, M Zingale, D.Q Lamb, et al. Flash: An adaptive mesh hydrodynamics code for modeling astrophysical thermonuclear flashes. *The Astrophysical Journal Supplement Series*, 131(1):273, 2000.
- [252] J Allison, K Amako, J Apostolakis, P Arce, M Asai, T Aso, et al. Recent developments in geant4. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, 2016.
- [253] A.P Robinson and M Sherlock. Magnetic collimation of fast electrons produced by ultraintense laser irradiation by structuring the target composition. *Physics of Plasmas*, 14(8), 2007.
- [254] R.A Fonseca, L.O Silva, F.S Tsung, V.K Decyk, W Lu, C Ren, et al. Osiris: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators. In *Computational Science—ICCS 2002: International Conference Amsterdam, The Netherlands, April 21–24, 2002 Proceedings, Part III 2*, pages 342–351. Springer, 2002.

- [255] R Lehe, M Kirchen, I.A Andriyash, B.B Godfrey, and J.L Vay. A spectral, quasi-cylindrical and dispersion-free particle-in-cell algorithm. *Computer Physics Communications*, 203:66–82, 2016.
- [256] Y.L Klimontovich. Relativistic transport equations for a plasma 1. *Soviet Physics JETP*, 10(3):524–530, 1960.
- [257] C.K Birdsall and A.B Langdon. *Plasma physics via computer simulation*. CRC press, 2018.
- [258] K Yee. Numerical solution of initial boundary value problems involving maxwell’s equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, 1966.
- [259] J.P Boris. Relativistic plasma simulation-optimization of a hybrid code. In *Proc. Fourth Conf. Num. Sim. Plasmas*, pages 3–67, 1970.
- [260] S Jalas, M Kirchen, P Messner, P Winkler, L Hübner, J Dirkwinkel, et al. Bayesian optimization of a laser-plasma accelerator. *Physical Review Letters*, 126(10):104801, 2021.
- [261] S.J Dann, C.D Baird, N Bourgeois, O Chekhlov, S Eardley, C.D Gregory, et al. Laser wakefield acceleration with active feedback at 5 hz. *Physical Review Accelerators and Beams*, 22(4):041303, 2019.
- [262] F Irshad, S Karsch, and A Döpp. Multi-objective and multi-fidelity bayesian optimization of laser-plasma acceleration. *Physical Review Research*, 5(1):013063, 2023.
- [263] J Goodman, M King, E.J Dolier, R Wilson, R.J Gray, and P McKenna. Optimization and control of synchrotron emission in ultraintense laser–solid interactions using machine learning. *High Power Laser Science and Engineering*, 11:e34, 2023.
- [264] C.E Rasmussen and C.K Williams. *Gaussian processes for machine learning*, volume 1. Springer, 2006.

- [265] Cambridge service for data-driven discovery (csd3) high-performance computing cluster [online, accessed 1 march 2024]. available from: <https://www.hpc.cam.ac.uk/high-performance-computing>.
- [266] F Pedregosa, G Varoquaux, A Gramfort, V Michel, B Thirion, O Grisel, et al. Scikit-learn: Machine learning in python. *The Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [267] A Forrester, A Sobester, and A Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley and Sons, 2008.
- [268] J Nocedal and S.J Wright. *Numerical optimization*. Springer, 1999.
- [269] M.J Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal*, 7(2):155–162, 1964.
- [270] M.J Powell. A method for minimizing a sum of squares of non-linear functions without calculating derivatives. *The Computer Journal*, 7(4):303–307, 1965.
- [271] J.A Nelder and R Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965.
- [272] A L’Heureux, K Grolinger, H.F Elyamany, and M.A Capretz. Machine learning with big data: Challenges and approaches. *IEEE Access*, 5:7776–7797, 2017.
- [273] V Picheny, T Wagner, and D Ginsbourger. A benchmark of kriging-based infill criteria for noisy optimization. *Structural and Multidisciplinary Optimization*, 48:607–626, 2013.
- [274] <http://www.sfu.ca/ssurjano/hart6.html>.
- [275] C.T Chu, S Kim, Y.A Lin, Y.Y Yu, G Bradski, K Olukotun, et al. Map-reduce for machine learning on multicore. *Advances in Neural Information Processing Systems*, 19, 2006.

- [276] G Hughes. On the mean accuracy of statistical pattern recognizers. *IEEE Transactions on Information Theory*, 14(1):55–63, 1968.
- [277] J.M Bernardo and A.F Smith. *Bayesian theory*, volume 405. John Wiley and Sons, 2009.
- [278] <https://towardsdatascience.com/bayes-rule-applied-75965e4482ff>.
- [279] D.K Duvenaud. Automatic model construction with Gaussian processes. PhD thesis, University of Cambridge, 2014.
- [280] B Shahriari, K Swersky, Z Wang, R.P Adams, and N De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2015.
- [281] M Hoffman, E Brochu, and N De Freitas. Portfolio allocation for bayesian optimization. In *UAI*, pages 327–336, 2011.
- [282] T Kluge, T.E Cowan, A Debus, U Schramm, K Zeil, and M Bussmann. Electron temperature scaling in laser interaction with solids. *Physical Review Letters*, 107(20):205003, 2011.
- [283] A Sgattoni, P Londrillo, A Macchi, and M Passoni. Laser ion acceleration using a solid target coupled with a low-density layer. *Physical Review E*, 85(3):036405, 2012.
- [284] E d’Humières et al. Optimization of laser-target interaction for proton acceleration. *Physics of Plasmas*, 20(2), 2013.
- [285] D.J Stark, L Yin, B.J Albright, and F Guo. Effects of dimensionality on kinetic simulations of laser-ion acceleration in the transparency regime. *Physics of Plasmas*, 24(5), 2017.
- [286] M Coury, D.C Carroll, A.P Robinson, X.H Yuan, C.M Brenner, M Burza, et al. Influence of laser irradiated spot size on energetic electron injection and proton acceleration in foil targets. *Applied Physics Letters*, 100(7), 2012.

- [287] A Henig, D Kiefer, K Markey, D.C Gautier, K.A Flippo, S Letzring, et al. Enhanced laser-driven ion acceleration in the relativistic transparency regime. *Physical Review Letters*, 103(4):045002, 2009.
- [288] C.A Palmer, N.P Dover, I Pogorelsky, M Babzien, G.I Dudnikova, M Ispiryan, et al. Monoenergetic proton beams accelerated by a radiation pressure driven shock. *Physical Review Letters*, 106(1):014801, 2011.
- [289] A McIlvenny, H Ahmed, C Scullion, D Doria, L Romagnani, P Martin, et al. Characteristics of ion beams generated in the interaction of ultra-short laser pulses with ultra-thin foils. *Plasma Physics and Controlled Fusion*, 62(5):054001, 2020.
- [290] L.O Silva, M Marti, J.R Davies, R.A Fonseca, C Ren, F.S Tsung, et al. Proton shock acceleration in laser-plasma interactions. *Physical Review Letters*, 92(1):015002, 2004.
- [291] M Borghesi, A Schiavi, D.H Campbell, M.G Haines, O Willi, A.J Mackinnon, et al. Proton imaging detection of transient electromagnetic fields in laser-plasma interactions. *Review of Scientific Instruments*, 74(3):1688–1693, 2003.
- [292] C.B Curry, C.A Dunning, M Gauthier, H.G Chou, F Fiuza, G.D Glenn, et al. Optimization of radiochromic film stacks to diagnose high-flux laser-accelerated proton beams. *Review of Scientific Instruments*, 91(9), 2020.
- [293] Archer2 high-performance computing cluster [online, accessed 1 march 2024]. available from: <https://www.archer2.ac.uk/>.